When (not) to Publicize Inspection Results*

Eilon Solan[†] and Chang Zhao[‡]

April 27, 2023

Abstract

We consider a dynamic inspection problem between a principal and several agents. The principal observes the full inspection history, whereas each agent only observes inspections imposed on himself. When inspection resources are limited, the inspection intensities for agents are negatively correlated, and hence each agent cares not only about his own inspection history, but also about the inspection histories of the other agents. In such cases, should the principal publicly reveal past inspection history, or should she conceal this information? We show that the principal benefits from concealing inspection history. Nevertheless, this benefit becomes less significant as the number of agents increases, and disappears in the limit case with a continuum of agents.

JEL classification: C73; D82; D83

Declarations of interest: none.

1 Introduction

A utility-maximizing agent (he) can often increase his utility by breaking some rules: given the opportunity, firms will violate regulations that increase their costs, workers will shirk on their tasks, and taxpayers will under-report their income. To reduce the rate of violations, it is customary for the principal (she) to use inspections as a means to discipline agents. The principal's inspection scheme determines which agents are to be inspected in each period, as well as the information on past inspections that will be revealed to the agents.

*The research of the authors was supported by ISF grants #217/17 and #2510/17, National Natural Science Foundation of China #72273066, and Global Integrity. This document is an output from a project funded by UK Aid from the UK Government for the benefits of developing countries. The views expressed are not necessarily those of UK government's official policies.

[†]School of Mathematical Sciences, Tel-Aviv University, Israel. Email address: eilons@post.tau.ac.il

[‡]School of International Trade and Economics, University of International Business and Economics, China. Email address: zhaochangtd@hotmail.com

There are two channels to affect agents' beliefs on past inspections. The first operates by concealing an agent's inspection result from himself: that is, an agent who is inspected is not informed of his own inspection outcome. This channel has been well studied in the literature,[1] and Fuchs (2007) shows that the principal is best off keeping the agent uninformed of the intermediate outcome realizations until the last period, since this allows for the reusability of punishments. Nevertheless, concealing an agent's own inspection result from himself is possible *only* when the performance measure is subjective. Scenarios like environmental inspection or tax auditing do not fit into this category. In this paper, we focus on contexts where the performance evaluation is objective, and each agent knows his own inspection history. In these cases, the principal has another channel to affect agents' beliefs on past inspections, namely, by strategically revealing the inspection history of the other agents.

At a first glance, if each agent's payoff depends only on the actions he takes and the inspection intensity he faces, the inspection outcomes for other agents are irrelevant to the agent. This is not true when inspection resources are limited and full compliance is not implementable. Indeed, in such cases, the inspection intensities for agents are negatively correlated — from the perspective of agent i, a lower inspection intensity for other agents implies that more inspection resources are left for agent i, and hence he is less willing to violate. When the principal uses an inspection scheme that is history dependent, the strategic revelation of past inspection histories constitutes an important policy instrument.

In reality, more often than not, no information about inspection outcomes is publicly available. Yet the past few decades have witnessed a considerable increase in the amount of public information provided by inspection authorities. Back in the 1980s, the US started adopting the "public disclosure programs" that mandate public reporting of firms' compliance with environmental regulations. Since then, similar programs have been developed in many countries.[2]

Despite the growing popularity of such public disclosure programs, their benefits are open to debate, and no universal and firm conclusions can be drawn from the empirical ground. Some empirical works support the view that disclosure has a positive impact, but there are also some doubts about this conclusion, due to lack of data prior to the introduction of the disclosure program,[3] the implementation of other programs concurrently with the disclosure

[1]Levin (2003) and MacLeod (2003) were the first to analyze this environment.

[2]For instance, Canada's National Pollutant Release Inventory (1992), Indonesia's PROPER program (1995), Philippines' EcoWatch (1997), Australia's National Pollutant Inventory (2000), Europe's Pollutant Emission Register (2000), and China's GreenWatch (2000).

[3]The US Environmental Protection Agency reported an overall 43% decrease of national release of toxics under the US Toxics Release Inventory (TRI). However, since trend data prior to TRI's introduction are not available, one cannot infer that TRI is responsible for the entire reduction in emissions.

program,[4] and the non-universality of the positive results[5] (Folmer and Tietenberg, 2007).

From a theoretical perspective, there is a large literature devoted to the effect of publicizing inspection results. Yet, most of these works have abstracted away from the direct interaction between inspectors and agents. This literature shows that with public disclosure, firms face additional pressure for compliance from third parties, e.g., neighbouring communities, consumers, investors, or stockholders. This helps to discipline the behavior of firms and yields a positive result (see, e.g., Hamilton, 1995, Konar and Cohen, 1997, Lanoie, Laplante, and Roy, 1997, Tietenberg and Wheeler, 2001, Foulon, Lanoie, and Laplante, 2002, Blackman, Afsah, and Ratunanda, 2004, Stephan, 2002).

We contribute to this literature by focusing on the interaction between inspectors and agents, and showing that contrary to previous positive results, when the number of agents is small, public disclosure of inspection histories may hurt the principal. This being said, the advantage of concealing inspection histories becomes less significant as the number of agents increases. In the limiting game with a continuum of agents, the advantage of concealing information disappears. We mainly deal with the model where the principal has a commitment power. The no-commitment case is discussed in Section 3.2.

Formally, we compare two realistic monitoring structures: public monitoring and private monitoring.[6] Under *public monitoring*, the principal announces her observations after each period (that is, the identity of the inspected agents and their actions). Under *private monitoring*, the principal conceals her observations (that is, each agent only observes his own inspection history).

To illustrate the superiority of private monitoring, we start with an important feature of the optimal inspection scheme under public monitoring. Under public monitoring, it is optimal for the principal to adopt an inspection scheme that favors agents who successfully passed inspections (i.e., were found adhering) in early periods. Indeed, to improve upon the static scenario, a successfully passed inspection has to be accompanied by a future bonus to the agent: that is, the promise of the principal turning a blind eye with positive probability in the future, so that the agent can violate without penalty. The feature that greater compliance typically leads to less enforcement is well recognized in both theory and practice.[7] When

[4]In China, together with the introduction of the GreenWatch, conventional regulations were more strongly established, and other programs for emission reduction were implemented concurrently.

[5]Canada's National Pollutant Release Inventory (NPRI) was established in 1992. In an analysis of trends since 1997, the 2000 NPRI found that the total releases of 17 toxic substances were up slightly (4.5%).

[6]The problem of optimal information revelation is discussed in Section 4.3.

[7]Gray and Deily (1996) analyze the air pollution data in the US steel industry and find that less regulatory attention is imposed on the steel plants with greater compliance. Dubin and Wilde (1988) study the IRS data set and observe a similar pattern: a higher compliance level induces lower audit rates.

the principal plans her activities according to a pre-allocated budget, less enforcement on certain agents implies greater enforcement on other agents. This implies that agents who are inspected *less* often in early periods are subject to *greater* enforcement later on, and they violate less in the long run.

Under private monitoring, using a similar inspection scheme, but adding a small probability to no inspection in early periods, the principal manipulates agents' beliefs to her advantage: this way, with a positive probability each agent assigns a high probability to the event that the other agents were inspected in early periods and hence he himself will face high inspection intensities in the future. Since all agents follow the same reasoning, in such eventuality all agents violate less in the long run.

From a technical perspective, the advantage of private monitoring over public monitoring is that it allows the principal to transfer inspection resources across histories, and attain a more efficient use of it. We illustrate this important point by an example. Suppose that agent i adheres if and only if the probability he is inspected is at least $\frac{3}{4}$, and suppose that there are two histories h and h', which differ only in the actions of some other agent j. Suppose further that at h (resp., h') agent i is inspected with probability 1 (resp., 0). Under public monitoring, agent i adheres at h and violates at h'. Under private monitoring, agent i cannot distinguish between h and h'. If these two histories constitute an information set for agent i, and if, given that the information set is reached, agent i assigns probability at least $\frac{3}{4}$ to h, then at this information set agent i adheres. In particular, he adheres at h'.

Thus, the potential superiority of hiding information (private monitoring) stems from pooling some histories together, so that an agent's incentive constraint has to be satisfied only in expectation rather than state by state. This pooling is advantageous only if at some histories the inspection resource is inevitably superfluous (hence agents' incentive constrains are slack), and the pooling allows a more efficient use of this redundant resource as it can average out some other histories where the inspection resource is scarce.

The above effect is more significant when the number of agents is small. As the number of agents increases, under public monitoring, the ability to finely divide agents into groups (with agents in the same group being treated similarly) allows the principal to attain a more efficient use of the inspection resource at any given history, and the gain from pooling histories decreases. In the limit case with a continuum of agents, the perfect divisibility of agents allows public monitoring to attain the most efficient use of the inspection resource, so that agents' incentive constraints are binding at any history, and there is no gain from pooling histories.

The key point of our paper is that sometimes the principal wants to pool information to keep agents in the dark, so as to attain a better use of her resources.[8] This manipulation works under various conditions, which include, among others, a negative correlation between the resources available for inspecting each agent,[9] a commitment power of the principal,[10] and the inability of agents to communicate and figure out the past actions of the principal.[11] This result has important policy implications. As mentioned before, a large literature argues that public monitoring has an important advantage over private monitoring, as it makes firms face additional pressure for compliance from third parties. We show that the call for publicly disclosing more information may miss out some of the tradeoffs that arise under limited inspection resource. Especially when having a few agents, the flexibility in manipulating agents' beliefs offered by private monitoring may outweigh the said advantage of public monitoring. Yet as the number of agents increases, the advantage of private monitoring disappears, and public monitoring becomes unequivocally preferable.

Related literature.

Our paper is related to information design problems where the designer can send private signals to agents. In *static settings*, Bergemann and Morris (2016) and Taneva (2019) relate the optimal information disclosure to the best Bayes correlated equilibrium from the sender's perspective, and propose mechanisms to compute the optimal utility of the sender. In particular, Bergemann and Morris (2016) show that more information reduces the set of outcomes by imposing more incentive constraints. Arieli and Babichenko (2019) study the private information disclosure in a specific Bayesian persuasion model of product adoption, and analyze how information should be revealed optimally by the firm in order to maximize its revenue as a function of its utility and the utilities of the consumer.

Even though in the static model there is a clear pattern that more information reduces the set of outcomes and hence hurts the designer, this is not necessarily the case in repeated games. In dynamic settings, Matsushima (1991) and Bhaskar and Van Damme (2002) provide examples of repeated games where public monitoring is superior to private monitoring. In these examples, public disclosure is beneficial because it implies a larger set of strategies for the players, since they can condition their future actions on more information.

In contrast, Phelan and Skrzypacz (2012) provide dynamic examples in which less information benefits incentives. In these examples, players do not observe each other's action

[8] A similar idea also appears in Crawford (2003), where *sophisticated* players can mimic the action of *mortal* players (who always play a certain fixed action) so as to confuse the opponent.

[9] This condition is satisfied if the inspection resource is limited, or if the marginal inspection cost is increasing (see Section 3.3).

[10] In Section 3.2 we discuss the no-commitment case.

[11] In Remark 1 we discuss agents' (in)ability to communicate.

directly. Instead, they observe a noisy signal of it. In such cases, a less accurate signal serves to make the threat of punishment credible: if a player knows too well the action of the opponent, then even if he observes unexpected signals, the incentive to punish the opponent can be insufficient. More noise on the signal helps in generating more uncertainty on the opponent's past actions, which makes it easier to satisfy incentive constraints.

In our model, each agent interacts mainly with the principal, and his incentive to punish a deviation of another agent is less relevant. The change on the monitoring structure thus affects agents' behavior in a way different from Phelan and Skrzypacz (2012): in the two-period game, under private monitoring, the equilibrium play in the second period can be a correlated equilibrium of the stage game (with private histories as the correlation device), while under public monitoring, the equilibrium play in the second period is necessarily a Nash equilibrium of the stage game.

The fact that private histories can constitute an appropriate correlation device for the second-period play to support the efficient correlated equilibrium has long been observed in the private strategy literature (see, e.g., Lehrer, 1991, and Mailath, Matthews, and Sekiguchi, 2002). Lehrer (1991) studies two-player infinite undiscounted repeated game with a specific public monitoring structure; Mailath, Matthews, and Sekiguchi (2002) provide an example with two agents and two periods. We study finitely-repeated inspection problems that involve multiple agents. In addition to the above driving force, there exist others that favor public monitoring (see Remark 2 and Theorem 5). We identify conditions under which private (or public) monitoring dominates.

Andrews and Barron (2016) study a dynamic relational contract problem between a principal and several agents, and show that concealing information may be beneficial to the principal. In their paper, monetary transfers between the principal and each agent are allowed. If the principal can commit to a strategy, then the first-best outcome can be implemented with a simple stationary contract. Andrews and Barron (2016) therefore focus on no-commitment cases (the outputs produced by agents are non-contractable) and construct a dynamic allocation rule — the Favored Supplier Allocation — that attains the first-best outcome whenever any allocation rule does. In such cases, information does not play a crucial role, and the principal does not benefit from disclosing/concealing information on past play. They show that in the two-agent case, when the first-best outcome cannot be attained, full disclosure may not be optimal and concealing some information may yield a superior outcome to the principal.

Similarly to ours, this last result is related to the literature on correlated equilibrium: The private histories can constitute an appropriate correlation device for future play. A key difference between our paper and Andrews and Barron (2016) is that we do not allow monetary transfers between players. Therefore, even in the commitment case, concealing information can be superior. Also, we study cases that involve more than two agents. In particular, in the no-commitment case, we show that when the number of agents is large, in sharp contrast to the two-agent problem, *public* monitoring can be strictly superior to private monitoring.

Kandori and Obara (2006) study the difference between public and private strategies, and show that the latter may be superior. They consider a repeated prisoner's dilemma game with public signals on the stage outcomes. In their model, the signal is rather insensitive to a deviation at the cooperative point, but it is quite sensitive to one's deviation when the opponent is playing D. In such cases, private strategies (which depend not only on public signals, but also on players' own actions in the past) can achieve a more efficient outcome than public strategies (which depend solely on the history of publicly observable signals). This is because private strategies allow players to use additional information (i.e., their own past actions) to improve the efficiency of punishments. In contrast, in our model private monitoring implies that an agent has *less* information than in the model with public monitoring, since he can observe the actions of others only in the latter case.

Our paper is also related to repeated moral hazard problems that involve private evaluation. In those problems, there is an agent who takes an action, which generates a (noisy) signal to the principal. The principal does not observe the action of the agent, and the agent does not observe the signal to the principal. Levin (2003) and MacLeod (2003) were the first to analyze this environment, yet their analysis is essentially static. Fuchs (2007) extends the analysis to a dynamic environment, and analyzes the optimal timing to reveal the principal's private signal. A single-agent problem is studied and it is shown that the principal is best off keeping the agent uninformed of the intermediate outcome realizations until the last period.

In our paper, similarly to Fuchs (2007), concealing inspection histories makes the use of the reward more efficient. Nevertheless, there is a substantial difference: Unlike Fuchs (2007), we study multiple-agent problems, and the uncertainty no longer stems from the agent's own inspection history, but rather from *other* agents' inspection histories. Moreover, unlike Fuchs (2007), where monetary transfer is allowed, in our paper the reward can only take the form of violations and is bounded by 1 in each period. These two differences in the setup imply the following differences in the equilibrium analysis. First, in our model with multiple agents and no monetary transfer, it is no longer clear that concealing information is always better (we elaborate on this point in Remark 2 and Theorem 5). Second, for concealing information to be beneficial, a negative correlation between agents is required, and hence the inspection cost has to be convex in the number of inspections. Third, the effect of information now depends crucially on the number of agents.

The paper is organized as follows. Section 2 studies the benchmark model with one inspector and two agents, and Section 3 generalizes the benchmark result. Discussion appears in Section 4. Proofs are relegated to the Appendix.

2 Benchmark model

We start with a two-period game between a *principal* and two *agents*.[12] In every period each agent decides whether to Adhere or to Violate. The principal has limited inspection resources and she can costlessly inspect at most one agent in each period (i.e., the principal has one inspector at her disposal). The action of the inspected agent is perfectly observed by the principal.

The action set of agent i, $i = 1, 2$, is $\mathscr{A}_i := \{A, V\}$, where A stands for adhering and V for violating. The action set of the principal is $\mathscr{A}_0 := \{I_1, I_2, \emptyset\}$, where $\emptyset$ stands for no inspection, and I_1 (resp., I_2) stands for inspecting Agent 1 (resp., Agent 2). Throughout the paper, whenever a variable refers to the principal we add the subscript 0, and whenever it refers to Agent 1 or 2 we add the subscript 1 or 2, respectively.

2.1 The stage payoff

The gain of each agent from adhering is normalized to 0, and his gain from an undetected violation is normalized to 1. The agent's loss from a detected violation is denoted by c. The value of c is exogenously given, say, by legal constraints. Consequently, the stage payoff function of agent i, denoted u_i, is given by

$$u_i(a) = \begin{cases} 0, & \text{if } a_i = A, \\ -c, & \text{if } a_i = V \text{ and } a_0 = I_i, \\ 1, & \text{if } a_i = V \text{ and } a_0 \neq I_i, \end{cases} \tag{1}$$

where $a = (a_0, a_1, a_2)$ is the vector of actions played at the current stage. The principal loses 1 for each violation, detected or undetected. Therefore, the principal's loss function, denoted ℓ_0, is given by

$$\ell_0(a) = \begin{cases} 2, & \text{if } a_1 = a_2 = V, \\ 1, & \text{if } a_i = A \text{ and } a_j = V \text{ for } i \neq j, \\ 0, & \text{if } a_1 = a_2 = A. \end{cases} \tag{2}$$

[12] In Section 3.1 we extend our results to games that involve more agents. In Section 3.4 we discuss games with a larger number of periods.

We assume that the principal does not observe undetected violations, hence she does not observe her realized payoff. The form of the principal's loss function implies that we focus on the principal's incentive to deter violations, rather than on her incentive to collect penalties. Indeed, in many cases, it is difficult, if not impossible, to compare the damage caused by the violation behavior with the monetary penalties, as they are in different dimensions.[13] In Section 4.1 we discuss scenarios where the principal benefits from collecting penalties.

In the one-shot game, an agent adheres (resp., violates) as long as the probability he is inspected is larger (resp., smaller) than $\frac{1}{1+c}$. Therefore, when the fine for violation is sufficiently large, specifically, if $c \geq 1$, the principal can deter both agents from violating by inspecting each agent with probability $\frac{1}{2}$. Similarly, in the repeated game, when $c \geq 1$, the stationary strategy that inspects each agent with probability $\frac{1}{2}$ at every period guarantees that rational agents will always adhere. In the sequel we will focus on the nontrivial case $c < 1$.

Assumption 1. $0 < c < 1$.

When c is less than 1, in the one-shot game an agent adheres whenever the inspection probability is larger than $\frac{1}{1+c} > \frac{1}{2}$. The best the principal can do is to deter one agent from violating and allow the other agent to violate.

2.2 Monitoring structure and histories

We now turn to the two-period game. If the principal inspects an agent, the action of that agent is perfectly observed by the principal. Thus, the principal observes at the end of each period a private signal y_0, drawn from a signal space $Y_0 = \{V_1, A_1, V_2, A_2, \emptyset\}$, with the interpretation that the signal $\emptyset$ means no agent was inspected, and the signal A_i (resp., V_i) means that agent i was inspected and found adhering (resp., violating). As for the agents' observations, we consider two monitoring structures.

In one structure, the monitoring is *public*: The identity of the inspected agent (if any) as well as the outcome of the inspection are announced to all players. In other words, the signal received by the principal is publicly observed by the agents, and it constitutes the *public history*. The public monitoring game is denoted by G^{pub}.

In the other structure, the monitoring is *private*: Any uninspected agent does not know whether the principal inspected the other agent and, if the other agent was inspected, what was the outcome of the inspection. Thus, an agent only observes whether he himself is inspected or not. Formally, if agent i is inspected, he receives the same signal as the principal

[13] For instance, in the environmental control problems, the damage caused by a marine oil spill cannot be compared with a monetary fine.

(that is, either A_i or V_i). Otherwise, agent i receives the signal N_i, meaning that he was not inspected. The private monitoring game is denoted by G^{pri}. Figure 1 provides an illustration for the difference in signals (marked in red) under the two monitoring structures.

Action	Principal	I_1	I_1	I_2	I_2	$\emptyset$
	Agent 1	A	V	A	V	A
	Agent 2	A	A	A	V	A
Signal (**Public** monitor.)	Principal	A_1	V_1	A_2	V_2	$\emptyset$
	Agent 1	A_1	V_1	$\boldsymbol{A_2}$	$\boldsymbol{V_2}$	$\boldsymbol{\emptyset}$
	Agent 2	$\boldsymbol{A_1}$	$\boldsymbol{V_1}$	A_2	V_2	$\boldsymbol{\emptyset}$
Signal (**Private** monitor.)	Principal	A_1	V_1	A_2	V_2	$\emptyset$
	Agent 1	A_1	V_1	$\boldsymbol{N_1}$	$\boldsymbol{N_1}$	$\boldsymbol{N_1}$
	Agent 2	$\boldsymbol{N_2}$	$\boldsymbol{N_2}$	A_2	V_2	$\boldsymbol{N_2}$

Figure 1: Examples for signals under different monitoring structures.

We further assume that the players have at their disposal a public correlation device, which outputs at the beginning of every period t a random signal ζ^t that is uniformly distributed on the interval $[0, 1]$ and independent of past play and past random signal. The public correlation device is crucial in some parts of the proofs. On the equilibrium path, however, the use of the correlation device is minimal.[14]

For player i, the information available in period t, denoted h_i^t, is the $(t-1)$-period history of the outcome of the correlation device, his private signals and past actions. The set of finite-length private histories is denoted by $\mathscr{H}_i$.

2.3 Strategies and payoffs

A (behavior) strategy of the principal is a function from the set $\mathscr{H}_0$ of her private histories to the set $\Delta\mathscr{A}_0$ of her mixed actions,

$$\sigma_0 : \mathscr{H}_0 \to \Delta\mathscr{A}_0.$$

Denote by $\mathscr{B}_0$ the set of strategies of the principal in the repeated game.

It is assumed that the principal publicly announces her entire inspection strategy at the beginning of the game and she is able to commit to it (we discuss the no-commitment case in Section 3.2). Since the principal's entire strategy is known to the agents from the outset

[14]In the main model where the principal has a commitment power, the correlation device is not used in the two-period game, and it is used in at most one period in T-period games with $T \geq 3$.

of the game, a strategy for an agent assigns a mixed action to every strategy of the principal and every private history of the agent. Formally, for $i = 1, 2$, a strategy of agent i is a function

$$\sigma_i : \mathscr{B}_0 \times \mathscr{H}_i \to \Delta\mathscr{A}_i,$$

where $\Delta\mathscr{A}_i$ is the set of mixed actions of agent i. Note that under public monitoring, it is without loss of generality to assume that agents use only public histories and ignore their private information (that is, their uninspected past actions).[15]

Each player's payoff is the discounted sum of his stage payoffs, where δ is the common discount factor. Formally, in period t, the action profile a^t yields the payoff $u_i(a^t)$ to agent i, and the loss $\ell_0(a^t)$ to the principal. A play $\mathbf{a} = (a^1, a^2)$ thus determines two stage payoffs/losses to each player. Denote agent i's discounted payoff under strategy profile σ by

$$v_i(\sigma) := \mathbf{E}_\sigma \left[\sum_{t=1}^{2} \delta^{t-1} u_i(a^t) \right], \tag{3}$$

and the principal's discounted loss under the strategy profile σ by

$$L_0(\sigma) := \mathbf{E}_\sigma \left[\sum_{t=1}^{2} \delta^{t-1} \ell_0(a^t) \right]. \tag{4}$$

Recall that agents observe their realized stage-game payoff, while the principal does not. This is because the principal knows only the action of the inspected agent, so she is unaware of the damage inflicted on her by the uninspected agent.

2.4 Equilibrium concept

Denote the two-period game following the announcement of an inspection strategy σ_0 by $\Gamma^{\text{pub}}(\sigma_0)$ in the public monitoring game and by $\Gamma^{\text{pri}}(\sigma_0)$ in the private monitoring game.

Under public monitoring, all participants observe the same public signal, and we study *perfect public equilibria* (PPEs) of the game $\Gamma^{\text{pub}}(\sigma_0)$. That is, deviations of the two agents are non-profitable after every public history. Formally, given a player i, for each history h^t and each strategy σ_i, player i's *continuation strategy* given history h^t, denoted $\sigma_i|_{h^t}$, is defined by $\sigma_i|_{h^t}(h^\tau) := \sigma_i(h^t h^\tau)$, for every $h^\tau \in \mathscr{H}$. The pair of strategies (σ_1, σ_2) is a PPE of the subgame $\Gamma(\sigma_0)$ if for every public history $h^t \in \mathscr{H}$, the pair of strategies $(\sigma_1|_{h^t}, \sigma_2|_{h^t})$ is a Nash equilibrium of $\Gamma(\sigma_0|_{h^t})$.

[15] Conditional on the public history, each agent's private information is independent of the private information of the other agent. Therefore, for any equilibrium strategy in which agents use their private information, we can find a public-strategy equilibrium which yields the same outcome.

Under private monitoring, each participant observes his or her own private signal, and we study *sequential equilibria* (SEs) of the game $\Gamma^{\text{pri}}(\sigma_0)$. That is, deviations of the agents are non-profitable after every private history. Formally, a *system of beliefs* is a function μ that assigns to each information set w of the game tree a probability distribution on the set of histories that lead to w. Given a strategy profile σ, a system of beliefs μ is *consistent* if there is a sequence of completely mixed strategy profiles $(\sigma_k)_{k\in\mathbb{N}}$, converging pointwise to σ, such that the associated conditional beliefs $(\mu_k)_{k\in\mathbb{N}}$ converge pointwise to the system of beliefs μ. A strategy profile and a system of beliefs form a sequential equilibrium (SE) if each player, after every private history of his, plays a best response given beliefs that are consistent with his private history.

By Assumption 1, $0 < c < 1$, and under both public and private monitoring, no inspection mechanism can ensure a violation-free behavior of the agents. Indeed, an agent who is inspected with probability less than $\frac{1}{1+c}$ in the first period can always obtain a positive payoff by violating in the first period and adhering in all subsequent periods. In fact, if both agents react myopically in the first period and adhere in all subsequent periods, they guarantee themselves[16] a total payoff of $1 - c$. Since the principal's loss is equal to or larger than the agents' gain,[17] a loss of $1 - c$ for the principal is inevitable.

Fact 1. When there are two agents and one inspector, the equilibrium loss of the principal is at least $1 - c$, regardless of the monitoring structure and the length of the game.

2.5 Equilibrium analysis

Denote by L^{pub} the lowest loss of the principal among all PPEs under public monitoring, and by L^{pri} the lowest loss of the principal among all SEs under private monitoring. If $L^{\text{pri}} < L^{\text{pub}}$ (resp., $L^{\text{pri}} > L^{\text{pub}}$), the principal's loss under private monitoring is lower (resp., higher) than her loss under public monitoring. We write Private $\succ$ Public (resp., Private $\prec$ Public) for this case, and say that private monitoring is superior (resp., inferior) to public monitoring for the principal. The notation Private $\approx$ Public corresponds to the case $L^{\text{pri}} = L^{\text{pub}}$. The next theorem provides conditions under which Private $\succ$ Public.

Theorem 1. *Suppose there are two agents and one inspector. If $\delta > 1 - c$, then Private $\succ$ Public in the two-period game.*

[16] By this strategy agent i guarantees himself a payoff $\max\{0, 1 - (1 + c)p_i\}$, where p_i is the inspection probability for agent i. Since $p_1 + p_2 \leq 1$, regardless of the choice of p_1 and p_2, the two agents guarantee themselves a total payoff of $1 - c$.

[17] The principal's loss exceeds agent's gain if and only if the agent is inspected while he violates; in such cases the inspection hurts the agent without benefiting the principal.

Proof. See Appendix A.1. □

The condition $\delta > 1 - c$ is a simplifying assumption, which guarantees that the principal can induce both agents to adhere in the first period (under both monitoring structures).[18] The rest of this section is devoted to providing an intuitive explanation to Theorem 1. We first discuss the structure of the optimal PPE under public monitoring. We then explicitly solve an example and show that there exists an SE that is superior to the optimal PPE.

Preliminary: Properties of the optimal PPE. First, to identify the optimal PPE, without loss of generality we can restrict attention to equilibria that satisfy the following conditions:[19]

(i) Whenever an agent is indifferent between adhering and violating, he adheres.
(ii) On the equilibrium path, whenever an agent violates, he is inspected with probability 0.
(iii) An agent who is inspected and found violating in the first period is punished in the most severe way: he is inspected with probability 1 in the second period.

Intuitively, Part (i) holds because the correlation device can be used to mimic lotteries performed by the agents. To illustrate part (ii), recall that the principal cannot completely deter violations. When an agent is supposed to violate in equilibrium, a positive probability of inspection hurts the agent without benefiting the principal. If this happens in the second period, it weakens the agent's incentive to adhere in the first period and hurts the principal. Consequently, in equilibrium the principal never detects violations, and any detected violation implies that the inspected agent deviated. This further implies that, without loss of generality one can assume that an agent who is inspected and found violating is punished in the most severe way, which is the content of Part (iii). In the rest of this paper we focus on PPEs that satisfy these conditions.

In period 2, an agent adheres if and only if the probability of inspection is at least $\frac{1}{1+c}$. In period 1, however, agent i may adhere when the inspection probability p_i is lower than $\frac{1}{1+c}$, provided that he is compensated in the second period (by being allowed to violate) if he passes the inspection in period 1. The lower p_i, the higher the reward required in period 2.

[18]If $\delta < 1 - c$, then using the second-period violation as a reward, the principal cannot deter both agents from violating in the first period. This is true under both monitoring structures. Even in this case, it is still possible that Private $\succ$ Public, since private monitoring may lead to higher compliance level in the second period.

[19]Results in this subsection are reminiscent of results in Solan and Zhao (2021), and their proofs are omitted (Solan and Zhao, 2021, study games with infinitely many periods, but their proofs can be modified to fit games with T periods for any finite $T \geq 2$).

Formally, suppose agent i is inspected in period 1 with probability p_i. If agent i violates, he obtains

$$v_i(V) = \big((1-p_i) - c \cdot p_i\big) + \delta \cdot \Big(p_i \cdot v_i^V + (1-p_i)v_i^{NI}\Big), \tag{5}$$

where $v_i^V = 0$ (resp., v_i^{NI}) is agent i's continuation payoff if he is inspected and found violating (resp., if he is not inspected). The first term on the right-hand side of Eq. (5) represents agent i's current-period gain from violation, and the second term represents his future payoff, depending on whether he is inspected in the current period. If agent i adheres, he obtains

$$v_i(A) = 0 + \delta \cdot \Big(p_i \cdot v_i^A + (1-p_i)v_i^{NI}\Big),$$

where v_i^A is his continuation payoff if he is inspected and found adhering.

By adhering rather than violating, agent i loses in the first period, but gains in the second period if his adhering behavior is observed by the principal. If the latter effect is sufficiently strong, that is, the agent's continuation payoff upon being found adhering is sufficiently higher than his continuation payoff upon being found violating (which is zero for PPEs that satisfy condition (iii)), then the agent is better off adhering. As the next proposition shows, the following function $f : (0,1] \to \mathbb{R}_+$ measures this difference (see Figure 2):

$$f(p) := \begin{cases} \frac{1-p-cp}{p\delta}, & \text{if } 0 < p < \frac{1}{1+c}, \\ 0, & \text{if } \frac{1}{1+c} \leq p \leq 1. \end{cases} \tag{6}$$

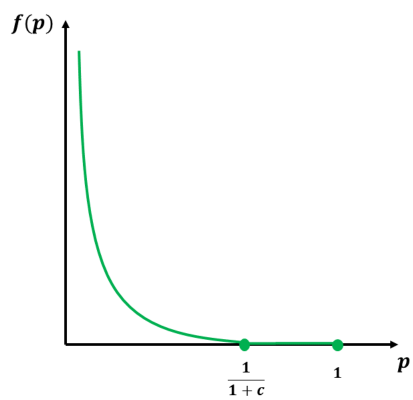

Figure 2: The function $f(p)$.

Proposition 1. *Suppose that σ is a PPE under public monitoring. Then agent i adheres in the first period if and only if $v_i^A \geq f(p_i)$.*

Proposition 1 asserts that agent i who is inspected in the first period with probability p_i adheres in that period if and only if his continuation payoff upon *being inspected and* *found adhering* is at least $f(p_i)$. Thus, $f(p_i)$ is the minimal reward to an agent who is inspected with probability p_i that ensures he adheres in the current period. The fact that f is non-increasing reflects the property that in a given period, inspection probabilities and future rewards are substitute in deterring an agent from violating. Indeed, the higher the probability of inspection, the less attractive it is to violate, hence a lower compensation is needed to ensure that the agent adheres.

Another property of the optimal PPE is that the optimal response of an agent in a given period depends only on his continuation payoff *if he is inspected* in that period. In contrast, the agent's continuation payoff *if he is not inspected*, which is the same regardless of his current-period action, does not affect the agent's incentive compatibility constraint.

This result has an important implication. To improve upon the static game and have both agents adhere in the first period, the agent who is inspected is subject to a certain reward in the future. Whereas for the uninspected agent, his continuation payoff can be set as low as possible (subject to the constraint on the inspection resource), without affecting the first-period incentives. Therefore, *the agent who is not inspected in early periods will be treated less favorably* — he will face higher inspection probability and violate less in the second period.

Optimal PPE in an example with public monitoring. To illustrate the optimal PPE, let us study a special case where $c = 0.85$ and $\delta = 0.9$. In the one-shot game, an agent adheres whenever the inspection probability he faces is higher than $\frac{1}{1+c} = 0.54$, and hence it is impossible for the principal to deter both agents from violating.

Under public monitoring, the optimal inspection scheme, denoted σ^*, adopts the following structure (see Figure 3). Under σ^*, no agent violates in period 1 and only one agent violates in period 2 — the principal's loss is δ.

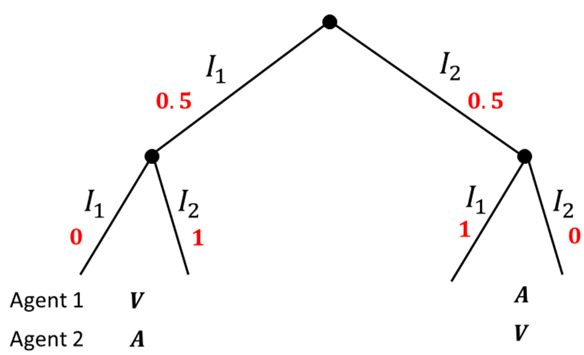

Figure 3: Optimal equilibrium under public monitoring.

- Strategy of the principal:

- In the first period, inspect each agent with probability 0.5.
- If the inspected agent is found adhering in the first period, then in the second period he is inspected with probability 0, and the other agent is inspected with probability 1.
- Punishment: if the inspected agent is found violating in the first period, he is inspected with probability 1 in the second period.

- Strategy of agent i, $i = 1, 2$:
  - Agent i adheres in the first period.
  - If agent i is inspected in the first period and found adhering, he violates in the second period.
  - In all other scenarios agent i adheres in the second period.

That is, in the first period the principal randomly chooses one agent to inspect; if this agent played A in the first period, then inspect the other agent in the second period; if this agent played V in the first period, then keep inspecting the same agent in the second period. This way, both agents play A in the first period,[20] and the agent being inspected first, who is found to play the principal's preferred action, is not inspected in the second period (so that he can play V as a reward). In equilibrium, being inspected is good news for an agent and not being inspected is bad news. Since the principal announces her observations after the first period, one violation in the second period is unavoidable (we will show that this is not the case under private monitoring).

An SE with private monitoring that yields a superior outcome. Now we turn to *private monitoring*, and show that the principal can attain a superior outcome. In particular, while making both agents adhere in the first period, with a positive probability both agents adhere also in the second period. This can be done, for instance, by the following inspection strategy, which is similar to the previous strategy, except that it assigns a positive probability to no inspection in the first period (see Figure 4).[21] Let ϵ be a sufficiently small positive number.

- Strategy of the principal:

[20] Since a successfully passed inspection in the first period leads to a reward of 1 in the second period, and since $1 > f(0.5)$, both agents are better off adhering in the first period (see Proposition 1).

[21] We omit the specification of the belief system in an SE when agents' beliefs over the private histories of the other players can be computed by Bayes' rule.

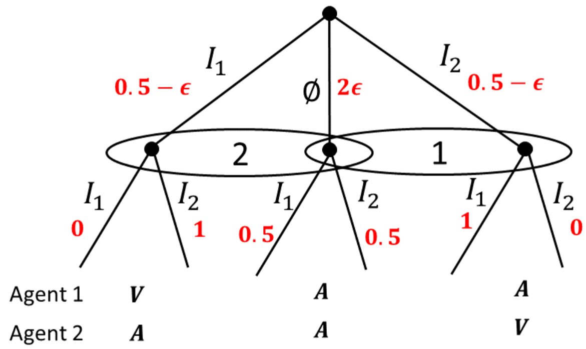

Figure 4: An inspection strategy under private monitoring.

- In the first period, the principal inspects each agent with probability $\frac{1}{2}-\epsilon$. With probability 2ϵ no agent is inspected.
- If no agent is inspected in the first period, then in the second period both agents are inspected with probability $\frac{1}{2}$.
- Otherwise, follow the same instructions as in σ^*.

- The strategies of the agents are similar to σ^*: on the equilibrium path both agents adhere in the first period, and the agent who is inspected in the first period violates in the second period.

Under private monitoring, the principal benefits from assigning a positive probability (2ϵ) to *no inspection* in the first period. The difference between this equilibrium and σ^* (the optimal PPE under public monitoring) lies on the eventuality that no agent is inspected in the first period. If this happens, each agent is inspected with probability $\frac{1}{2}$ in the second period, and, in fact, each agent is better off violating. But since an agent under private monitoring cannot distinguish between the event "no one is inspected in the first period" (which leads to an inspection probability of $\frac{1}{2}$) and the event "the other agent is inspected in the first period" (which leads to an inspection probability of 1), he considers the average situation, and is better off adhering if ϵ is not too large.[22]

From a technical perspective, the superiority of private monitoring stems from pooling some histories together, so that an agent's incentive compatibility constraints have to be satisfied only in expectation rather than state by state. This reduces the number of incentive constraints, which benefits the principal.

[22]First, ϵ has to be small (in our example, $\epsilon \leq 0.13$) so that agents are still better off adhering in the first period (that is, $1 \geq f(\frac{1}{2}-\epsilon)$). Second, ϵ has to be small so that if an agent is not inspected in the first period, the expected inspection probability he faces in the second period is at least $\frac{1}{1+c}$.

In the inspection context, the above logic can be made more concrete. As argued before, under the optimal PPE, having the opponent being inspected is bad news for an agent—this agent will face higher inspection intensity and violate less in the future. Under private monitoring, by assigning a positive probability to *no inspection* in the first period, the principal induces both agents to believe with sufficiently high probability that the opponent was inspected, and as a result both agents violate less in the future.

Remark 1. For the principal to manipulate agents' beliefs as described above, it is required that agents cannot communicate and figure out the past actions of the principal. Therefore, our results better fit scenarios under which agents cannot communicate, e.g., for legal reasons. If agents can communicate, they can possibly improve their payoffs by truthfully reporting their past histories. Indeed, cheap-talk may help the agents because coordination is desirable in the second period (see, e.g., Farrell and Gibbons, 1989). The proper analysis of the cheap-talk problem (and its variant where agents have lexicographic preferences, i.e., given an agent obtains the same payoff, he prefers the outcome where the other agent obtains a lower payoff) is an interesting and challenging subject, which we leave for future studies.

Remark 2. Is private monitoring *always* weakly superior to public monitoring? The answer is positive in all cases that we study in this paper, as long as the principal has a commitment power.[23] Yet, this result is not straightforward. Indeed, public histories are finer than private histories (see Figure 1), and hence an agent's set of public strategies is larger than his set of private strategies.[24] As a result, under public monitoring the principal has more flexibility in adjusting agents' behavior according to the true history, and this flexibility may allow a more efficient use of the inspection resource. Suppose, for instance, that under private monitoring Agent 1 faces a low inspection probability at history $h_1 = N_1$ (Agent 1 is not inspected), and hence he violates at both histories $\emptyset$ (no agent is inspected) and I_2 (Agent 2 is inspected). Under public monitoring, facing the same expected inspection probability, it is possible that Agent 1 adheres at $\emptyset$ and violates only at I_2. See Appendix A.2 for a detailed discussion.

Section 3 generalizes Theorem 1 in several aspects. Section 3.1 deals with games that involve more agents and inspectors. Section 3.2 studies the no-commitment case. Section 3.3 replaces the resource constraint by a fixed cost for each inspection in each period. Section 3.4 discusses games with a larger number of periods.

[23]When the principal lacks the ability to commit, private monitoring can be strictly inferior to public monitoring, see Section 3.2.

[24]In Appendix A.2 we provide an example where a strategy profile under public monitoring cannot be replicated by any strategy profile under private monitoring.

3 Generalizations of the benchmark model

3.1 More agents and inspectors

In the previous section, we assumed that there are two agents and one inspector. In this section, we consider the model with more inspectors and agents. We will show that when the number of agents is large relative to the number of inspectors, the logic of Theorem 1 carries over, and private monitoring can be strictly better than public monitoring. Nevertheless, as the number of agents and inspectors increases, the superiority of private monitoring over public monitoring becomes less significant. In the limit case with an infinite number of agents, the two monitoring structures yield the same optimal outcome.

Suppose there are $n \geq 3$ agents and $m \geq 1$ inspectors: m is an integer that indicates the maximum number of agents that can be inspected in each period. For every real number x denote by $\lfloor x \rfloor$ the largest integer smaller than or equal to x. As before, L^{pub} (resp., L^{pri}) is the minimum equilibrium loss for the principal under public monitoring (resp., private monitoring).

Theorem 2. *Consider the two-period game with n agents and m inspectors. Suppose $n \geq \lfloor (2+c)m \rfloor$. (i) For almost all (c, δ), Private $\succ$ Public. (ii) $L^{\text{pub}} - L^{\text{pri}} \leq \delta$, and the ratio between $L^{\text{pub}} - L^{\text{pri}}$ and the equilibrium loss goes to zero as $n \to \infty$.*

Proof. See Appendix A.3. □

Part (i) of Theorem 2 asserts that when the number of agents is large relative to the number of inspectors, private monitoring is strictly better than public monitoring for almost all (c, δ) (i.e., except a parameter set with Lebesgue measure zero). The condition $n \geq \lfloor (2+c)m \rfloor$ is not necessary for the superiority of private monitoring, but it significantly simplifies the analysis.[25]

Part (ii) of Theorem 2 asserts that even though private monitoring is superior to public monitoring, the magnitude of its net benefit is bounded above by δ. We show that the equilibrium loss in this case is at least $n - \lfloor (1+c+\delta)m \rfloor + \delta n - \delta \lfloor (1+c)m \rfloor$, which is increasing in n. Therefore, when the number of agents is small, the advantage of private monitoring over public monitoring can be significant relative to the equilibrium loss. Yet, when the number of agents is large, the upper-bound δ is small compared with the equilibrium

[25] In fact, we can construct examples where $n < \lfloor (2+c)m \rfloor$ and private monitoring is strictly superior to public monitoring. Nevertheless, for n and m that violate the conditions of Theorem 2, it is generally difficult to identify the optimal equilibrium under public monitoring, and we are unable to provide a complete comparison between the two monitoring schemes in such cases.

loss, and the difference between the two monitoring structures is negligible. The intuition for this result will be more transparent after we present the next theorem on the non-atomic model, so we postpone its discussion.

Note that L^{pub} and L^{pri} capture the *total* loss of the principal from n agents' violations. By Theorem 2, the difference between the *average* loss of the principal per agent in the two monitoring structures, $\frac{1}{n}(L^{\text{pri}} - L^{\text{pub}})$, is bounded by $\frac{\delta}{n}$, which goes to zero as n increases. As we will see in the next theorem, in the limit case with a continuum of agents (the non-atomic model), the two monitoring structures yield the same optimal outcome — the average loss of the principal is the same in both monitoring structures. Thus, the difference between private and public monitoring in the non-atomic game is the limit of the differences in finite games, and the rate of convergence is $\mathcal{O}(\frac{1}{n})$.

We next discuss the non-atomic case. Suppose the set of agents is $[0, 1]$ and the principal can inspect at most a fixed fraction of the agents, which we denote by $\alpha \in (0, 1)$. The fraction α corresponds to $\frac{m}{n}$ in the finite model. The formal definition of the non-atomic game is presented in Appendix A.4.

Theorem 3. *In the non-atomic model, Private* $\approx$ *Public for every* $0 < \alpha < 1$.

Proof. See Appendix A.5. □

To prove Theorem 3, we first show that regardless of the monitoring structure, the number of violations is bounded below by some lower bound $\underline{L}$.[26] We then construct an equilibrium under public monitoring, where the principal attains this lower-bound loss. In this equilibrium, the inspection intensity for an agent depends only on his own inspection result, and hence it can be mimicked under private monitoring.

We next provide the intuition why the benefit from manipulating information disappears in the non-atomic case. The equilibrium that attains the lower-bound violations is described at the end of this section.

Recall that the potential superiority of hiding information (private monitoring) stems from pooling some histories together, so that an agent's incentive constraint has to be satisfied only in expectation rather than state by state. This pooling is advantageous only if at some histories the inspection resource is inevitably superfluous (hence agents' incentive constraints are slack), and the pooling allows a more efficient use of this redundant resource as it can average out some other histories where the inspection resource is scarce.

[26]The value of $\underline{L}$ varies with α. When α is large, $\underline{L} = 0$; when α is moderately small, $\underline{L} = 1 - (1 + c)\alpha$; and when α is very small, $\underline{L} = [1 - (1 + c + \delta)\alpha] + [\delta - \delta(1 + c)\alpha]$.

When there is a finite number of agents, under public monitoring, inspection resources are necessarily wasted at certain histories. This may happen because after inducing some agents to adhere, the residual inspection resource is not enough to deter another violation, and hence is left unused; or because some agents are allowed a free violation (as they successfully passed inspections in the past), and hence the principal stays idle deliberately. Private monitoring allows a transfer of the inspection resource across histories, and attains a more efficient use of it.

In the non-atomic game, the inspection resource is efficiently used already under public monitoring: first, since agents are perfectly divisible, the principal can put every bit of the inspection resource to use, so there is no undesired inefficiency. Second, the principal can finely divide agents into different groups, so that the size of the set of agents who are allowed a free violation in the future can perfectly match the available inspection resource, and hence there is no deliberate inefficiency. As a result, public monitoring attains the most efficient equilibrium outcome, and there is no gain from changing the monitoring structure.

The above logic applies also to the case where there is a finite but large number of agents: in such a case, the ability of the principal to finely divide agents into groups enables a more efficient use of the inspection resource under public monitoring, and as a result the advantage of private monitoring shrinks.

We turn to describe the equilibrium that attains the lower-bound violations under public monitoring (see Figure 5): the population is divided into three groups. Agents in Group 1 are inspected in the first period with probability $p^* = f^{-1}(1)$ and they are treated dynamically — a successfully passed inspection in the first period is rewarded with a free violation in the second period. Agents in Groups 2 and 3 are inspected with probability $\frac{1}{1+c}$ and 0, respectively, and they are treated myopically — the second-period inspection intensities they face are independent of the first-period history. In the second period, except those agents in Group 1 who are allowed a free violation, the principal inspects as many other agents as possible with probability $\frac{1}{1+c}$. As a best response, in the first period agents in Groups 1 and 2 adhere, and agents in Group 3 violate. In the second period, agents who are inspected with probability $\frac{1}{1+c}$ adhere, and all other agents violate. It turns out that this simple structure can yield the lower-bound violations $\underline{L}$, so long as we choose the size of the groups carefully (as a function of α).

Remark 3. In Theorem 2(ii) and Theorem 3, in line with other sections, we focus on the comparison between public and private monitoring. In fact, stronger versions of these results hold. In Theorem 2(ii), let L^{opt} be the minimum equilibrium violations *across all* monitoring

<table>
<tr><th></th><th>Prob. of inspection</th><th>behavior</th><th></th><th>Prob. of inspection</th><th>behavior</th></tr>
<tr><td rowspan="2">Group 1</td><td rowspan="2">p^*</td><td rowspan="2">A</td><td>Group 1 inspected</td><td>0</td><td>V</td></tr>
<tr><td>Group 1 uninspected</td><td rowspan="3">Inspect as many agents as possible with probability $\frac{1}{1+c}$.</td><td rowspan="3">An agent adheres iff he is inspected with probability $\frac{1}{1+c}$.</td></tr>
<tr><td>Group 2</td><td>$\frac{1}{1+c}$</td><td>A</td><td>Group 2</td></tr>
<tr><td>Group 3</td><td>0</td><td>V</td><td>Group 3</td></tr>
<tr><td colspan="3">Period 1</td><td colspan="3">Period 2</td></tr>
</table>

Figure 5: The equilibrium that attains the lower-bound loss.

structures.[27] If $n \geq \lfloor (2+c)m \rfloor$, then $L^{\text{pub}} - L^{\text{opt}} \leq \delta$. This implies that when there is a large number of agents, public monitoring is nearly optimal, and the benefit from manipulating information is negligible. In the limit non-atomic game, public monitoring is optimal and there is no gain from changing the monitoring structure.

3.2 No commitment

In line with previous studies on inspection problems, our paper focuses on games where the principal has a commitment power.[28] The analysis of the no-commitment case is another interesting and challenging subject.[29] In this section, we show that private monitoring still has its advantage in the no-commitment case. Nevertheless, an opposite driving force that favors public monitoring emerges: public monitoring enables agents to better monitor the principal's actions and prevents the principal from deviating. We show that the advantage of private monitoring dominates when the number of agents is small, whereas the advantage of public monitoring dominates when the number of agents is large.

[27]We always assume that the inspection results are objective, so that each agent knows his own inspection history. The monitoring structure only changes an agent's knowledge about other agents' inspection histories.

[28]In reality, it is not so infrequent that the inspection agency is able to announce and commit to a certain inspection scheme in advance. Examples include speed deterrence programs that announce and implement greater policy presence on certain highways (Eeckhout, Persico, and Todd, 2010), and safety inspections managed by the Federal Aviation Administration that conduct periodic aircraft inspections on civil aviation (FAR 91.409b).

[29]Applications that fit better the no-commitment scenario include inspections on workplace safety such as the Mine Safety and Health Administration, restaurant hygiene inspections, and firms' health and safety inspections (see Jin and Leslie, 2009, Levine, Toffel, and Johnson, 2012, Dechenaux and Samuel, 2014). Such inspections are typically conduced "on-surprise" and without the principal announcing and committing to a certain inspection plan in advance.

The no-commitment model. In the no-commitment case, we assume that a strategy of a player is a function from the set of his or her information sets to the set of his or her mixed actions. In particular, we assume that the principal does not announce her inspection strategy in advance, even though in equilibrium agents correctly anticipate the principal's strategy.

Under public monitoring, we study PPE, which requires that deviations of the principal and the agents are non-profitable after every public history. Under private monitoring, we study SE, which requires that deviations of the principal and the agents are non-profitable after every private history given their beliefs. The formal definitions of PPE and SE in this game are standard, see, e.g., Mailath and Samuelson (2006) and Maschler, Solan, and Zamir (2020).

The key difference between the above definition of equilibrium and that in the commitment case (see Section 2.4) is the principal's incentive constraint: in the no-commitment case, deviations of the principal present a new condition that should be taken into account. Since the principal's payoff depends only on agents' actions (in particular, inspections are costless up to capacity), the principal is indifferent about whom to inspect in the second period. As a result, the principal's incentive to deviate is relevant only in the first period, where her action affects agents' second-period behavior.

The next result extends Theorem 1 to the no-commitment case.

Theorem 4. *Suppose there are two agents and one inspector, and the principal has no commitment power. Then in the two-period game, Private $\succ$ Public if $\delta > (1-c) \cdot \frac{3(1+c)}{c}$.*

Proof. See Appendix A.6. □

Theorem 4 asserts that when there are two agents and one inspector, private monitoring can still be superior to public monitoring, even though under a smaller set of parameters compared with the commitment case.

We first argue that in the no-commitment case there is a new driving force that favors public monitoring, and hence Private $\succ$ Public is more difficult to hold. Indeed, the principal with no-commitment is relatively well disciplined (in the sense that her deviations can be more easily detected) under public monitoring than under private monitoring. For instance, the strategy profile depicted in Figure 3 can still be supported as a PPE: agents will punish the principal by violating in the second period if they observe that the principal deviates to "no inspection" in the first period.

In contrast, under private monitoring, since agents' observations are more restricted, most deviations of the principal go unnoticed.[30] For instance, the strategy profile depicted in Figure 3 cannot be supported as an SE, because the principal is better off deviating from the prescribed strategy and inspecting no agent in the first period. Such a deviation is profitable, as it induces both agents to adhere in the second period. Moreover, it cannot be detected since an uninspected agent cannot rule out the possibility that the other agent was inspected in the first period.

We next argue that the driving force that favors private monitoring still exists in the no-commitment case, and as a result private monitoring can still be superior under some parameters. Recall that in the commitment case, private monitoring is advantageous because it enables the principal to manipulate agents' beliefs to her advantage: an agent's future payoff is higher if he passed an inspection, and lower if the other agent passed an inspection. By assigning a positive probability to no-inspection in the first period, the principal makes both agents believe that (with a high probability) the other agent was inspected in the first period, and consequently both agents violate less in the second period.

In the no-commitment case, we can show that to prevent the principal from deviating, any equilibrium under private monitoring requires that each agent's continuation payoff does not depend on whether this agent is inspected in that period. This restriction seems to undermine the benefit of private monitoring altogether, since now there is no longer any payoff boost from a passed inspection. Nevertheless, with the help of the public correlation device, we are able to restore the desired incentive.

In our construction, with the use of the correlation device, if an agent, say Agent 1, is inspected in period 1 and passes the inspection, then he gets a free pass in the second period when the realization of the correlation device is in $[0, \frac{1}{3}]$. If Agent 1 is not inspected in period 1, he gets a free pass in the second period when the realization of the correlation device is in $[\frac{1}{3}, \frac{2}{3}]$. This way, on the one hand, if the realization of the correlation device is in $[0, \frac{1}{3}]$, the agent enjoys a payoff boost from a passed inspection. This feature makes private monitoring attractive: no inspection triggers second-period adherence from both agents in some contingencies, a feature that is absent in even the best public monitoring PPE. On the other hand, this agent's expected continuation payoff remains the same regardless of whether he is inspected or not in the first period. This ensures that the principal is indifferent between inspecting the agent or not. See Figure 11 in Appendix A.6 for the detailed construction.

The next theorem deals with the non-atomic game. Recall that in each period the

[30]The only exceptions are histories under which the principal is supposed to (not) inspect an agent with certainty.

maximal fraction of agents the principal can inspect is α.

Theorem 5. *Suppose the principal has no commitment power. In the non-atomic problem with two periods, if* $\frac{1}{2+2c+\delta} < \alpha < \frac{1}{1+c}$*, then Private* $\prec$ *Public. Otherwise, Public* $\approx$ *Private.*

Proof. See Appendix A.7. □

To prove the optimality of public monitoring in Theorem 5, we show that the lower bound $\underline{L}$ identified in the commitment case can be attained as a PPE outcome in the no-commitment case,[31] and therefore no other monitoring structure (even with the use of the correlation device) can improve on public monitoring.

The intuition is similar to that of Theorem 3. For private monitoring to be superior to public monitoring, it must be the case that pooling some histories together is advantageous for the principal. This happens only if at some histories the agents' incentive constraints are slack (so that some inspection resources are wasted). This holds true for the two-agent problem in Theorem 4. When the set of agents is $[0, 1]$, the ability to perfectly divide agents into groups enables the most efficient use of the inspection resource under public monitoring, and agents' incentive constraints are binding at any history. As a result, there is no longer any gain from pooling histories together.

We now explain why public monitoring is sometimes strictly superior to private monitoring. Recall that under public monitoring, by observing the proportion of agents that are inspected, a deviation of the principal can be easily detected and punished. However, under private monitoring, since agents do not observe the proportion of inspected agents, they cannot detect deviations of the principal. To deter the principal from not inspecting an agent i who is supposed to be inspected, agent i's second-period payoff at history N_i (where he is not inspected in the first period) has to be no less than his second-period payoff at history I_i (where he is inspected in the first period). This happens in particular when agent i's second-period payoff is independent of his inspection history.

When α is large ($\alpha \geq \frac{1}{1+c}$) or small ($\alpha \leq \frac{1}{2+2c+\delta}$), we can show that the PPE that attains the lower bound $\underline{L}$ has the feature that each agent's action in the second period is independent of his inspection history. The same strategy profile forms an equilibrium also under private monitoring, as the principal has no incentive to deviate.

Indeed, when α is large, the inspection resource is superfluous, and under the optimal equilibrium both agents adhere in both periods, regardless of the history. When α is small, the inspection resource is scarce, and under the optimal equilibrium each agent who is supposed to adhere in the first period gets a free pass in the second period, regardless of

[31] Since every equilibrium under no-commitment forms an equilibrium under commitment, the value $\underline{L}$ also constitutes a lower bound on the number of violations in the no-commitment case.

whether he is inspected in the first period (as long as he is not found violating). Such an inspection scheme seems inefficient at first glance, as it gives an agent a reward (that is, a free violation in the second period) even if this agent is not inspected. Nevertheless, since the inspection resource is scarce and a large number of violations in the second period is unavoidable, such "generous" inspection strategy is optimal.

When α is in the intermediate range ($\frac{1}{2+2c+\delta} < \alpha < \frac{1}{1+c}$), the PPE that attains the lower bound $\underline{L}$ requires agents' continuation payoffs to be history dependent: an agent's continuation payoff following no inspection has to be lower than his continuation payoff following a successfully passed inspection. In particular, a "generous" inspection strategy that gives all agents who are supposed to adhere a free pass (even if they are not inspected in the first period) is no longer optimal, since some inspection resources are wasted in the second period. The optimal PPE in this case fails to be an equilibrium under private monitoring, as the principal benefits from deviating to "no inspecting" in the first period. As a result, public monitoring is strictly superior to private monitoring.

Remark 4. We assume that the principal is subject to a resource constraint, and the inspection within the constraint costs zero. Suppose, instead, that inspection is costly. Then the principal who lacks the ability to commit has an incentive to shirk (that is, to deviate to no inspection). It turns out that in such cases, violations by both agents in both periods constitute the unique equilibrium regardless of the monitoring structure.[32] The argument relies heavily on the assumption that the goal of the principal is to deter violations, rather than to collect penalties (in Section 4.1 we discuss cases where the principal benefits from catching violations). The full defection outcome here is in sharp contrast to the case where the principal has a commitment power, which is studied in the next section.

3.3 Costly inspection

In the benchmark model, it is assumed that the principal is subject to a resource constraint: she can inspect at most one out of the two agents in each period. Suppose now that the principal no longer has a resource constraint, but rather that there is a fixed cost for each inspection in each period. Formally, suppose that in each period, the cost for the first inspection is r_1 and that for the second inspection is r_2, where $r_1 \leq r_2$. The benchmark model corresponds to the special case $r_1 = 0$ and $r_2 = \infty$.

[32]The principal is always better off deviating to no inspection in the second period. Therefore, regardless of the history, both agents violate in period 2. Since the agents' first-period actions do not affect their second-period payoff, they play a Nash equilibrium of the one-shot game in period 1 — again, violation by both agents.

Suppose $r_1 = r_2$. Then regardless of the monitoring structure, it is optimal for the principal to treat the agents independently, and hence changing the monitoring structure in this case does not benefit the principal (see Appendix A.8). This implies that the positive effect of private monitoring in Theorem 1 stems from the tension between agents that results from the non-linearity of the inspection cost.

An interesting question is, therefore, when $r_1 < r_2$ yet the difference is relatively small, whether private monitoring is superior to public monitoring. The next result shows that the logic of the superiority of private monitoring applies also to this case. For this result, we focus on values r_1 and r_2 that satisfy a *divergence condition*, which states that in the one-shot game, the principal is best off deterring only one agent from violating. In fact, the divergence condition not only rules out the use of the second inspection in the one-shot game, it also rules out the second inspection in the optimal PPE under the two-period game (see Appendix A.9). The case where the divergence condition is violated is discussed in Appendix A.10.

Definition 1. The inspection costs (r_1, r_2) satisfy the *divergence condition* if $r_1 < 1 + c$ and $r_2 > 1 + c + \frac{c(1+c-r_1)}{1-c}$.

Note that for pairs (r_1, r_2) that satisfy the divergence condition, as r_1 increases to $1 + c$, the lower bound on r_2 decreases to $1 + c$. That is, r_1 and r_2 can be arbitrarily close yet satisfy the divergence condition. The next result extends Theorem 1, and follows the same line of logic.

Theorem 6. *Suppose there are two agents and (r_1, r_2) satisfy the divergence condition. If $\delta > 1 - c$, then Private $\succ$ Public in the two-period game.*

Proof. See Appendix A.9. □

Theorem 6 implies that the key driving force of the superiority of private monitoring is the increasing marginal cost of inspection: since the first inspection is cheaper than the second one, if an agent is inspected with a lower intensity, then there are more "cheap" inspection resources left for the other agent, and hence the other agent is more likely to face a higher level of scrutiny. A pooling of the inspection resource across histories is advantages only when such tension between agents exists. Theorem 1 studies the extreme case $(r_1 = 0, r_2 = \infty)$, and Theorem 6 generalizes this result to all parameters that satisfy the divergence condition.

3.4 T-period games with $T > 2$

In this section we generalize the two-period benchmark model to games that involve $T > 2$ periods. We show that when players are relatively patient, private monitoring is superior to public monitoring as long as public monitoring is not the optimal information revelation scheme for the principal. It is well known that the analysis of dynamic games with private monitoring is challenging (see, e.g., Kandori, 2002). This is because players do not necessarily know each other's continuation strategies, as they have different information. Hence they have to compute the beliefs about what the opponents are going to do, and these beliefs become fairly complex over time. In this section we propose a novel approach to tackle this problem.[33]

The next theorem provides conditions under which Private $\succ$ Public in the T-period game. For any $T \geq 2$, let L_T^{pub} and L_T^{pri} be the principal's minimum equilibrium loss under public and private monitoring, respectively, in the T-period game. By Fact 1, both L_T^{pub} and L_T^{pri} are bounded below by $L_L := 1 - c$.

Theorem 7. *Suppose there are two agents and one inspector, and $\delta > \frac{1-c^2}{c}$. In the T-period game, if $L_T^{\text{pub}} = L_L$, then Private $\approx$ Public. If $L_T^{\text{pub}} > L_L$, then Private $\succ$ Public.*

Proof. See Appendix B.1. □

The condition $\delta > \frac{1-c^2}{c}$ guarantees that Private $\succ$ Public for $T = 2$, and it serves to simplify the analysis.

In the T-period game, if public monitoring attains the lower-bound loss L_L, then clearly private monitoring cannot improve on this outcome. We show that the optimal PPE can be mimicked by an SE with the same path, and hence Private $\approx$ Public in this case. If, however, the optimal PPE under public monitoring *does not* attain L_L, then by Theorem 7, there are equilibria under private monitoring that yield the principal a loss lower than L_T^{pub}.

The intuition of the superiority of private monitoring in the T-period game is similar to that in the two-period game: private monitoring allows a transfer of the inspection resource across histories, and attains a more efficient use of it.

To construct an SE under private monitoring that is superior to the best PPE outcome, we propose a novel approach that resembles backward induction. We first construct the optimal PPE under public monitoring, and show that this equilibrium has an SE under private monitoring with the same path. We then modify the last two periods of this SE: in the new SE, the agents' beliefs on past plays are identical to their beliefs under the original

[33]In the literature on private monitoring, block strategies are sometimes used in the construction of SE to avoid the difficult task of keeping track of agents' beliefs in long horizons (e.g., Fuchs, 2007, and Sugaya, 2012). Even though this approach is helpful in establishing the Folk Theorem in games with private monitoring, it does not help with our problem: to compare private and public monitoring, it is not sufficient to show that private monitoring can attain the best feasible payoff when $\delta \to 1$. Instead, we need to show that for any fixed δ, private monitoring yields an outcome superior to the best equilibrium under public monitoring. Since the optimal equilibrium under public monitoring is not a block construction, restricting attention to block constructions cannot help us finding a superior equilibrium in private monitoring.

SE in all $T-1$ periods, and differ only in the last period. This approach significantly simplifies the analysis of agents' incentives and makes the proof tractable.

In Appendix A.11 we provide conditions under which $L_T^{\text{pub}} = L_L$, and discuss the effect of T and δ on L_T^{pub}.

4 Discussion

4.1 The principal's additional gain from catching violations

In this paper we focus only on the principal's incentive to deter violations, and leave aside her potential gain from catching violations. We now consider the case where the principal benefits from catching violations, e.g., for the concern of public opinion or collecting fines. In the commitment case, as long as the gain from catching a violation is small compared with the damage caused by the violation, the principal's incentive remains essentially the same and the logic of Theorem 1 applies. This is because inspecting an agent who violates has a positive direct effect on the principal's payoff, yet it reduces the agent's gain from violation; this in turn affects negatively the agent's incentive in earlier periods, thereby inflicting a negative indirect effect on the principal's payoff.

In the no-commitment case, however, assuming that the principal benefits from catching violations has a drastic impact on the equilibrium outcome. The principal now has another deviation, which is to catch those she is supposed to give a free pass. It turns out that such incentive prevents the principal from disciplining agents: in all equilibria of the game with two agents and one inspector, under both public and private monitoring, both agents violate in both periods (see Appendix A.12). The intuition is that if in equilibrium one agent adheres and another one violates, then the principal is better off deviating from the prescribed strategy, and shifting all inspection resources from the adhering agent to the violating agent (in an attempt to catch violations), which upsets the equilibrium. It follows that rewarding an agent who successfully passes an inspection is difficult,[34] and this restricts the principal's ability to use future inspection intensities to discipline agents.

4.2 Optimal inspection scheme under private monitoring

In Section 1, to show that private monitoring is superior to public monitoring in the benchmark model, we constructed one SE under private monitoring that is superior to the optimal PPE under public monitoring. A question arises regarding the structure of the opti-

[34]This is possible only if both agents violate in the rewarding stage.

mal SE under private monitoring. This is the content of the current section. For simplicity, in this section we exclude the use of the public correlation device.

To identify the optimal SE, we need to use the first-period inspections optimally so as to generate efficient private signals to deter most violations in the second period. A difficulty is that the first-period inspections cannot be arbitrary, since they have to take into account the agents' first-period incentives.

Recall that to induce an agent to adhere in the first period, there are two options. The first is to inspect the agent with sufficiently high probability (at least $\frac{1}{1+c}$), so that adhering becomes a myopic best response. The second is to inspect the agent with a lower probability and use future violations as a reward for a successfully passed inspection. In the two-period game, the maximum reward is one free violation in the second period. This imposes a lower bound $p^* = f^{-1}(1) = \frac{1}{1+c+\delta}$ on the first-period inspection that satisfies incentive constraints.

We will show that under private monitoring, when c is small, the principal's minimum loss is $2\delta p^*$: each agent is inspected with probability p^* in the first period, and the one who is inspected in the first period violates in the second period. When c is large, the principal's equilibrium loss drops to δp^*: one agent is always inspected with probability $\frac{1}{1+c}$ and he always adheres; whereas the other agent is inspected with probability p^* in the first period, and he violates in the second period if he is inspected in the first period. Thus, an increase on the punishment level c allows the principal to better discipline agents.

In line with Theorem 1, we focus on the case where $\delta > 1-c$. This simplifying assumption guarantees that the principal can induce both agents to adhere in the first period. Divide the parameter space into three regions, as shown in Figure 6. The next proposition characterizes the optimal inspection scheme under private monitoring.

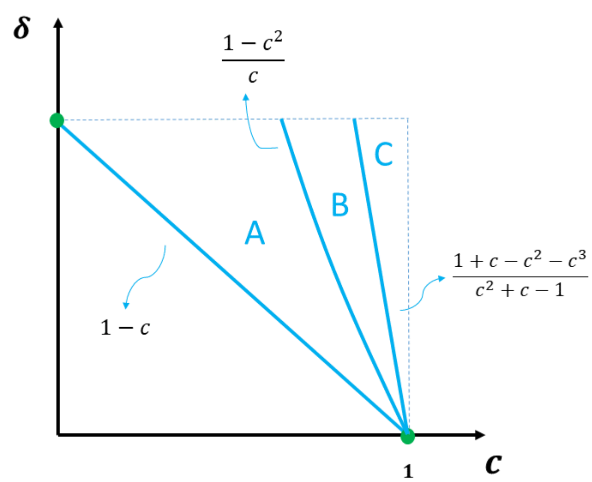

Figure 6: The parameter space.

Proposition 2. *(i) For (c, δ) in regions A and B, the optimal inspection scheme under private monitoring (as long as no violation is detected) follows Figure 7. (ii) For (c, δ) in region C, the optimal inspection scheme under private monitoring (as long as no violation is detected) follows Figure 8. (iii) Punishment: If an agent is found violating in period 1, he is inspected with probability 1 in period 2.*

Proof. See Appendix A.13. □

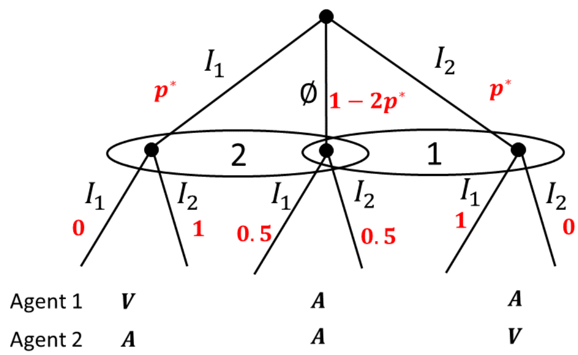

Figure 7: Optimal SE for parameters in regions A and B.

For parameters in regions A and B, under the inspection scheme depicted in Figure 7, both agents are inspected with probability p^* in the first period and they adhere. In the second period, the agent who is inspected in the first period faces inspection intensity zero and he violates. The agent who is not inspected in the first period considers the average situation and is better off adhering. The loss of the principal in this equilibrium is $2\delta p^*$: each of the agents violates in the second period with probability p^*.

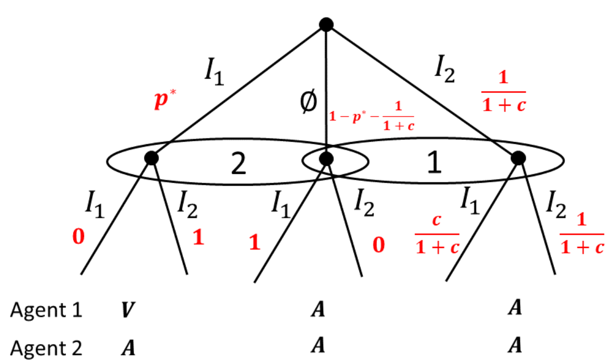

Figure 8: Optimal SE for parameters in region C.

For parameters in region C, under the inspection scheme depicted in Figure 8, Agent 2 is treated myopically and he always adheres (that is, the expected inspection intensity he faces is at least $\frac{1}{1+c}$ regardless of his private history). As for Agent 1, he adheres in the first period in exchange for the chance of a free violation in the second period. In the second period, Agent 1 violates if and only if he was inspected in the first period. The loss of the principal in this equilibrium is δp^*: Agent 2 always adheres and Agent 1 violates in the second period with probability p^*.

The key difference between the strategy profiles described in Figures 7 and 8 is that in the latter, Agent 2 adheres in the second period even if he is *inspected and found adhering* in the first period. This implies that a successfully passed inspection for Agent 2 in the first period is not rewarded, and hence to deter him from violating in the first period, an inspection probability of $\frac{1}{1+c}$ is needed. In fact, in the latter construction, Agent 2 is inspected with probability at least $\frac{1}{1+c}$ regardless of his private history. If c is small (parameters in regions A and B), $\frac{1}{1+c}$ is large, which implies that the inspection of Agent 2 takes too many inspection resources, and the latter construction cannot be supported as an equilibrium: for parameters in Region A, $1 - \frac{1}{1+c} < p^*$ and Agent 1 is better off violating in the first period. For parameters in Region B, Agent 1 adheres in the first period, but if he is not inspected in the first period, the expected inspection intensity he faces is

$$1 \cdot \frac{1 - p^* - \frac{1}{1+c}}{1 - p^*} + \frac{c}{1+c} \cdot \frac{\frac{1}{1+c}}{1 - p^*},$$

which is lower than $\frac{1}{1+c}$, and he is better off violating in the second period.

4.3 Optimal information revelation

In this paper we focus on two commonly observed monitoring schemes: private monitoring and public monitoring. Another interesting question concerns the identification of the optimal way for the principal to reveal information. In the non-atomic problem, we've stated in Remark 3 that public monitoring is optimal, and no other information-revelation scheme can improve on its outcome. This section is devoted to problems with finitely many agents, in particular, with one inspector and two agents, and we show that the optimal revelation scheme can outperform both public and private monitoring.

We first show that even though the lower-bound violations $1 - c$ (see Fact 1) cannot be attained under public and private monitoring, it can be supported by a more complicated mechanism of information revelation. We then discuss the practical issues that may prevent the implementation of the optimal information revelation scheme.

By a version of the revelation principle (see Forges, 1986), without loss of generality one can assume that in each period the principal privately sends to each agent a recommended action, and each agent follows the recommendation he gets.

One-shot game. In the one-shot game, when private signals are available, the principal can use the following method to reduce her loss to $1-c$. To this end, the principal sends recommendations as detailed in Table 1. For instance, with probability $1-c$ the principal inspects Agent 2, Agent 1 is told to violate, and Agent 2 is told to adhere. Note that Agent 2 is told to adhere in all three rows of the table, and hence the recommendation of adhering is not informative for Agent 2.

Probability	Principal	Agent 1	Agent 2
$1-c$	I_2	V	A
$\frac{c}{1+c}$	I_1	A	A
$\frac{c^2}{1+c}$	I_2	A	A

Table 1: Recommendations in the one-shot game.

When Agent 1 gets the signal to violate, he knows that he is not inspected, and hence he is better off violating. If Agent 1 gets the signal to adhere, then with probability $\frac{\frac{c}{1+c}}{\frac{c}{1+c}+\frac{c^2}{1+c}} = \frac{1}{1+c}$ he is inspected and hence he is better off adhering. It can be verified that Agent 2 is also better off following the recommended actions. When both agents follow the recommendations, their payoffs are $(v_1 = 1-c, v_2 = 0)$, and the principal's loss attains the lower bound.

Under the optimal information revelation mechanism, whenever an agent is told to violate, the agent knows that he will not be inspected. When an agent is told to adhere, he is not sure whether he will be inspected or not: the expected probability that the agent will be inspected is $\frac{1}{1+c}$.

Intuitively, when the principal follows the recommendations in Table 1, Agent 2 always adheres and $\frac{1}{1+c}$ inspection resources are allocated to him. As a result, $\frac{c}{1+c}$ inspection resources are left for Agent 1. Agent 1 can be disciplined on an event E if on that event, the conditional probability he is inspected is $\frac{1}{1+c}$. This implies that $P(E) = c$. On the complement of E, Agent 1 violates.

Note the connection between the current exercise and the second-period of Figure 4 (private-monitoring game with $T = 2$). In the two-period private-monitoring game, the agents' private history in the first period serves as a private signal in the second period. However, unlike in the current case where private signals are costless, in the model with private monitoring, sending signals by private history is "costly" for the principal: the probabilities according to which agents are inspected in the first period affect not only the agents' second-period beliefs, but also their first-period incentives.

In Appendix A.14 we analyze the two-period game, and show that the optimal information-revelation rule in the two-period game is not a simple repetition of the solution to the one-shot game. This is because the principal can use the second period to better discipline agents in the first period.

To attain the optimal information-revelation outcome, the principal has to *privately* communicate with each agent regarding her intent to inspect *before* agents take actions.[35] There are several practical difficulties that may prevent the implementation of this approach. First, ethical concerns arise when the inspector privately tells an agent what to do, and, in particular, instructs an agent to violate. Such a behavior could be perceived by the public as collusion. Moreover, in some cases, this may facilitate real collusion: to implement the optimal outcome, the inspector has to draw a lottery that determines the identity of the inspected agent (and send private signals accordingly) before agents take actions. If some of the officials in the regulatory department are corrupt, a message leak is very likely.[36]

Such concerns are absent in our main model. First, the principal only *publicly* reveals past inspection histories (or not), and there is no need to privately communicate with agents. Second, even though the principal's future inspection plan is pre-announced (in the commitment case), the lottery that determines the identity of the inspected agent can be drawn in the last minute prior to inspection (and hence after the agents take actions), so the risk of a message leak is absent.

This may be the reason why in practice, both pre-announced inspections (analogous to the commitment case) and surprise inspections (analogous to the no-commitment case) are commonly used and accepted (with or without a public disclosure of past inspection results).[37] Inspections under which the principal privately gives tip-offs to agents are rarely seen.

Another interesting question is whether there is an information design approach that relies on public signals only, and still beats both public and private monitoring. We study the two-period benchmark case in Appendix A.15 and show that when focusing on public signals, the principal's ability to manipulate agents' beliefs is significantly restricted: for

[35] For instance, to implement the payoffs in Table 1, with probability $1 - c$ the principal privately advices Agent 1 to violate and Agent 2 to adhere.

[36] Message leaks from a regulatory department are not so infrequent. For example, the New York Department of Transportation conducts inspections of school bus companies to test the road-worthiness of their buses. School-bus supervisors at the department of transportation in New York, however, were later known to accept bribes from firms in exchange for the information on upcoming inspections (Von Zielbauer, 2009).

[37] Applications that fit the commitment case include speed deterrence programs that announce and implement greater policy presence on certain highways (Eeckhout, Persico, and Todd, 2010), and safety inspections managed by the Federal Aviation Administration, which conducts periodic aircraft inspections on civil aviation (FAR 91.409b). On the other hand, inspections on workplace safety such as inspections by the Mine Safety and Health Administration, restaurant hygiene inspections, and firms' health and safety inspections are typically conducted "on-surprise" and fit better the no-commitment scenario (see Jin and Leslie, 2009, Levine, Toffel, and Johnson, 2012, Dechenaux and Samuel, 2014).

most parameters, public signals do not benefit the principal. For a small set of parameters, public signals can be advantageous. Nevertheless, the desirable level of public disclosure is sensitive to changes in the parameters, and its analysis is rather complicated.

5 Conclusion

In this paper we study a model with a principal who tries to use her limited inspection resource to induce adherence from multiple agents. The principal observes the full inspection history, whereas each agent only observes what happens to himself. To use her inspection resource most efficiently, the principal decides not only on the inspection scheme (i.e., whom to inspect after each history), but also whether to reveal information about past inspection histories.

We show that when the number of agents is small and the principal has commitment power, the principal can better discipline agents if she does not publicly disclose inspection histories. The key observations are as follows. First, to improve upon the static game, a successfully passed inspection has to come with a future payoff bonus to the agent; here, the promise of the principal turning a blind eye with positive probability, so that the agent can violate without penalty. As such, being inspected is good news for an agent. When the total inspection resource is limited, less enforcement on certain agents implies greater enforcement on other agents. Consequently, having the other agents being inspected in early periods is bad news for an agent: he is subject to greater enforcement and hence violates less in the long run. By keeping past inspections private and sometimes not inspecting anybody, the principal can keep all agents on their toes in the long run: they now all have some lingering doubt that their opponents may have just successfully passed an inspection and therefore fear increased scrutiny themselves.

Several key assumptions are crucial for the above logic. First, there must exist a negative correlation between the inspection intensities for agents. This happens if the inspection resource is limited, or if the marginal inspection cost is increasing. If, instead, the inspection cost is linear in the number of inspections, it is optimal for the principal to treat each agent independently, and all monitoring structures yield the same optimal outcome.

Second, the superiority of private monitoring is significant only when there is a small number of players. When the numbers of agents and inspectors are larger, public monitoring attains a more efficient use of the inspection resource, and the gain from changing the monitoring structure decreases.

Finally, the principal's commitment power allows her to take better advantage from concealing information (that is, using private monitoring). This is because when the principal cannot commit, under private monitoring she has too many potential deviations that may upset a desirable equilibrium. As a result, public monitoring may work better in this scenario. The importance of commitment has the following policy implication: in environmental protection problems or auditing problems, it is beneficial to have all inspection rules made public. This way, by concealing information about past inspections, the principal can further improve her payoff by manipulating agents' beliefs.

References

[1] Isaiah Andrews and Daniel Barron. "The allocation of future business: Dynamic relational contracts with multiple agents". *American Economic Review* 106.9 (2016), pp. 2742–59.

[2] Itai Arieli and Yakov Babichenko. "Private Bayesian persuasion". *Journal of Economic Theory* 182 (2019), pp. 185–217.

[3] Robert Aumman. "Mixed and behavior strategies in infinite extensive games, advances in game theory". *Annals of Mathematics Studies* (1964), pp. 627–650.

[4] Dirk Bergemann and Stephen Morris. "Bayes correlated equilibrium and the comparison of information structures in games". *Theoretical Economics* 11.2 (2016), pp. 487–522.

[5] Venkataraman. Bhaskar and Eric Van Damme. "Moral hazard and private monitoring". *Journal of Economic Theory* 102.1 (2002), pp. 16–39.

[6] Allen Blackman, Shakeb Afsah, and Damayanti Ratunanda. "How do public disclosure pollution control programs work? Evidence from Indonesia". *Human Ecology Review* 11.3 (2004), pp. 235–246.

[7] Vincent P. Crawford. "Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions". *American Economic Review* 93 (2003), pp. 133–149.

[8] Emmanuel Dechenaux and Andrew Samuel. "Announced vs. surprise inspections with tipping-off". *European Journal of Political Economy* 34 (2014), pp. 167–183.

[9] Jeffrey A. Dubin and Louis L. Wilde. "An empirical analysis of federal income tax auditing and compliance". *National Tax Journal* (1988), pp. 61–74.

[10] Jan Eeckhout, Nicola Persico, and Petra E. Todd. "A theory of optimal random crackdowns". *American Economic Review* 100.3 (2010), pp. 1104–1135.

[11] Joseph Farrell and Robert Gibbons. "Cheap talk can matter in bargaining". *Journal of Economic Theory* 48 (1989), pp. 221–237.

[12] Henk Folmer and Tom Tietenberg. *International Yearbook of Environmental and Resource Economics 2006/2007*. Edward Elgar Publishing, 2007.

[13] Francoise Forges. "An approach to communication equilibria". *Econometrica* (1986), pp. 1375–1385.

[14] Jérôme Foulon, Paul Lanoie, and Benoît Laplante. "Incentives for pollution control: regulation or information?" *Journal of Environmental Economics and Management* 44.1 (2002), pp. 169–187.

[15] William Fuchs. "Contracting with repeated moral hazard and private evaluations". *American Economic Review* 97.4 (2007), pp. 1432–1448.

[16] Wayne B. Gray and Mary E. Deily. "Compliance and enforcement: Air pollution regulation in the US steel industry". *Journal of Environmental Economics and Management* 31.1 (1996), pp. 96–111.

[17] James T. Hamilton. "Pollution as news: Media and stock market reactions to the toxics release inventory data". *Journal of Environmental Economics and Management* 28.1 (1995), pp. 98–113.

[18] Ginger Zhe Jin and Phillip Leslie. "Reputational incentives for restaurant hygiene". *American Economic Journal: Microeconomics* 1.1 (2009), pp. 237–267.

[19] Michihiro Kandori. "Introduction to repeated games with private monitoring". *Journal of Economic Theory* 102 (2002), pp. 1–15.

[20] Michihiro Kandori and Ichiro Obara. "Efficiency in repeated games revisited: The role of private strategies". *Econometrica* 74.2 (2006), pp. 499–519.

[21] Shameek Konar and Mark A. Cohen. "Information as regulation: The effect of community right to know laws on toxic emissions". *Journal of Environmental Economics and Management* 32.1 (1997), pp. 109–124.

[22] Paul Lanoie, Benoit Laplante, and Maite Roy. *Can Capital Markets Create Incentives for Pollution Control?* The World Bank, 1997.

[23] Ehud Lehrer. "Internal correlation in repeated games". *International Journal of Game Theory* 19.4 (1991), pp. 431–456.

[24] Jonathan Levin. "Relational incentive contracts". *American Economic Review* 93.3 (2003), pp. 835–857.

[25] David I. Levine, Michael W. Toffel, and Matthew S. Johnson. "Randomized government safety inspections reduce worker injuries with no detectable job loss". *Science* 336.6083 (2012), pp. 907–911.

[26] W. Bentley MacLeod. "Optimal contracting with subjective evaluation". *American Economic Review* 93.1 (2003), pp. 216–240.

[27] George J. Mailath, Steven A. Matthews, and Tadashi Sekiguchi. "Private strategies in finitely repeated games with imperfect public monitoring". *Contributions to Theoretical Economics* 2.1 (2002), pp. 1–23.

[28] George J. Mailath and Larry Samuelson. *Repeated games and reputations: long-run relationships*. Oxford university press, 2006.

[29] Mic Maschler, Eilon Solan, and Shmuel Zamir. *Game theory*. Cambridge University Press, 2020.

[30] Jordi Massó. "Undiscounted equilibrium payoffs of repeated games with a continuum of players". *Journal of Mathematical Economics* 22.3 (1993), pp. 243–264.

[31] Hitoshi Matsushima. "On the theory of repeated games with private information: Part I: anti-folk theorem without communication". *Economics Letters* 35.3 (1991), pp. 253–256.

[32] Christopher Phelan and Andrzej Skrzypacz. "Beliefs and private monitoring". *Review of Economic Studies* 79.4 (2012), pp. 1637–1660.

[33] Hamid Sabourian. "Anonymous repeated games with a large number of players and random outcomes". *Journal of Economic Theory* 51.1 (1990), pp. 92–110.

[34] Eilon Solan and Chang Zhao. "Dynamic inspection under resource constraints". *Games and Economic Behavior* 129 (2021), pp. 476–491.

[35] Mark Stephan. "Environmental information disclosure programs: They work, but why?" *Social Science Quarterly* 83.1 (2002), pp. 190–205.

[36] Takuo Sugaya. "The Folk theorem in repeated games with private monitoring". PhD thesis. Princeton University, 2012.

[37] Ina Taneva. "Information design". *American Economic Journal: Microeconomics* 11.4 (2019), pp. 151–85.

[38] Tom Tietenberg and David Wheeler. "Empowering the community: Information strategies for pollution control". *Frontiers of Environmental Economics* (2001), pp. 85–120.

[39] Paul Von Zielbauer. "School bus inspector is sentenced in bribery scheme". *The New York Times* (June 2009).

A Appendix A. Proofs

A.1 Proof of Theorem 1

Under public monitoring, since the second-period equilibrium outcome must be a Nash equilibrium of the stage game, one violation in the second-period is unavoidable. Consequently, the optimal PPE yields the principal a loss no less than δ. It is therefore sufficient to show that under private monitoring, there exists an SE that yields the principal a loss lower than δ. Consider the strategy profile $\sigma_\epsilon = \big(\sigma_{\epsilon 1}, \sigma_{\epsilon 2}\big)$ displayed in Figure 4. The strategy σ_ϵ is similar to σ^{pub}, except that it assigns a positive probability to no inspection in the first period.

We now specify the proper beliefs of the players. In each period an agent has only two actions: Adhere and Violate. Let $\sigma^k_{\epsilon i}$ be the strategy that chooses the action prescribed by $\sigma_{\epsilon i}$ with probability $\frac{k-1}{k}$, and the other action with probability $\frac{1}{k}$. Naturally, $\lim_{k\to\infty} \sigma^k_{\epsilon i} = \sigma_{\epsilon i}$. The reader can verify that the sequence of beliefs $(\mu_{\epsilon,k})_{k\in\mathbb{N}}$ induced by $(\sigma^k_\epsilon)_{k\in\mathbb{N}}$ converges as k goes to ∞. Denote the limit belief system by μ_ϵ:

$$\mu_\epsilon(h) := \lim_{k\to\infty} \mu_{\epsilon,k}(h), \quad \forall h \in \mathscr{H}. \tag{7}$$

Under μ_ϵ, if an agent, say, agent i, is inspected in the first period, he believes that he will be inspected with probability 0 in the second period. While if agent i is not inspected in the first period, he assigns probability $\frac{\frac{1}{2}-\epsilon}{\frac{1}{2}+\epsilon}$ to the event that the other agent j was inspected in the first period and subsequently agent i will be inspected with probability 1 in the second period; and agent i assigns probability $1-\frac{\frac{1}{2}-\epsilon}{\frac{1}{2}+\epsilon}$ to the event that no one was inspected in the first period and subsequently he will be inspected with probability $\frac{1}{2}$ in the second period. Thus, if agent i is not inspected in the first period, he assigns probability $\frac{1}{1+2\epsilon}$ to the event that he will be inspected in the second period.

We now argue that the strategy profile σ_ϵ and the belief system μ_{σ_ϵ} constitute an SE, provided ϵ is sufficiently small.

We start by considering the second period. If an agent is inspected in the first period and found adhering, then according to μ_ϵ, he believes he will be inspected in the second period with probability 0, and hence violating in the second period is the best response. As for the agent who is not inspected in the first period, then according to μ_{σ_ϵ}, he assigns probability $\frac{1}{1+2\epsilon}$ to himself being inspected in the second period. When ϵ is not too large (in fact, for $\epsilon \le \frac{c}{2}$), the expected inspection probability this agent faces is larger than $\frac{1}{1+c}$, and hence adhering in the second period is the best response.

Consider now the first period. Since $\delta > 1-c$, we have $1 > f(\frac{1}{2})$. Since the function f

is decreasing, for sufficiently small ϵ (in fact, for $\epsilon \leq \frac{1}{2} - f^{-1}(1)$), we have $1 > f(\frac{1}{2} - \epsilon)$, and hence adhering is the best response for both agents in the first period.

To summarize, for $\epsilon \leq \min\{\frac{c}{2}, \frac{1}{2} - f^{-1}(1)\}$, the strategy profile σ_ϵ together with the belief system μ_ϵ is an SE in G_2^{pri}, and it yields the payoffs $v_1(\sigma_\epsilon) = (\frac{1}{2} - \epsilon)\delta$ and $v_2(\sigma_\epsilon) = (\frac{1}{2} - \epsilon)\delta$ to the agents. Here $v_1(\sigma_\epsilon) + v_2(\sigma_\epsilon) < \delta$, as desired.

A.2 The potential advantage of public monitoring

In this section, we argue that in principle, public monitoring has its advantage in repeated interactions, even though this advantage is dominated by the advantage of private monitoring when the principal has a commitment power.[38]

As shown in Figure 1, public histories are more informative than private histories. Therefore, agents have a larger set of strategies under public monitoring. As a result, equilibria under public monitoring need not be implementable under private monitoring. This point is illustrated by the following example. Consider the equilibrium under public monitoring given in Figure 9. We will argue that the same strategy profile is not feasible under private monitoring.

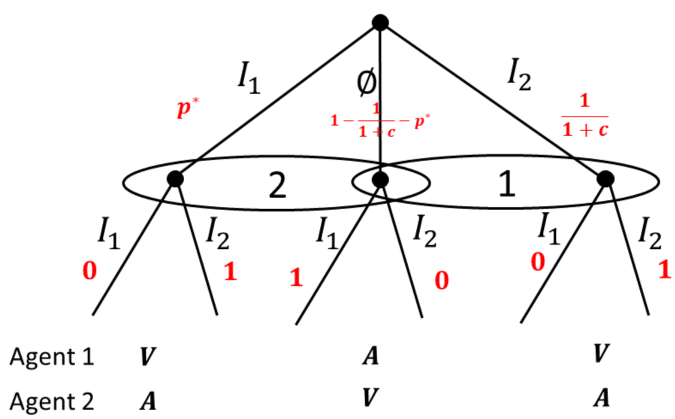

Figure 9: An inspection strategy under public monitoring.

In the equilibrium shown in Figure 9, in the first period Agents 1 and 2 are inspected with probabilities $p^* = \frac{1}{1+c+\delta}$ and $\frac{1}{1+c}$, respectively. If Agent 1 is inspected and found adhering in period 1, he is reward with no inspection in period 2. Since $f(p^*) = 1$, Agent 1 is better off adhering in period 1 (see Proposition 1). If Agent 2 is inspected and found adhering in

[38]When the principal has a commitment power, in all cases that we study in this paper, we can show that private monitoring is weakly better than public monitoring. When the principal does not have a commitment power, it is not difficult to construct examples where private monitoring is strictly worse than public monitoring (see Section 3.2).

period 1, the inspection intensity for him in period 2 is 1. Since $f(\frac{1}{1+c}) = 0$, no future reward is required, and Agent 2 also adheres in period 1. If no agent is inspected in the first period, the principal inspects Agent 1 with probability 1 in the second period.

This strategy profile forms an equilibrium under public monitoring. In this equilibrium, if Agent 1 is not inspected in period 1, his second period action depends on the public history: Agent 1 adheres if no agent is inspected, and Agent 1 violates if Agent 2 is inspected. Yet, under private monitoring, this behavior is not a valid strategy profile: Agent 1 cannot distinguish between the histories "no agent is inspected" and "Agent 2 is inspected", and hence his actions under these two histories must be the same.

This example shows that under public monitoring, the agent's future behavior can be better attuned with the true history. This flexibility may benefit the principal, as it enables a more efficient use of the inspection resource. Indeed, under private monitoring, because each agent's private history is coarse, if at the private history $h_1 = N_1$ the expected inspection intensity available for Agent 1 is low, then Agent 1 violates in all histories that are consistent with h_1 (that is, $\emptyset$ and I_2). Whereas under public monitoring, there is more flexibility in adjusting agent's behavior according to the true history, and it is possible to induce Agent 1 to adhere at $\emptyset$ and to violate only at I_2.

Note that the above effect is absent in Fuchs (2007), where monetary transfer between the principal and the agent is allowed, and the goal is to induce full compliance from the agent (with minimum reward). As a result, the key issue in Fuchs (2007) is to relax the incentive compatibility constraint of the agent. Private monitoring is better in this aspect, since incentive compatibility has to be satisfied in expectation rather than state by state. In our model, since monetary transfer is not allowed and there is limited inspection resource, violation is unavoidable. Hence, when designing the optimal scheme, a new driving force takes effect: In addition to considering the resource constraint, how to optimally give a free pass to the agents constitutes another important concern. Since public monitoring is more flexible in allocating violations, this gives the principal a potential advantage.

A.3 Proof of Theorem 2

Step 1: Optimal PPE under public monitoring.

We start by identifying the optimal inspection scheme under public monitoring. We will define a strategy profile σ^*, under which $\lfloor(1+c+\delta)m\rfloor$ agents adhere in the first period, and $\lfloor(1+c)m\rfloor$ agents adhere in the second period. We then show that it constitutes a PPE, and argue that it is optimal for the principal under public monitoring. Let σ^* be the following strategy profile (see Figure 10).

Denote

$$b = m - \frac{1}{1+c}\left(\lfloor(1+c+\delta)m\rfloor - m\right),$$

which is non-negative since $\delta < 1$. We will divide the n agents into three groups: $\lfloor(1+c+\delta)m\rfloor$ *1-disciplined* agents, $\lfloor(1+c)b\rfloor$ *2-disciplined agents*, and $n - \lfloor(1+c+\delta)m\rfloor - \lfloor(1+c)b\rfloor$ *dummy* agents. In the first period, exactly m 1-disciplined agents are inspected and all 1-disciplined agents adhere, while none of the other agents (2-disciplined or dummy) is inspected and they all violate. The subset of m 1-disciplined agents who are inspected is selected uniformly among all subsets of 1-disciplined agents of size m, hence the probability that each 1-disciplined agent is inspected is $\frac{m}{\lfloor(1+c+\delta)m\rfloor}$.

If some 1-disciplined agents who are inspected in the first period are found violating, then the one who has the minimal index is punished, by being inspected in the second period with probability 1. The other agents are not inspected in the second period and they violate. Since under σ^* all 1-disciplined agents adhere in the first period, the eventuality that is described in this paragraph occurs with probability 0 under σ^*. We next describe the behavior of σ^* in the second period, assuming all 1-disciplined agents that were inspected in the first period adhered.

In the second period, each 1-disciplined agent who was inspected in the first period (and found adhering) is not inspected, while each 1-disciplined agent who was not inspected in the first period is inspected with probability $\frac{m}{\lfloor(1+c)m\rfloor}$. In addition, each 2-disciplined agent is inspected with probability $\frac{m}{\lfloor(1+c)m\rfloor}$ (and they all adhere), and each dummy agent is inspected with probability 0 (and they all violate).

	number	Prob. of inspection	behavior		number	Prob. of inspection	behavior
1-disciplined	$\lfloor(1+c+\delta)m\rfloor$	$\frac{m}{\lfloor(1+c+\delta)m\rfloor}$	A	1-disc. Inspected	m	0	V
				1-disc. uninspected	$\lfloor(1+c+\delta)m\rfloor - m$	$\frac{m}{\lfloor(1+c)m\rfloor}$	A
2-disciplined	$\lfloor(1+c)b\rfloor$	0	V	2-disciplined	$\lfloor(1+c)b\rfloor$	$\frac{m}{\lfloor(1+c)m\rfloor}$	A
dummy	...	0	V	dummy	...	0	V
Period 1				Period 2			

Figure 10: Structure of σ^*.

The inspection strategy σ_0^* satisfies the resource constraint. Indeed, in the first period exactly m (1-disciplined) agents are inspected. In the second period, it can be verified that the sum of the 1-disciplined uninspected agent and 2 disciplined agent is $\lfloor (1+c)m \rfloor$, and hence the resource constraint is satisfied.

We next argue that σ^* is a PPE. Indeed, under σ^*, all agents who violate are inspected with probability 0. In period 1, each 1-disciplined agent is inspected with probability $\frac{m}{\lfloor (1+c+\delta)m \rfloor} \geq \frac{1}{1+c+\delta} = f^{-1}(1)$, and a successfully passed inspection leads to a free violation in period 2, therefore each 1-disciplined agent adheres. In period 2, each agent who adheres is inspected with probability $\frac{m}{\lfloor (1+c)m \rfloor} > \frac{1}{1+c}$ and they adhere. The number of agents who adhere in the first period is $\lfloor (1+c+\delta)m \rfloor$ and the number of agents who adhere in the second period is $\lfloor (1+c)m \rfloor$.

We can now show that the PPE σ^* is optimal for the principal under public monitoring. In the second period, an agent adheres if and only if the inspection probability he faces is at least $\frac{1}{1+c}$. Therefore, under any PPE, the number of agents who adhere in the second period does not exceed $\lfloor (1+c)m \rfloor$. We next turn to the first period. In the first period, the maximum reward for a successful inspection is 1 (a free violation in the second period), and hence to deter an agent from violating in the first period, he has to be inspected with probability no less than $f^{-1}(1) = \frac{1}{1+c+\delta}$. Therefore, in any PPE, the number of agents who adhere in the first period does not exceed $\lfloor (1+c+\delta)m \rfloor$. We conclude that under public monitoring, the PPE σ^* is optimal.

Step 2: Proof of part (i) of Theorem 2.

We here show that private monitoring can yield an outcome superior to σ^*. Suppose that (c, δ) satisfy $\lfloor (1+c+\delta)m \rfloor < (1+c+\delta)m$ and $\lfloor (1+c)m \rfloor < (1+c)m$. These two conditions hold in particular when c and $c+\delta$ are irrational, hence they hold for almost all (c, δ).

Consider a strategy profile σ' under private monitoring that is identical to σ^*, except that in the first period the principal assigns a positive probability ϵ to no inspection. If no inspection is conducted in the first period, in the second period each agent is inspected with probability 0 (this inspection probability is not important). An argument similar to the one in the proof of Theorem A.1 can be applied to show that σ' forms an SE, provided ϵ is sufficiently small.

Since $\lfloor (1+c+\delta)m \rfloor < (1+c+\delta)m$, a slight decrease in the inspection intensity in the first period does not affect the agents' first-period IC constraint, hence under σ' the 1-disciplined agents adhere in the first period. If a 1-disciplined agent is not inspected in the first period, then in the second period he assigns a probability close to one to the event that the continuation play is the one indicated by σ^* where he faces inspection intensity $\frac{m}{\lfloor (1+c)m \rfloor} > \frac{1}{1+c}$, and a probability close to zero to the event that no agent was inspected in the first period in which case he faces an inspection intensity 0. When ϵ is sufficiently small, this agent is better off adhering in the second period. A similar argument shows that each 2-disciplined agent adheres in the second period when ϵ is sufficiently small. Consequently, under σ', in the first period $\lfloor(1+c+\delta)m\rfloor$ agents adhere, as in σ^*. In the second period, with probability $1-\epsilon$ the number of agents who adhere is $\lfloor(1+c)m\rfloor$, and with probability ϵ the number of agents who adhere is $\lfloor(1+c)m\rfloor + m > \lfloor(1+c)m\rfloor$. This improves the principal's payoff.

Step 3: Proof of part (ii) of Theorem 2.

First, we claim that, regardless of the monitoring structure, the expected number of adhering agents in the second period is at most $(1+c)m$. To this end, define the following random variables for each agent i:

$$X_i = \text{agent } i \text{ is inspected in period 2,}$$

$$Y_i = \text{agent } i \text{ adheres in period 2.}$$

Fix some PPE. Because of the budget constraint,

$$\sum_{i=1}^{n} X_i \leq m. \tag{8}$$

Denote by $\mathscr{F}_i$ agent i's information at period 2. Then

$$Y_i = 1 \text{ if } P(X_i = 1|\mathscr{F}_i) > \frac{1}{1+c},$$

$$Y_i = 0 \text{ if } P(X_i = 1|\mathscr{F}_i) < \frac{1}{1+c},$$

$$Y_i = \{0,1\} \text{ if } P(X_i = 1|\mathscr{F}_i) = \frac{1}{1+c}.$$

In particular,

$$Y_i \leq (1+c)P(X_i = 1|\mathscr{F}_i). \tag{9}$$

Therefore,

$$\begin{aligned}\sum_{i=1}^{n} E\left[Y_i\right] &\leq (1+c)\sum_{i=1}^{n} E\left[P(X_i=1|\mathscr{F}_i)\right]\\ &= (1+c)\sum_{i=1}^{n} E\left[E[X_i|\mathscr{F}_i]\right]\\ &= (1+c)\sum_{i=1}^{n} E[X_i]\\ &\leq (1+c)m.\end{aligned} \tag{10}$$

In Eq. (10), the first inequality follows from Eq. (9); the first equality holds since $(X_i)_{i=1}^n$ are Bernoulli random variables; the second equality holds because of the law of iterated expectations; and the second inequality follows from Eq. (8).

Next we turn to the first period. Since in the two-period game the reward for a successfully past inspection is at most 1 (i.e., a free violation in the second period), the minimum inspection intensity that can induce adherence in the first period is $p^* = f^{-1}(1) = \frac{1}{1+c+\delta}$. Therefore, the number of agents who adhere in the first period is at most $\lfloor(1+c+\delta)m\rfloor$. Hence, the upper bound of adherence in any monitoring structure is $\delta(1+c)m+\lfloor(1+c+\delta)m\rfloor$. As shown in Step 1, the optimal adherence in public monitoring is $\delta\lfloor(1+c)m\rfloor+\lfloor(1+c+\delta)m\rfloor$. The difference between the optimal PPE outcome and the upper bound of adherence in any monitoring structure is at most δ. Part (ii) of Theorem 2 follows.

A.4 A formal definition of the non-atomic game

In Section 2 we defined the game with finitely many inspectors and agents. Here we formally define the game Γ with a continuum of agents and inspectors. Repeated games with a continuum of players have been studied, e.g., by Sabourian (1990), and Massó (1993).

The players in Γ are the principal and the continuum of agents who are represented by the interval $[0,1]$. An action of the principal is the set of inspected agents. Therefore, the action set of the principal is the collection $\mathcal{F}_\alpha$ of all Borel measurable subsets of $[0,1]$ whose Lebesgue measure is at most α. The set of feasible actions for each agent is $\{A,V\}$. The set of feasible joint actions of the agents is the set $\mathcal{L}$ of measurable functions from $[0,1]$ to $\{A,V\}$. To allow for mixed strategies, we need to endow $\mathcal{F}_\alpha$ and $\mathcal{L}$ with sigma-algebras.

To this end, we consider the set $\mathcal{F}$ of all Borel measurable subsets of $[0,1]$. The set $\mathcal{F}_\alpha$ is a subset of $\mathcal{F}$ and $\mathcal{L}$ is equivalent to $\mathcal{F}$ (an element $f\in\mathcal{L}$ corresponds to the set of agents who adhere under f). It is therefore sufficient to define a sigma-algebra on $\mathcal{F}$, and equip $\mathcal{F}_\alpha$ and $\mathcal{L}$ with the induced sigma-algebras. Consider then the sigma-algebra on $\mathcal{F}$ generated by all sets of the form

$$\{G \subseteq [0,1] \colon G \text{ is Borel measurable and } \mu(G) \geq c\},$$

where μ ranges over all measures on $[0,1]$ and c ranges over $[0,1]$.

A *pure strategy* for the principal is a set $D_1 \in \mathcal{F}_\alpha$ (the set of agents who are inspected in the first period) together with a measurable function $d_2 : \mathcal{F}_\alpha \times \mathcal{L} \mapsto \mathcal{F}_\alpha$ (that indicates the set of agents who are inspected in the second period). A pure strategy profile for the agents is a function $a_1 \in \mathcal{L}$ (the agents' behavior in the first period), together with a measurable function $a_2 : \mathcal{F}_\alpha \times \mathcal{L} \mapsto \mathcal{L}$ (the agents' behavior in the second period). The strategies of the principal and the agents must depend only on the information they possess. Therefore, we require that $d_2(D_1, a_1) = d_2(D_1, a_1')$ whenever $a_1 = a_1'$ on D_1: the inspection strategy in the second period cannot depend on the first period behavior of agents that were not inspected in the first period. For the agents, the condition imposed on strategy profiles depends on the monitoring structure. When monitoring is public, we require that $a_2(t; D_1, a_1) = a_2(t; D_1, a_1')$ for every $t \in [0,1]$ such that $a_1 = a_1'$ on $D_1 \cup \{t\}$: agent t's choice in the second period can depend on the first period's choices of the principal, the inspected agents, and agent t himself (even if he is not inspected in the first period). When monitoring is private, we require that $a_2(t; D_1, a_1) = a_2(t; D_1', a_1')$ for every $t \in [0,1]$ such that $a_1(t) = a_1'(t)$ and $t \in D_1$ if and only if $t \in D_1'$: agent t's choice in the second period depends on his action in the first period and on whether he was inspected in the first period.

The payoff function of each agent is similar to his payoff function in the game with finitely many agents and inspectors, while the payoff function of the principal is the integral of the function in Eq. (2) w.r.t. the Lebesgue measure.

In the spirit of Aumman (1964), a *mixed strategy* is a measurable function that assigns a pure strategy to each number in $[0,1]$, with the interpretation that the real number in $[0,1]$ is selected according to the uniform distribution, and then the player executes the corresponding pure strategy. Since we did not define pure strategies for single players, but rather strategy profiles, this definition gives rise to correlated strategies for the agents. We will study equilibria where the principal uses a mixed strategy and the agents use a pure strategy profile, hence they do not randomize and the issue of correlation does not arise.

In the game with public monitoring, the concept of PPE is defined in the standard way. The definition of SE in the game with private monitoring involves the concept of assessment and the convergence of a sequence of mixed strategies. Since the set of players has the cardinality of the continuum, this definition is more complex than the definition in finite games. Since the extension of the concept of SE to games with a continuum of players is not the focus of the paper, we will study Nash equilibria in the game with private monitoring.

In our construction of strategies for the principal, we will often use the following, for $0 < b < a$: given a measurable set G of agents of size a, each agent in G is selected with probability b. This verbal description translates into the following formal construction: Let $\varphi : [0, a] \to G$ be a measurable measure-preserving bijection. The principal uniformly selects a number $x \in [0, a]$, and then inspects all agents $\{\varphi(t)\colon t \in [x, x + b]\}$, where addition is modulo a.

A.5 Proof of Theorem 3

The non-atomic game is formally defined in Appendix A.4.

First consider the case $1 \geq (2+c)\alpha$. Under public monitoring, a proof similar to the one in Section A.3 applies, and the optimal PPE outcome yields the principal the upper bound of adherence $\delta(1+c)\alpha + (1+c+\delta)\alpha$. The same strategy profile also constitutes a Nash equilibrium under private monitoring, and hence Private $\approx$ Public in this case.

If $1 \leq (1+c)\alpha$, then inspecting each agent with probability $\frac{1}{1+c}$ in each period completely deters violations, and is optimal under both monitoring schemes.

We finally focus on the case $(1+c)\alpha < 1 < (2+c)\alpha$. Given the strategy profile of the agents, denote by G_a the group of agents who adhere in period 1, and by G_v the group of agents who violates in period 1. Let $|G_a|$ and $|G_v|$ be the Lebesgue measure of the two groups, respectively. Necessarily, $|G_a| + |G_v| = 1$.

For each agent $i \in G_a$, denote by $p(i)$ the inspection probability for agent i in period 1. Because of the resource constraint, $\int_{i \in G_a} p(i)\ \mathrm{d}i \leq \alpha$.

Under both monitoring schemes, each agent who adheres in the first period (that is, agents in G_a) has to obtain a continuation payoff of no less than $f\big(p(i)\big)$ if he is inspected in period 1. This implies that the total discounted number of violations is no less than $1 - (1+c)\alpha$:

$$\begin{aligned}
&\overbrace{|G_v|}^{\text{first period violations}} \quad + \quad \overbrace{\int_{i \in G_a} \delta \cdot p(i) \cdot f\big(p(i)\big) di}^{\text{lower bound on the second period violations}} \\
&\geq |G_v| + \int_{i \in G_a} \big[1 - (1+c)p(i)\big] di \\
&= |G_v| + |G_a| - (1+c) \int_{i \in G_a} p(i) di \\
&\geq 1 - (1+c)\alpha.
\end{aligned} \tag{11}$$

This lower bound is not the same as the one obtained in Step 3 in Appendix A.3. In particular, the strategy profile σ^* characterized in Step 1 in Appendix A.3 cannot be implemented as a PPE in the current case. This is because under σ^*, the second-period compliance level is $(2+c)\alpha$, which is larger than the size of the population in the current case, since $1 < (2+c)\alpha$. Therefore, if the principal minimizes the first-period violations (that is, she inspects as many agents as possible with probability $\frac{1}{1+c+\delta}$), then she has to reward too many agents in the second period, and hence some inspection resources are wasted in period 2. To better use the second-period resources, the principal treats some agents myopically by inspecting them with a high probability $\frac{1}{1+c}$ already in period 1. This way, fewer agents adhere in period 1, but more agents adhere in period 2.

We next construct a PPE under public monitoring that yields the lower bound $1-(1+c)\alpha$. We then argue that this outcome can also be implemented by a Nash equilibrium under private monitoring. Note that if the inspector uses a myopic strategy, the minimum loss in each period is $1-(1+c)\alpha$, which yields a total loss of $(1+\delta)\cdot\big(1-(1+c)\alpha\big)$.

It turns out that the following strategy profile $\sigma_{\text{inf}}^{\text{pub}}$ is optimal under public monitoring. Let $t_1 = (1+c+\delta)\big(1-(1+c)\alpha\big)$, and $t_2 = t_1 + (1+c)\big(\alpha - [1-(1+c)\alpha]\big)$. Note that t_2 is smaller than 1 because $(1+c)\alpha < 1$, and it is larger than t_1 because $(2+c)\alpha > 1$. Agents are divided into three groups: $G_1 = [0, t_1]$, $G_2 = [t_1, t_2]$, and $G_3 = [t_2, 1]$.

- Strategy of the principal:
  - In the first period
    * Agents in G_1 are inspected with probability $\frac{1}{1+c+\delta}$.
    * Agents in G_2 are inspected with probability $\frac{1}{1+c}$.
    * Agents in G_3 are inspected with probability 0.
  - In the second period
    * Agents in G_1 who were inspected and found adhering in period 1 are inspected with probability 0.
    * All other agents are inspected with probability $\frac{1}{1+c}$, irrespective of the history.
    * Punishment: If some agents in G_1 are insepcted in the first period are found violating, then the one who has the minimal index is punished by being inspected in the second period with probability 1. The other agents are not inspected in the second period.[39]

[39]Since under $\sigma_{\text{inf}}^{\text{pub}}$ all agents in G_1 adhere in the first period, the eventuality that is described in this paragraph occurs with probability 0 under $\sigma_{\text{inf}}^{\text{pub}}$.

- Strategy of the agents:
  - In the first period, agents in G_1 and G_2 adhere. Agents in G_3 violate.
  - In the second period, agents in G_1 who were inspected in period 1 violates. All other agents adhere.

The inspection strategy $\sigma^{\text{pub}}_{\text{inf},0}$ satisfies the budget constraint. Indeed, in period 1, the required inspection resource is

$$\frac{1}{1+c+\delta} \cdot |G_1| + \frac{1}{1+c} \cdot |G_2| = \alpha.$$

The size of the set of agents in G_1 who are inspected in period 1 is $\frac{1}{1+c+\delta} \cdot |G_1| = 1-(1+c)\alpha$. Therefore, in period 2, the size of the set of agents who are inspected with probability zero is $1-(1+c)\alpha$, and the set of remaining agents has the size $(1+c)\alpha$. The inspection resource required in period 2 is $\frac{1}{1+c} \cdot (1+c)\alpha = \alpha$. Thus, the budget constraint is met in both periods.

We now verify that the agents' strategies in $\sigma^{\text{pub}}_{\text{inf}}$ are best response to the principal's strategy $\sigma^{\text{pub}}_{\text{inf},0}$. This is clear for agents in G_2 and G_3 (in both periods), and for agents in G_1 (in the second period). Regarding agents in G_1 in the first period, since their continuation payoff is 1 upon a successfully passed inspection, and since they are inspected in period 1 with probability $\frac{1}{1+c+\delta} = f^{-1}(1)$, adhering is a best response for them in period 1.

We finally compute the number of violations under $\sigma^{\text{pub}}_{\text{inf}}$. It can be verified that

$$L_0(\sigma^{\text{pub}}_{\text{inf}}) = |G_3| + \delta \cdot \big(1-(1+c)\alpha\big) = 1-\alpha(1+c).$$

As shown in Eq. (11), this is the lower bound on the amount of violations in equilibrium, and hence $\sigma^{\text{pub}}_{\text{inf}}$ is optimal.

Note that agents' strategies in $\sigma^{\text{pub}}_{\text{inf}}$ depend only on (i) which group the agent belongs to, and (ii) the agent's own inspection history. Therefore, the same strategy profile constitutes a Nash equilibrium under private monitoring, and it is optimal as well.

A.6 Proof of Theorem 4

We first illustrate the superiority of private monitoring in this case with an example. Take $c = 0.96$ and $\delta = 0.9$, so that the requirement in Proposition 4 is satisfied. In the one-shot game, an agent adheres if and only if he is inspected with probability at least $\frac{1}{1+c} = 0.51$. Under public monitoring, in the optimal PPE, both agents adhere in the first period, and one of the agents violates in the second period (the identity of the violating agent depends on the first-period history).

We now construct an equilibrium $\sigma_{\text{nc}}^{\text{pri}}$ under private monitoring that yields the principal a better payoff (see Figure 11). In this equilibrium, both agents adhere in the first period. In the second period, the agents' expected payoffs are $(v_1 = \frac{1}{3}, v_2 = \frac{1}{3})$, regardless of the principal's first-period action.

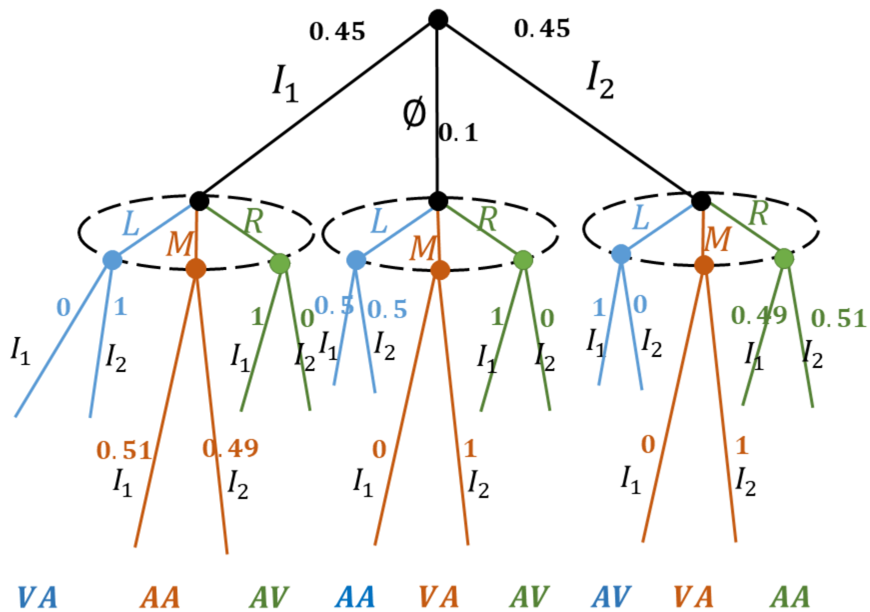

Figure 11: Strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ under private monitoring with no-commitment (example).

- In the first period inspect each agent with probability 0.45. Both agents adhere.
- A public correlation device is used at the beginning of the second period. With probability $\frac{1}{3}$ the correlation device shows L (left), with probability $\frac{1}{3}$ it shows M (middle), and with probability $\frac{1}{3}$ it shows R (right).
  - Suppose Agent 1 is inspected in the first period and found adhering.
    - If the correlation device shows L, then Agent 1 (resp., Agent 2) is inspected with probability 0 (resp., probability 1). Agent 1 violates and Agent 2 adheres.
    - If the correlation device shows M, then Agent 1 (resp., Agent 2) is inspected with probability 0.51 (resp., probability 0.49). Both agents adhere.
    - If the correlation device shows R, then Agent 1 (resp., Agent 2) is inspected with probability 1 (resp., probability 0). Agent 1 adheres and Agent 2 violates.

The behavior of the inspector and agents in the remaining cases is as displayed in Figure 11.

- Punishment: if an agent is found violating in the first period, then in the second period he is inspected with probability 1 and the other agent is inspected with probability 0. Agent i adheres and the other agent violates.

Note that Agent 1 cannot distinguish between the history "no agent is inspected" and "Agent 2 is inspected". Therefore, if, say, Agent 1 is not inspected in the first period and the public correlation device shows R, then Agent 1 assigns probability $\frac{0.1}{0.1+0.45}$ that he will be inspected with probability 1, and with the remaining probability he will be inspected with probability 0.49. This yields an expected inspection probability of 0.58 to Agent 1, and he is better off adhering. By a similar argument, together with the observation that $f(0.45) < \frac{1}{3}$, one can verify that under the strategy profile $\sigma_{\text{nc}}^{\text{pri}}$, the agents are taking best responses to the inspection strategy.

Since the principal has no commitment power, it remains to verify that the principal has no incentive to deviate. Since inspection is costless, the principal's payoff depends only on the agents' actions. In the second period the principal cannot benefit from deviation, since her deviation does not change the actions of the agents. Regardless of the principal's first period action, both agents adhere in the first period, and their second-period payoffs are $(v_1 = \frac{1}{3}, v_2 = \frac{1}{3})$, which yields the principal a loss $\frac{2}{3}$ in the second period. Therefore, the principal has no incentive to deviate also in the first period. This equilibrium is superior to the optimal PPE under public monitoring, as claimed.

The formal proof of Proposition 4 follows the example. Consider the strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ shown in Figure 12, where $\widehat{p} := f^{-1}(\frac{1}{3})$. It can be verified that if $\delta > (1-c) \cdot \frac{3(1+c)}{c}$, then $\widehat{p} < \frac{1}{2}$. The interpretation of the figure is the same as the one in the example. The differences from the example are: (1) in the first period, each agent is inspected with probability $\widehat{p} = f^{-1}(\frac{1}{3})$ and both agents adhere; (2) if Agent 1 is inspected in the first period and the public correlation device shows M, then Agent 1 is inspected with probability $\frac{1}{1+c}$, and Agent 2 is inspected with probability $\frac{c}{1+c}$; (3) if Agent 2 is inspected in the first period and the public correlation device shows R, then Agent 1 is inspected with probability $\frac{c}{1+c}$ and Agent 2 is inspected with probability $\frac{1}{1+c}$. The agents' actions under each history are the same as in the example.

Since in the first period both agents adhere, and in the second period both agents adhere with a positive probability, the strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ yields the principal a better payoff than the optimal PPE under public monitoring. We next verify that it is an equilibrium.

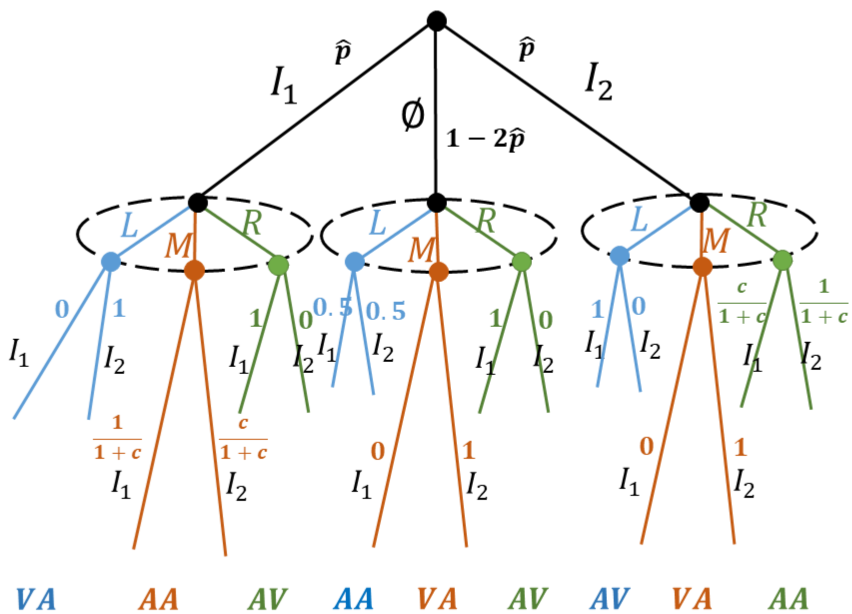

Figure 12: Strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ under private monitoring with no-commitment.

For the strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ to be an SE, we first study the conditions under which the agents have no incentive to deviate. For the second period, these conditions are

$$\frac{\widehat{p}}{1-\widehat{p}} \cdot 1 + \frac{1-2\widehat{p}}{1-\widehat{p}} \cdot 0.5 \geq \frac{1}{1+c}, \tag{12}$$

and

$$\frac{\widehat{p}}{1-\widehat{p}} \cdot \frac{c}{1+c} + \frac{1-2\widehat{p}}{1-\widehat{p}} \cdot 1 \geq \frac{1}{1+c}. \tag{13}$$

Eq. (12) always holds, and Eq. (13) holds if $\delta > (1-c) \cdot \frac{3(1+c)}{c}$, as assumed in the proposition.

Since $f(\widehat{p}) = 1$, in the first period both agents are better off adhering. We now verify that the principal cannot benefit from deviating. Since no matter what the action of the principal in the first period is, the agents' second period payoffs are always $(v_1 = \frac{1}{3}, v_2 = \frac{1}{3})$, the principal cannot benefit from deviation. In the second period, since the principal's action does not affect her own payoff, deviation is also non-profitable, as desired. This completes the proof that when $\delta > (1-c) \cdot \frac{3(1+c)}{c}$, the strategy profile $\sigma_{\text{nc}}^{\text{pri}}$ shown in Figure 12 is an SE.

A.7 Proof of Theorem 5

Case 1: $\alpha \geq \frac{1}{1+c}$. In this case there are sufficient resources to inspect all agents with probability $\frac{1}{1+c}$, and thereby deter violations in both periods regardless of the monitoring structure. In particular, Public $\approx$ Private.

Case 2: $\alpha \leq \frac{1}{2+2c+\delta}$. By an argument similar to the one in Step 3 in Appendix A.3, the expected number of adherences in the second period is at most $(1+c)\alpha$, and the number of adherences in the first period is at most $(1+c+\delta)\alpha$. That is, the upper bound of adherence in any monitoring structure is $\delta(1+c)\alpha + (1+c+\delta)\alpha$.

We next construct for the model with public monitoring a PPE that yields the principal the upper bound of adherences, and then show that a similar strategy profile constitutes a Nash equilibrium under private monitoring. Note that since the principal has no commitment power, the strategy profile σ^* characterized in Step 1 in Appendix A.3 is no longer an equilibrium under private monitoring: In σ^*, each 1-disciplined agent who was inspected in period 1 violates in period 2, and each 1-disciplined agent who was not inspected in period 1 adheres in period 2. Therefore, the principal benefits by deviating and inspecting no 1-disciplined agent in the first period. To fix this problem, we need to modify the inspection scheme so that the principal's second-period payoff is independent of her first-period action.

Let σ^*_{nc} be the following strategy profile: Denote $t_1 = (1+c+\delta)\alpha$ and $t_2 = (1+c+\delta)\alpha + (1+c)\alpha$. Divide the population $[0,1]$ into three groups: agents in $[0,t_1]$ are *1-disciplined agents*, agents in $[t_1,t_2]$ are *2-disciplined agents*, and agents in $[t_2,1]$ are *dummy agents*. The assumption $1 \geq (2+2c+\delta)\alpha$ guarantees that the upper bound of the 2-disciplined agent, t_2, is at most 1.

In the first period, all 1-disciplined agents are inspected with probability $\frac{1}{1+c+\delta}$ and they all adhere, while each of the other agents is inspected with probability 0 and they all violate.

As in Step 1 in Appendix A.3, if a 1-disciplined agent is inspected and found violating, then he is punished by being inspected in the second period with probability 1. This situation does not occur in equilibrium. We next describe the behavior of σ^*_{nc} in the second period, assuming all 1-disciplined agents adhere in the first period.

In the second period, each 1-disciplined agent is inspected with probability zero (and they all violate), *regardless of whether he is inspected in the first period.* Each 2-disciplined agent is inspected with probability $\frac{1}{1+c}$ (and they all adhere). Each dummy agent is inspected with probability 0 (and they all violate). By an argument similar to the one in Step 1 in Appendix A.3, σ^*_{nc} forms a PPE under public monitoring. One can verify that this PPE yields the principal the upper bound of adherence.

In this construction, the principal cannot change the number of violations by deviating from her prescribed strategy $\sigma^*_{nc,0}$. Since in equilibrium the inspection intensity for each agent depends only on which group he belongs to, the same strategy profile constitutes a Nash equilibrium under private monitoring. Therefore, Public $\approx$ Private in this case.

Case 3: $\frac{1}{2+2c+\delta} < \alpha < \frac{1}{2+c}$. Like in the previous case, the expected number of adherences in the second period is at most $(1+c)\alpha$, and the number of adherences in the first period is at most $(1+c+\delta)\alpha$.

Consider the non-atomic public monitoring version of the strategy profile that is similar to σ^* in Step 1 of Section A.3. The only difference is that to prevent the principal from deviating, if the agents observe that in the first period the total mass of the agents who are inspected is not equal to m, then they punish the principal by violating in period 2 (as a best response, the principal inspects no agent in this eventuality). By an argument similar to the one in Step 1 of Section A.3, it can be verified that the modified version of σ^* is a PPE of the no-commitment problem.

We next argue that under private monitoring, the lower-bound violations cannot be supported as an equilibrium. Suppose $\sigma^{\text{pri}}_{\text{nc}}$ is the optimal equilibrium under private monitoring. Let $\mathscr{A}$ be the set of agents who adhere in period 1 under $\sigma^{\text{pri}}_{\text{nc}}$. In the second period, it is without loss of generality to assume that an agent who violates is not inspected. This is because an inspected violation hurts the agent without benefiting the principal. Under $\sigma^{\text{pri}}_{\text{nc}}$, the continuation payoff of any agent if he is *not inspected* in period 1 (denoted v_{ni}) must be no less than his continuation payoff if he is *inspected and found adhering* in period 1 (denoted v_a). Indeed, if $v_{ni} < v_a$, then when the agent is not inspected in period 1, he violates less in period 2. The principal, whose goal is to minimize violations, benefits from reducing the inspection probability on this agent. Note that under private monitoring, this deviation of the principal cannot be observed by agents.

To generate the lower-bound violations, necessarily a mass of $(1+c+\delta)\alpha$ agents adhere in period 1. Denote the set of these agents by $\mathscr{A}$. Attaining the lower bound can be done only by inspecting each of the agents in $\mathscr{A}$ with probability $\frac{1}{1+c+\delta}$ in period 1,[40] and rewarding a successfully passed inspection with a free violation in the future. This implies that for any agent in $\mathscr{A}$, the continuation payoff $v_a \geq f(\frac{1}{1+c+\delta}) = 1$. By the argument above, $v_{ni} \geq v_a \geq 1$. Since the continuation payoff is bounded above by 1, we have $v_a = v_{ni} = 1$. Therefore, the second-period violation is at least $(1+c+\delta)\alpha$. This implies that the second period adherence is at most $1-(1+c+\delta)\alpha$, which is strictly lower than $(1+c)\alpha$ since

[40]If the inspection probability for any agent in $\mathscr{A}$ is lower than $\frac{1}{1+c+\delta}$, then necessarily this agent violates in period 1. If the probability is higher than $\frac{1}{1+c+\delta}$, then because of the resource constraint, some of the other agents in $\mathscr{A}$ must be inspected with probability lower than $\frac{1}{1+c+\delta}$ and violate.

$\alpha > \frac{1}{2+2c+\delta}$. Consequently, the lower bound on the number of violations cannot be attained under private monitoring.

Case 4: $\frac{1}{2+c} < \alpha < \frac{1}{1+c}$. By an argument similar to the one in Appendix A.5, the lower bound on the number of violations is $1-(1+c)\alpha > 0$ in both monitoring structures.[41] Moreover, under public monitoring, this lower bound can be attained as an equilibrium by the strategy profile $\sigma^{\text{pub}}_{\text{inf}}$ characterized in Appendix A.5 with slight modifications to prevent the principal from deviating: As long as $1-(1+c)\alpha$ agents in G_1 and $\alpha - [1-(1+c)\alpha]$ agents in G_2 are inspected in period 1, agents follow the strategy prescribed in $\sigma^{\text{pub}}_{\text{inf}}$ in the second period. Otherwise, each agent is inspected with probability zero and he violates in period 2. This modification guarantees that the principal has no profitable deviations.

It remains to verify that the lower bound on the number of violations cannot be attained as an equilibrium outcome under private monitoring.

Denote by G_a the set of agents who adhere in period 1, and by G_v the set of agents who violate in period 1. Let $|G_a|$ and $|G_v|$ be the Lebesgue measure of the two sets, respectively. Necessarily, $|G_a| + |G_v| = 1$.

Under the optimal Nash equilibrium, for each agent $i \in G_a$, denote by $p(i)$ the inspection probability for agent i in period 1. Because of the resource constraint, $\int_{i\in G_a} p(i) \leq \alpha$.

For agent i in G_a to adhere in period 1, upon being inspected and found adhering, this agent's continuation payoff has to be at least $f\big(p(i)\big)$. Moreover, as argued in the previous case, since the principal has no commitment power, to form an equilibrium under private monitoring, this agent's continuation payoff upon no-inspection has to be at least $f\big(p(i)\big)$ as well.

This implies that the expected number of violations under a Nash equilibrium is not smaller than

$$\begin{aligned}
&\overbrace{|G_v|}^{\text{first period violations}} + \overbrace{\int_{i\in G_a} \delta \cdot f\big(p(i)\big)di}^{\text{lower bound on the second period violations}} \\
&> |G_v| + \int_{i\in G_a} \delta \cdot p(i) \cdot f\big(p(i)\big)di \\
&\geq 1-(1+c)\alpha,
\end{aligned} \tag{14}$$

where the second inequality in Eq. (14) follows from Eq. (11). That is, private monitoring cannot attain the lower-bound payoff $1-(1+c)\alpha$ in the current case. This completes the

[41] In Case 3, the number of violations $1-(1+c)\alpha$ is lower than the number of violations in σ^*_{nc} and cannot be attained in equilibrium.

proof of Theorem 5.

A.8 Linear inspection cost

In this section we discuss the case in which the inspection cost is linear in the number of inspections, and we argue that for an inspection problem that involves $n \geq 2$ agents, $m \geq 1$ inspectors, and $T \geq 2$ periods, both public monitoring and private monitoring yield the principal the same optimal payoff. Indeed, since inspection costs are linear, under both monitoring structures the principal is better off treating each agent independently, and hence the monitoring structure does not play a role in improving the principal's payoff. We next provide the detailed argument.

We first argue that the principal is better off treating each agent independently under public monitoring. Let σ^* be the optimal PPE. We turn to calculate the loss to the principal under σ^* that arises due to each agent: the expected loss due to violations of each agent plus the expected loss due to inspections made on each agent. Denote by i the agent who imposes the lowest loss v_i to the principal under σ^*. Since at each history agent i cares only about the inspection probability he faces, and since the cost of inspection is linear, by using public correlation device to mimic the original play path, we can construct a new inspection strategy under which agent i imposes the same loss L_0 and the inspection probability for agent i is independent of the inspection history for other agents. We next use the same inspection strategy to inspect the other $n-1$ agents, so that the principal's total loss is n times the loss due to agent i. This procedure necessarily yields the principal a loss no larger than the loss in σ^*, and hence it is optimal as well. Consequently, one of the principal's optimal inspection strategy is to treat the n agents independently.

We next argue that under private monitoring the principal is better off treating each agent independently. Consider an SE σ^{pri}. Suppose that under σ^{pri}, at a private history h_1 of Agent 1, the expected inspection probability Agent 1 faces is $\overline{p}$. Denote by $\mathscr{I}_1$ the information set containing h_1. Modify the principal's strategy so that at every play $a \in \mathscr{I}_1$, Agent 1 faces the inspection probability $\overline{p}$. Because of the linearity of the inspection cost, this modification does not change the average inspection cost. Suppose that we do the same modification simultaneously at all information sets of all agents, and denote the new inspection strategy by $\widehat{\sigma}_0^{\text{pri}}$. Under $\widehat{\sigma}_0^{\text{pri}}$, the inspection probability each agent faces at every information set is the same as under σ_0^{pri}, and hence his best response remains the same. This implies that $\left(\widehat{\sigma}_0^{\text{pri}}, (\sigma_1^{\text{pri}}, \sigma_2^{\text{pri}}, ..., \sigma_n^{\text{pri}})\right)$ generates the same loss to the principal as σ^{pri}, and with the same inspection costs.

Therefore, without loss of generality we can assume that under an SE of G_T^{pri}, at every information set $\mathscr{I}_i$ of agent i, all plays in $\mathscr{I}_i$ induce the same inspection probability for agent i. In other words, the inspection probability an agent faces in each period depends only on his private history, and is independent of the inspection for the other agent. Since in the single-agent problem the monitoring structure plays no role, it follows that when the inspection cost is linear, the two information structures yield the same expected loss to the principal.

A.9 Proof of Theorem 6

Since $\delta > 1-c$, we have $f(\frac{1}{2}) < 1$. In the two-period game, the optimal PPE under public monitoring, denoted σ^{pub}, is as described below (see Figure 13). In particular, it does not involve the use of the second inspection. This strategy is a variation of the one constructed for the case $r_1 = 0$ and $r_2 = \infty$ (or equivalently, $m = 1$ and $n = 2$), while minimizing the probability of inspection as much as possible (so as to minimize the inspection cost).

Let $p^* := f^{-1}(1)$. Note that $p^* < \frac{1}{2}$.

- In the first period, inspect each agent with probability p^*. Both agents adhere.
- If the agent, denoted agent i, is inspected in the first period and found adhering, then in the second period agent i is inspected with probability 0 and the other agent is inspected with probability $\frac{1}{1+c}$. agent i violates and the other agent adheres.
- If no agent is inspected in the first period, then in the second period Agent 1 is inspected with probability 0 and the other agent is inspected with probability $\frac{1}{1+c}$. agent i violates and the other agent adheres.
- Punishment: if an agent is found violating in the first period, then in the second period he is inspected with probability $\frac{1}{1+c}$ and the other agent is inspected with probability 0. agent i adheres and the other agent violates.

The PPE σ^{pub} is optimal because (1) in the last period it is optimal to let one agent violate and one agent adhere, and (2) it minimizes the inspection costs in the first period under the constraint that each agent is inspected with probability at least $f^{-1}(1)$.

The principal's total loss is

$$L_0(\sigma^{\text{pub}}) = \overbrace{0 + \delta \cdot 1}^{\text{loss from violations}} \quad + \overbrace{2p^* \cdot r_1 + \delta \cdot \frac{1}{1+c} \cdot r_1}^{\text{inspection cost}}. \tag{15}$$

We next consider the following strategy profile under private monitoring, denoted σ^{pri} (see Figure 14).

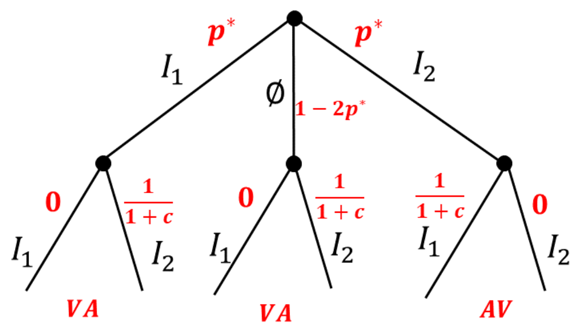

Figure 13: Optimal PPE under public monitoring.

- In the first period, inspect each agent with probability p^*. Both agents adhere.
- If agent i is inspected in the first period and found adhering, then in the second period agent i is inspected with probability 0 and the other agent is inspected with probability q_1. Agent i violates and the other agent adheres.
- If no agent is inspected in the first period, then in the second period each agent is inspected with probability q_2. Both agents adhere.
- Punishment: if an agent is found violating in the first period, then in the second period he is inspected with probability $\frac{1}{1+c}$ and the other agent is inspected with probability 0. agent i adheres and the other agent violates.

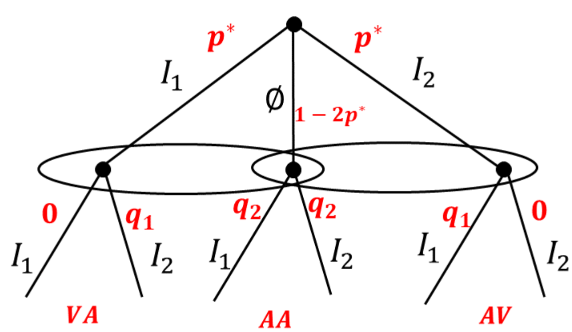

Figure 14: Strategy profile σ^{pri} under private monitoring.

We claim that, if

$$\frac{p^*}{1-p^*} \cdot q_1 + \frac{1-2p^*}{1-p^*} \cdot q_2 = \frac{1}{1+c}, \tag{16}$$

then σ^{pri} is an SE. Indeed, suppose that Eq. (16) holds. If agent i is not inspected in the first period, then he assigns probability $\frac{1}{1+c}$ to the event that he will be inspected in the second period, and hence adhering is the best response. Since $f(p^*) = 1$, in the first period adhering is also a best response.

Next we claim that there exist q_1 and q_2 that satisfy Eq. (16), where $\frac{1}{1+c} < q_1 < 1$ and $0 < q_2 < \frac{1}{2}$. Indeed, when $q_1 = 1$ and $q_2 = \frac{1}{2}$, the left-hand side of Eq. (16) is $\frac{1+c+\delta}{2c+2\delta}$, which is strictly larger than $\frac{1}{1+c}$. When $q_1 = \frac{1}{1+c}$ and $q_2 = 0$, the left-hand side of Eq. (16) is strictly smaller than $\frac{1}{1+c}$. By the continuity of the left-hand side of Eq. (16) with respect to q_1 and q_2, the claim follows. Suppose (q_1^*, q_2^*) is one solution with the desired properties.

We argue that when $q_1 = q_1^*$ and $q_2 = q_2^*$, the SE σ^{pri} yields the principal a loss lower than that in σ^{pub}. Indeed, the principal's loss in this case is

$$\begin{aligned} L_0(\sigma^{\text{pri}}) &= \overbrace{0 + \delta \cdot 2p^*}^{\text{loss from violations}} + \overbrace{2p^* \cdot r_1 + \delta \cdot \big(2p^* \cdot q_1 \cdot r_1 + (1-2p^*) \cdot 2q_2 \cdot r_1\big)}^{\text{inspection cost}} \\ &= \delta \cdot 2p^* + 2p^* \cdot r_1 + \delta \cdot 2r_1 \cdot \frac{1-p^*}{1+c}. \end{aligned} \tag{17}$$

The second equality in Eq. (17) follows from Eq. (16). Comparing Eq. (17) with Eq. (15), we see that

$$L_0(\sigma^{\text{pub}}) - L_0(\sigma^{\text{pri}}) = \delta \cdot (1-2p^*) \cdot \left(1 - \frac{r_1}{1+c}\right). \tag{18}$$

Since $r_1 < 1 + c$ by the divergence condition, $L_0(\sigma^{\text{pub}}) - L_0(\sigma^{\text{pri}}) > 0$, so private monitoring yields the principal a lower loss, as claimed.

A.10 Pairs (r_1, r_2) that violate the divergence condition

If $r_1 \geq 1 + c$, then under both private and public monitoring, in the second period, on the equilibrium path both agents are inspected with probability 0 and both violate. Indeed, suppose to the contrary, that on the equilibrium path, in the second period agent i is inspected with probability $p \geq \frac{1}{1+c}$. In this case, the principal benefits from reducing the inspection probability to 0. On the one hand, the additional violation in the second period is compensated by a lower inspection cost (since $1 < p \cdot r_1$), and on the other hand, this increases agent i's payoff on the equilibrium path and hence it does not alter his first-period best response action.

Therefore, the inspection probability an agent faces in the second period remains zero, regardless of which equilibrium path is selected in the first period. This implies that the monitoring structure does not play a role and hence both public and private monitoring yield the principal the same payoff.

If $r_2 \leq 1+c+\frac{c(1+c-r_1)}{1-c}$, then having both agents inspected and adhere is optimal already in the one-shot game. Since the inspection cost that induces full compliance is fixed regardless of the monitoring structure,[42] the principal can save inspection cost only by sacrificing the compliance level. In particular, there are cases where the principal could have profitably induce full compliance, but she deliberately allows agents to violate in the second period, as a cheaper means to deter violations in the first period. We find such cases less relevant to the main focus of this paper, which is to improve compliance by modifying the monitoring structure, and hence we omit the detailed analysis.

A.11 Conditions under which $L_T^{\text{pub}} = L_L$ in the T-period game

In this section we provide conditions under which $L_T^{\text{pub}} = L_L$ (or, alternatively, $L_T^{\text{pub}} > L_L$).

Proposition 3. *Suppose there are two agents and one inspector, and $\delta > \frac{1-c^2}{c}$. The set of (c, δ) under which $L_T^{\text{pub}} > L_L$ shrinks in T, and it converges to the area $\delta > \max\left(\frac{1-c^2}{c}, \frac{1+c}{2+c}\right)$.*

Proof. See Appendix B.1. □

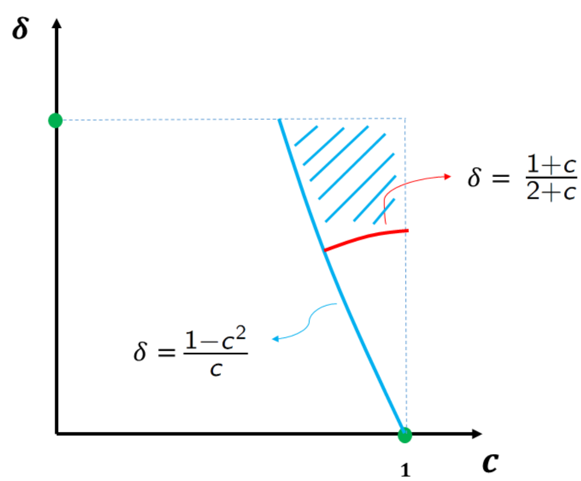

Figure 15: Parameter for which Private ≻ Public for every $T \geq 2$.

Proposition 3 asserts that, when T increases, the set of parameters (c, δ) under which public monitoring *does not* attain L_L shrinks, and as $T \to \infty$ it converges to the shaded area

[42] Under both monitoring structures, to induce full compliance, in the second period each agent has to be inspected with probability no less than $\frac{1}{1+c}$, regardless of the history. Therefore, an inspection cost of $r_1 + \frac{1-c}{1+c}r_2$ is inevitable in period 2. Since on the equilibrium path both agents obtain zero, to deter a violation in period 1, an inspection intensity $\frac{1}{1+c}$ is also needed (otherwise, the agent who adheres is subject to a positive continuation payoff). Consequently, an inspection cost of $r_1 + \frac{1-c}{1+c}r_2$ is also inevitable in the first period.

of Figure 15. Recall that by Theorem 7, $L_T^{\text{pub}} > L_L$ implies Private $\succ$ Public. Therefore, as a corollary, if $\delta > \max\left(\frac{1-c^2}{c}, \frac{1+c}{2+c}\right)$, Private $\succ$ Public in any T-period game[43]with $T \geq 2$.

We next discuss why a longer period of interaction benefits the principal under public monitoring (and hence allows her to attain the best feasible payoff L_L).

Consider a simplified scenario where the principal faces only one agent, and the inspection intensity available for this agent is $p < \frac{1}{1+c}$. In the one-shot game, the agent necessarily violates, inflicting a loss of $\gamma_1 := 1$ to the principal.

Now add one additional period to this game. In the first period of the two-period game, if the agent is inspected and found adhering[44] in period 1, he gets a continuation payoff of 1. Whereas if he is not inspected, the game continues as in the one-shot game. This scheme is optimal when the agent is not too impatient. Due to discounting, the loss of the principal is reduced to $\gamma_2 := \delta \cdot \big(p \cdot 1 + (1-p) \cdot \gamma_1\big) < \gamma_1$.

Suppose we add a third period to the two-period game. Again, in the first period, the agent's continuation payoff upon being inspected and found adhering[45] remains 1. Whereas if he is not inspected, the game continuous as a two-period problem. The principal's loss is further reduced to $\gamma_3 := \delta \cdot \big(p \cdot 1 + (1-p) \cdot \gamma_2\big) < \gamma_2$. This argument applies inductively when additional periods are added. Therefore, a larger T helps to delay violations, which lowers the principal's equilibrium loss, and makes it more likely that the lower-bound loss can be attained under public monitoring.

Note the difference between a large T and a large δ. When δ increases, the same number of future violations becomes more valuable for the agent. The lower-bound payoff $1 - c$ is hence more difficult to implement in such cases. As a result, a large δ makes Private $\succ$ Public easier to hold, and yields private monitoring a larger advantage.

A.12 Benefit from detecting violations

Consider the model with two agents and one inspector, and suppose the principal benefits from detecting violations. In this section, we argue that if the principal has no commitment power, then full defection in both periods is the unique equilibrium under both public and

[43]This result does not contradict the folk theorem. Indeed, even if the best feasible outcome can be attained in public monitoring as $\delta \to 1$, private monitoring may be superior for each sufficiently high discount factor. The degree of the superiority, however, may shrink to zero as $\delta \to 1$.

[44]If the agent is found violating in period 1, he is punished most severely by being inspected with probability 1 in the second period.

[45]To yield the agent a reward of 1, a violation in the third period is not enough — from the perspective of period 2, a violation in period 3 yields him a payoff of $\delta < 1$. Therefore, in addition to a violation in period 3, with a positive probability the agent is not inspected and violates already in the second period. This implementation requires the use of the public correlation device in period 2.

private monitoring.

First note that in the one-shot game, necessarily both agents violate in equilibrium. Indeed, suppose, instead, one agent adheres with a positive probability. Then this agent must be inspected with probability at least $\frac{1}{1+c}$. This implies that the other agent is inspected with a low probability and he violates. In this case, the principal is better off deviating from the prescribed strategy and inspecting the violator with probability 1, as it maximizes the probability of catching a violation.

Now consider the two period public monitoring problem. In the second period, agents necessarily play the unique Nash equilibrium (full defection) regardless of the history. Since agents' first-period actions do not affect their second-period payoff, they play the unique Nash equilibrium (again, full defection) also in the first period.

Next consider the private monitoring problem. For simplicity, we neglect the use of the public correlation device.[46] Recall that after the first period, there are five private histories for the principal: A_1, V_1, A_2, V_2, and $\emptyset$ (see Figure 1). The principal's second-period inspection strategy (and hence the agents' best responses) depend only on the principal's private history.[47] We next refer to the principal's private history simply as *history*.

Step 1: Deriving Figure 16.

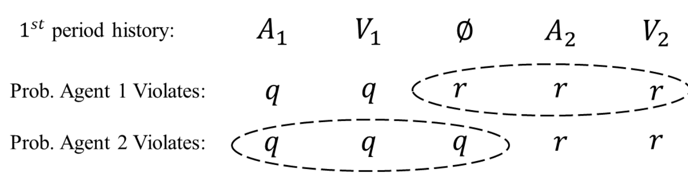

Figure 16: The probability agents violate at each history.

We now consider the probability of violation for each agent in the second period. Since Agent 1 cannot distinguish between histories in $\{\emptyset, A_2, V_2\}$, the probability he violates at these histories must be the same. The same argument applies to Agent 2 at histories $\{A_1, V_1, \emptyset\}$. See Figure 16 for an illustration, where r and q represents the violation probability for Agent 1 and 2, respectively.

We next argue that at histories in $\{A_1, V_1, A_2, V_2\}$, the probability that Agent 1 violates equals to the probability that Agent 2 violates.

Consider history A_1, and suppose to the contrary that at this history the probability that

[46] The proof in this section can be readily generalized to the game that involves public correlation device, since our argument applies to each realization of the public correlation device.

[47] Since the agent's private history is more coarse than that of the principal, the agent's action in period 2 depends on the signal he received, which determines his belief over the principal's private history.

Agent 1 violates is x_1, which is different from q. If $x_1 < q$, then Agent 2 violates with a probability higher than Agent 1. The principal, who does not have a commitment power and has an incentive to catch violations, is best off inspecting Agent 2 with probability 1. Since Agent 1 knowns that the real history is A_1, he knows that he is inspected with probability zero, and hence he violates with probability $x_1 = 1$, a contradiction to the assumption that $x_1 < q$. Suppose next $x_1 > q$. In this case, the principal is best off inspecting Agent 1 with probability 1. As a result, Agent 1 violates with probability $x_1 = 0$, a contradiction to the assumption that $x_1 > q$.

The same logic applies to histories V_1, A_2, and V_2. Note that this logic does not apply to history $\emptyset$, since even if an agent is inspected with probability 0 in this eventuality, he is uncertain about the true history, and hence he may still choose to adhere with a positive probability.

Step 2: The case $r \neq q$.

Suppose without loss of generality that $r < q$. Since $r < q \leq 1$, at histories A_2 and V_2 Agent 2 adheres with a positive probability. Since Agent 2 knows the true history in these events, he must be inspected with an inspection probability at least $\frac{1}{1+c}$. As for Agent 1, $r < 1$ implies that he adheres with a positive probability if he is not inspected in period 1. Since at histories A_2 and V_2, Agent 2 is inspected with probability at least $\frac{1}{1+c}$, in these histories Agent 1 is inspected with probability less than $\frac{c}{1+c}$. To induce Agent 1 to adhere, it must be the case that at history $\emptyset$, Agent 1 is inspected with probability at least $\frac{1}{1+c}$. However, if no agent is inspected in period 1, then Agent 2 violates with a higher probability than Agent 1 (since $q > r$), which implies that in the second period the principal is best off inspecting Agent 1 with probability 0, a contradiction. Therefore, $r < q$ cannot be an equilibrium outcome.

Step 3: The case $r = q$.

If $r = q = 1$, then all agents violate in period 2. This can be attained as an equilibrium with the principal inspecting each agent with probability $\frac{1}{2}$ regardless of the history. If $r = q < 1$, then regardless of the history, in the second period both agents adhere with a positive probability. Since an agents assigns a positive probability to adherence only if the expected inspection intensity he faces is at least $\frac{1}{1+c}$, in such an equilibrium necessarily each agent believes that he is inspected with probability at least $\frac{1}{1+c}$ regardless of the history. This contradicts the resource constraint. When $r = q = 1$, agents' first-period actions do not affect their second-period payoff, and hence they play the unique Nash equilibrium (both violating) in the first period. As a result, full defection in both period is the unique equilibrium under private monitoring.

A.13 Proof of Proposition 2

First note that for $\delta > 1 - c$, the principal can attain an equilibrium loss of δ by using the inspection scheme shown in Figure 3. Therefore, under the optimal inspection scheme with private monitoring, necessarily both agents adhere in period 1 — a violation in period 1 causes a loss of $1 > \delta$ to the principal, which is clearly not optimal.

There are two approaches to deter an agent from violating in the first period. The "myopic approach" is to inspect an agent with probability at least $\frac{1}{1+c}$. In this case, an agent adheres even in the absence of second-period incentives. The "dynamic approach" is to inspect an agent with probability lower than $\frac{1}{1+c}$ and to reward the agent by not inspecting him in the second period, provided he is inspected and found adhering in the first period. As long as the first-period inspection probability is at least $p^* = \frac{1}{1+c+\delta}$, the agent adheres in period 1. By Assumption 1, $\frac{2}{1+c} > 1$, and hence it is not feasible to deter both agents by the "myopic approach".

<u>Region A</u>: $1 - c < \delta < \frac{1-c^2}{c}$. In this case, $p^* + \frac{1}{1+c} > 1$. This implies that to deter both agents from violating in period 1, both agents have to be treated by the "dynamic approach" (the condition $1 - c < \delta$ guarantees that $2p^* < 1$). In such case, the minimum violation is $2\delta p^*$: if an agent is inspected with probability lower than p^* in period 1, then even with a free violation in the second period, the agent is better off violating in period 1, and the total number of violations is larger than $2\delta p^*$. If an agent is inspected with probability higher than p^* in period 1, since a successfully passed inspection in period 1 follows by a free violation[48] in period 2, the total number of violation is larger than $2\delta p^*$.

Consider the following strategy profile σ^A (see Figure 17) that yields the minimum number of violations.

- Strategy of the principal:
  - In the first period, the principal inspects each agent with the same probability p^*. With probability $1 - 2p^*$ no agent is inspected.
  - If no agent is inspected in the first period, then in the second period both agents are inspected with probability $\frac{1}{2}$.

[48]Here we use the fact that the agent being treated by the "dynamic approach" has to be rewarded in period 2, and the assumption that the public correlation device is excluded in period 2.

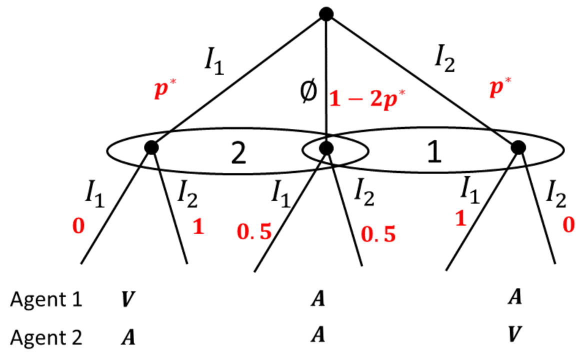

Figure 17: Region A: the structure of σ^A.

 – Otherwise, if the inspected agent is found adhering in the first period, then in the second period this agent is inspected with probability 0, and the other agent is inspected with probability 1.

 – Punishment: if the inspected agent is found violating in the first period, this agent is inspected with probability 1 in the second period.

- The strategies of the agents are: both agents adhere in the first period, and only the agent who is inspected and found adhering in the first period violates in the second period.

To verify that σ^A is an SE, the main step is to show that if an agent is not inspected in period 1, he adheres in the second period. Indeed, the expected inspection probability an agent faces if he is not inspected in period 1 is

$$1 \cdot \frac{p^*}{1-p^*} + 0.5 \cdot \frac{1-2p^*}{1-p^*} = \frac{1+c+\delta}{2(c+\delta)},$$

and it can be verified that $\frac{1+c+\delta}{2(c+\delta)} > \frac{1}{1+c}$, as desired.

Region C: $\delta \geq \frac{1+c-c^2-c^3}{c^2+c-1}$. In this case, $\delta > \frac{1-c^2}{c}$, and hence $p^* + \frac{1}{1+c} < 1$. Since under Assumption 1 it is impossible to deter both violations by the "myopic approach", the best the principal can achieve is to treat one agent with the "myopic approach", and the other agent with the "dynamic approach". Since in the "dynamic approach" a successful inspection in period 1 leads to a free violation in period 2, to minimize the number of violations, the principal minimizes the first-period inspection probability to p^*, and the minimum violation in this case is δp^*.

Consider the following strategy profile σ^C that yields the minimum number of violations δp^* (see Figure 18).

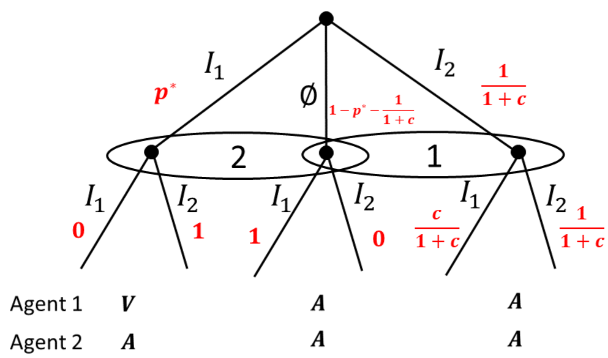

Figure 18: Region C: the structure of σ^C.

- Strategy of the principal:
  - In the first period, inspect Agent 1 with probability p^*, inspect Agent 2 with probability $\frac{1}{1+c}$. With the remaining probability, inspect no agent.
  - If Agent 1 is found adhering in the first period, then in the second period inspect Agent 1 and Agent 2 with probabilities 0 and 1, respectively.
  - If Agent 2 is found adhering in the first period, then in the second period inspect Agent 1 and Agent 2 with probabilities $\frac{c}{1+c}$ and $\frac{1}{1+c}$, respectively.
  - If no agent is inspected in the first period, then in the second period inspect Agent 1 and Agent 2 with probabilities 1 and 0, respectively.
  - Punishment: if the inspected agent is found violating in the first period, this agent is inspected with probability 1 in the second period.
- The strategies of the agents:
  - Both agents adhere in the first period.
  - If Agent 1 is inspected in the first period, he violates in the second period.
  - In all other scenarios, each agent adhere in the second period.

We next argue that σ^C is an SE. The key step is to show that in the second period, an agent who is not inspected in period 1 is better off adhering. If Agent 2 is not inspected in period 1, the expected inspection probability he faces in the second period is

$$1 \cdot \frac{p^*}{1-\frac{1}{1+c}} + 0 \cdot \frac{1-p^*-\frac{1}{1+c}}{1-\frac{1}{1+c}} = \frac{1+c}{c(1+c+\delta)}. \tag{19}$$

It can be verified that $\frac{1+c}{c(1+c+\delta)} > \frac{1}{1+c}$, and hence Agent 2 adheres.

If Agent 1 is not inspected in period 1, the expected inspection probability he faces in the second period is

$$1 \cdot \frac{1 - p^* - \frac{1}{1+c}}{1 - p^*} + \frac{c}{1+c} \cdot \frac{\frac{1}{1+c}}{1 - p^*}. \tag{20}$$

It can be verified that the quantity in Eq. (20) is larger than $\frac{1}{1+c}$ if and only if $\delta \geq \frac{1+c-c^2-c^3}{c^2+c-1}$. Therefore, σ^C is an SE that yields a loss of δp^* to the principal.

Region B: $\frac{1-c^2}{c} < \delta < \frac{1+c-c^2-c^3}{c^2+c-1}$. In this case, the strategy structure σ^A remains an SE, and it yields a discounted number of violations $2\delta p^*$. We next argue that this is the optimal outcome.

In the current case, $p^* + \frac{1}{1+c} \leq 1$, and hence in the first period the principal can treat one agent myopically and the other agent dynamically. Since $\delta < \frac{1+c-c^2-c^3}{c^2+c-1}$, when using the inspection scheme depicted in Figure 18, the quantity in Eq. (20) is lower than $\frac{1}{1+c}$. That is, if Agent 1 is not inspected in period 1, he violates in period 2. Therefore, σ^C is no longer an SE.

We next argue that when treating one agent myopically and the other agent dynamically in the first period, the discounted number of violations is always larger than $2\delta p^*$. Denote the inspection probabilities for Agent 1 and 2 in the first period by p_1 and p_2, respectively. Suppose without loss of generality that Agent 1 is treated dynamically (hence $p_1 \geq p^*$), and Agent 2 is treated myopically (hence $p_2 \geq \frac{1}{1+c}$). Upon a successful inspection, Agent 1 is rewarded with a free violation, which yields the number of violation $\delta p_1 \geq \delta p^*$. Therefore, the only possible scenario under which the total number of violations is less than $2\delta p^*$ is that except the above violation, both agents adhere. Suppose this is possible. Then the inspection scheme must have the structure depicted in Figure 19.

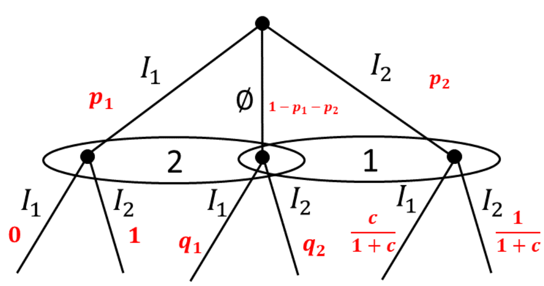

Figure 19: An inspection scheme for the principal.

Here, if Agent 1 is not inspected in the first period, the expected inspection probability for him is

$$q_1 \cdot \frac{1-p_1-p_2}{1-p_1} + \frac{c}{1+c} \cdot \frac{p_2}{1-p_1} \leq 1 \cdot \frac{1-p_1-p_2}{1-p_1} + \frac{c}{1+c} \cdot \frac{p_2}{1-p_1}$$
$$= 1 - \frac{1}{1+c} \cdot \frac{p_2}{1-p_1} \tag{21}$$
$$\underbrace{\leq 1 - \frac{1}{1+c} \cdot \frac{\frac{1}{1+c}}{1-p^*}}_{\text{eqiuvalent to (20)}} < \frac{1}{1+c}.$$

The last inequality in Eq. (21) holds because the quantity in Eq. (20) is smaller than $\frac{1}{1+c}$. This implies that under the inspection scheme depicted in Figure 19, Agent 1 is better off violating if he is not inspected in the first period. The discounted number of violations is hence δ. Therefore, when treating one agent dynamically and another agent myopically, it is impossible to attain violations less than $2\delta p^*$, and hence the principal prefers the SE σ^A. When treating both agent dynamically, σ^A is optimal, as discussed in Region A. Therefore, in the current case the optimal SE yields $2\delta p^*$ violations.

A.14 Optimal information revelation: two-period game

In the two period game, the lower-bound loss of the principal is $1-c$, see Fact 1. In this section we show that the optimal information revelation attains this lower bound.

An important implication of this exercise is to show that the optimal information revelation scheme in the two-period game is not a simple repetition of the solution to the one-shot game, as the latter yields the principal a loss of $(1+\delta)(1-c)$. Consider the following strategy profile.

- In the first period, the recommendations to the players are detailed in Table 2, with $x = \frac{(1-c)(1-\delta)}{2}$.

Probability	Principal	Agent 1	Agent 2
x	I_1	A	V
x	I_2	V	A
$0.5-x$	I_1	A	A
$0.5-x$	I_2	A	A

Table 2: First period recommendations.

- If Agent 1 is inspected in the first period and found adhering, then in the second period the principal recommends the agents to implement the payoff vector $(v_1 = 1-c, v_2 = 0)$, as described in Table 1.

- If Agent 2 is inspected in the first period and found adhering, then in the second period the principal recommends the agents to implement the payoff vector $(v_1 = 0, v_2 = 1-c)$, which is the analogue of Table 1.

- Punishment: if an agent, say agent i, is found violating in the first period, then in the second period he is inspected with probability 1 and the other agent is inspected with probability 0. agent i adheres and the other agent violates.

When the agents follow the recommended strategy, the principal loses $(1-\delta)(1-c)$ in the first period, and $\delta(1-c)$ in the second period. The total loss of the principal is thus $1-c$ and, as argued before, this is the lowest possible equilibrium loss in the two-period game.

We now verify that the proposed strategy profile constitutes an equilibrium of the two-period game. We start by showing that under the proposed information revelation scheme, both agents are better off following the recommended actions. Since the agents are symmetric, we only verify this for Agent 1.

If Agent 1 is told to violate in a certain period, then for sure he is not inspected in that period, and hence following the recommendation is optimal. Now consider the second period under the history A_1 (resp., A_2), and suppose Agent 1 is told to adhere. The expected inspection probability for Agent 1 in this case is $\frac{\frac{c}{1+c}}{\frac{c}{1+c}+\frac{c^2}{1+c}} = \frac{1}{1+c}$ (resp., $1-c+\frac{c^2}{1+c} = \frac{1}{1+c}$), and adhering is an optimal response. Finally, we turn to the first period. In the first period, if Agent 1 is told to adhere, then he faces an inspection probability $\frac{0.5}{1-x} = \frac{1}{1+c+(1-c)\delta} < \frac{1}{1+c}$. Since $f(\frac{0.5}{1-x}) = 1-c$, Agent 1 is better off adhering in the first period.

When players follow the recommended actions, the agents' payoffs are $(v_1 = \frac{1-c}{2}, v_2 = \frac{1-c}{2})$, and the principal's loss is $1-c$, which is the lower bound on her loss in a PPE/SE.

It is natural to ask whether the lower-bound payoff can be attained also in T-period games with $T > 2$. We conjecture that the answer is negative. The reason is that in the two-period game, even though the payoff vector $(v_1 = \frac{1-c}{2}, v_2 = \frac{1-c}{2})$ can be attained as an equilibrium payoff (as shown above), the payoff vector $(v_1 = 1-c, v_2 = 0)$ cannot. This is in contrast to the one-shot game, and prevents us from using induction to generalize the argument for the two-period game to games with a larger number of periods.

A.15 Information design that relies on public signals only

Consider the two-period problem with one inspector and two agents, and suppose monitoring is private. In this section, the principal has a commitment power and is allowed to send public signals regarding the true history. In particular, if the principal reveals *all* her past observations, the monitoring becomes public. The question is, whether by partially revealing the history, the principal can attain an outcome superior to both private and public monitoring.

As in Section 4.2, we assume that $\delta > 1 - c$ and exclude the use of the public correlation device. Moreover, since public histories may serve as a correlation device, to separate the effect of the public signal from that of the public correlation device, in this section we require the inspection strategy to be independent of the public signal.

The timeline is summarized as follows. The principal first publicly announces and commits to an inspection scheme and a public disclosure rule. Agents then take their first-period actions, and the principal conducts the inspection. Given the realization of players' first-period actions, the principal observes a private signal $y_0 \in Y_0 = \{V_1, A_1, V_2, A_2, \emptyset\}$, and sends a public signal about her observation. In the second period, agents take their actions according to their private inspection histories and the public signal (which together form the agents' private histories), and the principal conducts the inspection (which depends only on y_0 and not on the public signal). We study the sequential equilibrium of this game: deviations of the agents are non-profitable after every private history. Note that if the principal reveals nothing, the monitoring is private; if the principal fully reveals her observation, the monitoring is public.

As shown in Theorem 1, in the current setup, private monitoring is always weakly better than public monitoring. Our question therefore boils down to whether by partially revealing the history, the principal can attain an outcome better than the best equilibrium outcome under private monitoring (the latter is studied in Section 4.2).

As in Appendix A.13, we will distinguish between three cases, according to whether (c, δ) lies in Region A, B, or C in Figure 6. When (c, δ) lies in Regions A or C, the lower-bound violations derived in Appendix A.13 remains a lower-bound in the current case (the argument is similar to that in Appendix A.13 and hence omitted). Since private monitoring already attains this lower-bound (as shown in Appendix A.13), partial revelation of the history cannot further benefit the principal.

Consider next Region B in Figure 6. By Proposition 2, the minimum violation under private monitoring in this case is $2\delta p^*$. We next show by an example that by using public signals, the principal can attain a better outcome.

Consider the inspection scheme depicted in Figure 20. First, as a benchmark, suppose the principal does not reveal any history (that is, the monitoring is private). Agents' best responses are given in Figure 20.[49] Indeed, under private monitoring, if Agent 1 is not

[49]Note the difference between agents' actions in Figure 18 and Figure 20. In the former (parameters in Region C), if Agent 1 is not inspected in period 1, he adheres in period 2. Whereas in the latter (parameters

inspected in period 1, then the expected inspection probability he faces in period 2 is

$$1 \cdot \frac{1-p^* - \frac{1}{1+c}}{1-p^*} + \frac{c}{1+c} \cdot \frac{\frac{1}{1+c}}{1-p^*}. \tag{22}$$

This quantity is lower than $\frac{1}{1+c}$ for parameters in the current case (Region B), and hence Agent 1 violates.

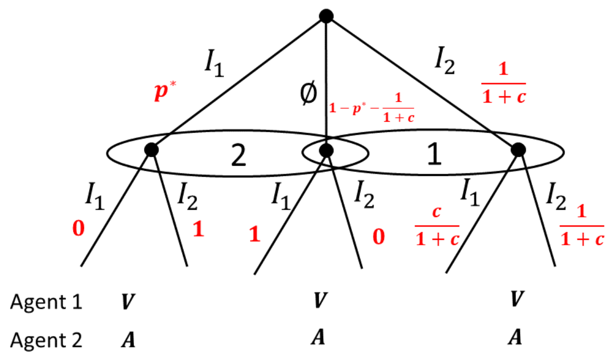

Figure 20: An equilibrium under Region B.

We next show that by partially sending public signals, the principal can reduce the number of violations to a quantity lower than $2\delta p^*$ (that is, the minimum equilibrium violation under private monitoring). Consider the following public signal: if the true history in period 1 is I_2, then with probability x the principal publicly announces I_2. Otherwise, the principal remains silence. If the principal announces I_2, then Agent 1 knows that he faces inspection probability $\frac{c}{1+c}$ in the second period, and he violates. If Agent 1 is not inspected in period 1 and the principal remains silence, then the inspection probability Agent 1 faces in period 2 is

$$1 \cdot \frac{1-p^* - \frac{1}{1+c}}{1-p^* - \frac{x}{1+c}} + \frac{c}{1+c} \cdot \frac{\frac{1}{1+c}\cdot(1-x)}{1-p^* - \frac{x}{1+c}}. \tag{23}$$

Compared with (22), expression (23) assigns a larger weight to inspection probability 1, and its value increases in x. Let $x^* := \frac{1+\delta+(1-\delta)c-(1+\delta)c^2-c^3}{(1-c)(1+c+\delta)}$. It can be verified that when $x \geq x^*$, (23) is (weakly) larger than $\frac{1}{1+c}$, and Agent 1 adheres in period 2 if he was not inspected in period 1 and the principal does not announce I_2. The total number of violations in this case is $\delta(p^* + \frac{x}{1+c})$.

Take, for instance, $(c = 0.85, \delta = 0.7)$. It can be verified that this pair of parameter lies in Region B, and $x^* = 0.29$. By setting $x = x^*$, the total number of violations is $\delta(p^* +$

in Region B), if Agent 1 is not inspected in period 1, he violates in period 2.

$\frac{x^*}{1+c}) = 0.39$. This quantity is lower than the minimum violation under private monitoring $2\delta p^* = 0.55$. Thus, a partial revelation of public signal improves the principal's payoff.

B Appendix B. Proofs for T-period games.

B.1 Proof of Theorem 7 and Proposition 3

In This section we provide the detailed analysis for the T-period game. Throughout this section we assume that there are two agents and one inspector, and $\delta > \frac{1-c^2}{c}$. Denote by G_T^{pub} (resp. G_T^{pri}) the T-period public monitoring game (resp. private monitoring game). Solan and Zhao (2021) studied the infinitely repeated game with public monitoring, and here we adapt their analysis to G_T^{pub}.

Preliminaries. Analogously to Section 2.5, without loss of generality we can restrict attention to equilibria that satisfy the following conditions:

(i) Whenever an agent is indifferent between adhering and violating, he adheres.

(ii) On the equilibrium path, whenever an agent violates, he is inspected with probability zero.

(iii) On the equilibrium path, an agent who is inspected and found violating is punished in the most severe way: he is inspected with probability 1 in all future periods.

The proof of this claim follows from results in Solan and Zhao (2021).

Denote by $\mathscr{E}_*(G_T^{\text{pub}})$ the set of all PPEs that satisfy conditions (i)–(iii) in the T-period game. Under every PPE $\sigma = (\sigma_0, \sigma_1, \sigma_2)$ in $\mathscr{E}_*(G_T^{\text{pub}})$, an agent violates if and only if he is inspected with probability 0, hence in the subgame following the announcement of σ_0, the agents' best response strategies in $\mathscr{E}_*(G_T^{\text{pub}})$ are uniquely determined. In the rest of the section we focus only on PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$. A result similar to Proposition 1 follows: An agent who is inspected with probability p adheres if and only if his continuation payoff upon *being inspected and found adhering* is at least $f(p)$.

Proposition 4. *Suppose that σ is a PPE in $\mathscr{E}_*(G_T^{\text{pub}})$. Then for every history h^k that occurs with positive probability, agent i adheres if and only if $v_i(\sigma|_{(h^1,A_i)}) \geq f\big(\sigma_0(I_i|h^1)\big)$.*

We are now ready to prove Theorem 7 and Proposition 3. For every $T \geq 1$, let γ_T be the minimum equilibrium payoff for Agent 1 in the T-period game, given that Agent 2 obtains zero.[50]

The rest of this appendix is structured as follows. Let $T \geq 2$. In Section B.1.1 we show that if $\gamma_{T-1} \leq \frac{1-c}{\delta}$, then in the T-period game public monitoring can attain the lower-bound loss $1-c$, and Private $\approx$ Public. In Section B.1.3 we show that if $\gamma_{T-1} > \frac{1-c}{\delta}$, then in the

[50] This value is the same regardless of the monitoring structure: Since Agent 2 obtains 0, he is necessarily inspected with probability $\frac{1}{1+c}$ in each period, regardless of the history. Therefore, the inspection intensity for Agent 1 depends only on his own inspection history, which is known to him. It follows that revealing past inspection results has no advantage in this case.

T-period game public monitoring *cannot* attain the lower bound $1-c$, and Private $\succ$ Public. Theorem 7 follows immediately from these two results.

In Section B.1.2 we provide the detailed characterization of γ_T and shows that the sequence $\{\gamma_T\}_{T\geq 1}$ is decreasing in T and converges to $\frac{(1-c)(1+c)}{1+c-\delta}$. Therefore, the set of parameters (c,δ) under which $\gamma_{T-1} > \frac{1-c}{\delta}$ is smaller as T increases. Moreover, as long as $\frac{(1-c)(1+c)}{1+c-\delta} > \frac{1-c}{\delta}$ (which is equivalent to $\delta > \frac{1+c}{2+c}$), for any $T \geq 2$ we have $\gamma_{T-1} > \frac{1-c}{\delta}$ and Private $\succ$ Public. Proposition 3 follows.

B.1.1 The case $\gamma_{T-1} \leq \frac{1-c}{\delta}$.

We will prove that in this case, in the T-period game the principal can attain the lower-bound loss $1-c$ already under public monitoring. To this end, we first present a technical result. For every $T \geq 2$, let $\Delta_T := 1 + \delta + \delta^2 + \cdots + \delta^{T-1}$. For every $T \geq 1$ and every $x \in [0, \Delta_T]$, let $g_T^{\text{pub}}(x)$ be the minimum payoff of Agent 1, where the minimum is taken over all PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ that yield Agent 2 the payoff x. Note that $g_T^{\text{pub}}(0) = \gamma_T$.

Proposition 5. *Suppose that in the T-period game, under public monitoring the payoff vector $(v_1 = x, v_2 = y)$ is a PPE outcome in $\mathscr{E}_*(G_T^{\text{pub}})$. Then for every $y' \in [y, \Delta_T]$, the payoff vector $(v_1 = x, v_2 = y')$ can be implemented as a PPE outcome in $\mathscr{E}_*(G_T^{\text{pub}})$.*

Proof. If $y = \Delta_T$, the result holds trivially. If $y < \Delta_T$, then there exist histories on the equilibrium path of σ under which Agent 2 adheres. To increase Agent 2's payoff, we choose a set of histories in which Agent 2 adheres, and we change the play at these histories: Agent 2 will not be inspected and will violate. This way we construct an equilibrium that yields Agent 2 a payoff higher than $g_T^{\text{pub}}(x)$ without affecting the payoff of Agent 1. To make the increment exactly equal to $y' - y$, we may need the use of the correlation device. □

Now consider the following PPE $\underline{\sigma}$ in G_T^{pub}.

- In the first period inspect each agent with probability $\frac{1}{2}$. Both agents adhere.
- If Agent 1 (resp., Agent 2) is inspected in the first period and found adhering, from the second period on the players implement a strategy profile that yields the payoffs $\big(f(\frac{1}{2}), 0\big)$ (resp., $\big(0, f(\frac{1}{2})\big)$).
- Punishment: If in the first period an agent is found violating, then in all future periods he is inspected with probability 1, he adheres, and the other agent violates.

The assumption $\gamma_{T-1} \leq \frac{1-c}{\delta} = f(\frac{1}{2})$ together with Proposition 5 imply that the payoff vectors $\big(f(\frac{1}{2}), 0\big)$ and $\big(0, f(\frac{1}{2})\big)$ can be achieved as PPE outcomes in G_{T-1}^{pub}. It can be verified that $\underline{\sigma}$ is a PPE that yields both agents the expected payoff $\frac{1}{2} \cdot \delta f(\frac{1}{2})$. The sum of the agents' payoffs is $\delta f(\frac{1}{2}) = 1 - c$, as desired.

B.1.2 Recursive formula for γ_T.

We here present a recursive formula for the sequence $\{\gamma_T\}_{T \geq 1}$.

Proposition 6. *Suppose there are one inspector and two agents. If* $\delta > \frac{1-c^2}{c}$, *then* $\gamma_1 = 1$, *and for every* $T \geq 2$,

$$\gamma_T = \begin{cases} \delta \cdot \gamma_{T-1}, & \text{if } \gamma_{T-1} \geq f(\frac{c}{1+c}), \\ \delta \cdot \left(\frac{1}{1+c} \cdot \gamma_{T-1} + \frac{c}{1+c} \cdot f(\frac{c}{1+c})\right), & \text{if } \gamma_{T-1} < f(\frac{c}{1+c}). \end{cases} \tag{24}$$

In particular, the sequence $\{\gamma_T\}_{T \geq 1}$ *is decreasing in* T *and converges to* $\frac{(1-c)(1+c)}{1+c-\delta}$.

Proof. Step 1: For every $T \geq 1$, to implement $\big(v_1 = \gamma_T, v_2 = 0\big)$ in G_T^{pub}, no correlation device has to be used in the first period.

Let $\sigma \in \mathscr{E}_*(G_T^{\text{pub}})$ be a PPE with corresponding payoff $(v_1 = \gamma_T, v_2 = 0)$, and suppose that σ uses the correlation device in the first period. Let (V_1, V_2) be the random variable that represents the expected outcome given the realization of the correlation device. Then (V_1, V_2) is a PPE outcome in G_T^{pub} almost surely (a.s.). Next, since Agent 2 can guarantee the payoff 0 by always adhering, $V_2 \geq 0$ a.s. Since the expectation of V_2 is $v_2 = 0$, we have $V_2 = 0$ a.s. Since γ_T is the lowest PPE outcome for Agent 1, when Agent 2's payoff is zero, we deduce that $V_1 \geq \gamma_T$ a.s. Since the expectation of V_1 is $v_1 = \gamma_T$, we have $v_1 = \gamma_T$ a.s. In particular, no correlation device is needed in the first period.

Let $\widehat{\gamma}_T$ be the minimum payoff that Agent 1 gets over all PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ that yield Agent 2 the payoff 0, provided that *Agent 1 adheres in the first period.* Naturally, $\gamma_T \leq \widehat{\gamma}_T$.

Step 2: For every $T \geq 2$, $\{\widehat{\gamma}_T\}_{T \geq 2}$ satisfies the following equation:

$$\widehat{\gamma}_T = \begin{cases} \delta \cdot \gamma_{T-1}, & \text{if } \gamma_{T-1} \geq f(\frac{c}{1+c}), \\ \delta \cdot \left(\frac{1}{1+c} \cdot \gamma_{T-1} + \frac{c}{1+c} \cdot f(\frac{c}{1+c})\right), & \text{if } \gamma_{T-1} < f(\frac{c}{1+c}). \end{cases} \tag{25}$$

To guarantee that Agent 2 obtains the payoff 0, Agent 2 has to be inspected with probability at least $\frac{1}{1+c}$ in all periods, and in particular in the first period. Denote the inspection probabilities for Agent 1 and Agent 2 in the first period by p_1 and p_2, respectively. Without loss of generality[51] we can assume that $p_1 + p_2 = 1$.

[51] We can always increase the inspection probability for Agent 2 to satisfy this condition.

The continuation payoff of Agent 1 satisfies the following conditions: (i) To ensure that Agent 2's continuation payoff is 0, Agent 1's continuation payoff is at least γ_{T-1}. (ii) Since Agent 1 adheres in the first period, by Proposition 6 this implies that Agent 1's continuation payoff when he is inspected is at least $f(p_1)$. Consequently,

$$\widehat{\gamma}_T \geq \min_{p_1,p_2} \Big(\delta \cdot \big(p_2 \cdot \gamma_{T-1} + p_1 \cdot \max\{f(p_1), \gamma_{T-1}\}\big)\Big), \tag{26}$$

where the minimum is taken over all (p_1, p_2) such that $p_2 \geq \frac{1}{1+c}$ and $p_1 + p_2 = 1$. Since the function $p \mapsto pf(p)$ is decreasing, the minimum on the right-hand-side of (26) is attained when $p_1 = \frac{c}{1+c}$ and $p_2 = \frac{1}{1+c}$.

We next claim that the minimum is attained as a PPE, so that equality holds in Eq. (26), which implies Eq. (25). Consider the following PPE: Agent 2 is inspected with probability $\frac{1}{1+c}$ in all periods and consequently he always adheres and obtains a payoff 0. Now we turn to Agent 1. In the first period Agent 1 is inspected with probability $\frac{c}{1+c}$. If Agent 1 is not inspected in the first period, he obtains a continuation payoff of γ_{T-1}. If Agent 1 is inspected in the first period, he obtains a continuation payoff $\max\{f(p_1), \gamma_{T-1}\}$ (by Proposition 5, this payoff can be implemented as a PPE outcome). This strategy constitutes a PPE that yields Agent 1 an expected payoff as shown in the right-hand side of Eq. (26).

Step 3: $\gamma_T = \widehat{\gamma}_T$ for all T.

As argued before, since $\delta > \frac{1-c^2}{c}$ (or equivalently $f(\frac{c}{1+c}) < 1$), in the two-period game $\gamma_2 = \delta < 1$. By Eq. (25) and since $\gamma_T \leq \widehat{\gamma}_T$, we have $\widehat{\gamma}_T < 1$ for every $T \geq 3$. Since we only consider PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ and since by Step 1 no correlation device is used in the first period, in the first period Agent 1 either obtains 0 (if he adheres) or 1 (if he violates). If under σ Agent 1 violates in the first period, he inflicts a loss of 1 on the principal, which is worse than the loss of γ_T. Consequently, for $T \geq 3$, γ_T agrees with $\widehat{\gamma}_T$ and satisfies the recursive formula shown in Eq. (24).

Step 4: The sequence $\{\gamma_T\}_{T\geq 1}$ is decreasing in T and converges to $\frac{(1-c)(1+c)}{1+c-\delta}$.

By Eq. (24), $\gamma_T = \delta\gamma_{T-1}$ as long as $\gamma_{T-1} \geq f(\frac{c}{1+c})$. Let T^* be the smallest integer that satisfies $\gamma_{T^*} \geq f(\frac{c}{1+c})$ and $\gamma_{T^*+1} < f(\frac{c}{1+c}) < 1$. Since $\gamma_1 = 1 > f(\frac{c}{1+c})$, such an integer exists and it is at least 1. By Eq. (24), $\gamma_{T^*+1} = \delta \cdot \gamma_{T^*} < \gamma_{T^*}$. We now verify that $\gamma_{T^*+2} < \gamma_{T^*+1}$.

Indeed,

$$\begin{aligned}\gamma_{T^*+2} &= \delta \cdot \left(\frac{c}{1+c} \cdot f(\frac{c}{1+c}) + \frac{1}{1+c} \cdot \gamma_{T^*+1}\right) \\ &< \delta \cdot \left(\frac{c}{1+c} \cdot \gamma_{T^*} + \frac{1}{1+c} \cdot \gamma_{T^*}\right) = \gamma_{T^*+1}.\end{aligned} \tag{27}$$

It now follows by induction that for every $T > T^*$, we have $\gamma_{T+2} < \gamma_{T+1} < 1$. The sequence $(\gamma_T)_{T\geq 1}$ is decreasing and bounded below by 0, hence converges to a limit, denoted L. This limit is a solution to $L = \delta \cdot \left(\frac{c}{1+c} \cdot f(\frac{c}{1+c}) + \frac{1}{1+c} \cdot L\right)$, which solves to $L = \frac{(1-c)(1+c)}{1+c-\delta}$. □

B.1.3 Case $\gamma_{T-1} > \frac{1-c}{\delta}$.

We now deal with the case $\gamma_{T-1} > \frac{1-c}{\delta}$. We first study the structure of the optimal PPE under public monitoring. We then show that the optimal PPE can be improved by an SE under private monitoring.

We first present a few technical results that are useful for the construction of the optimal PPE under public monitoring.

Proposition 7. *Let $T \geq 2$ and $x \in [0, \gamma_T]$. There is a PPE in $\mathscr{E}_*(G_T^{\text{pub}})$ that implements the payoff $\left(v_1 = x, v_2 = g_T^{\text{pub}}(x)\right)$ under which in the first period no correlation device is used and both agents adhere.*

Proof. See Appendix B.3. □

Proposition 8. *Fix $T \geq 2$. There is an optimal inspection strategy in the game G_T^{pub} under which the principal assigns probability zero to no inspection in those periods where at least one agent adheres. That is, $\sigma_0(\emptyset|h^t) = 0$ for every history h^t that satisfies $\sigma_1(A|h^t) = 1$ or $\sigma_2(A|h^t) = 1$.*

This result asserts that except those stages where both agents violate (in which case both agents are inspected with probability 0), in all other stages the principal cannot benefit from assigning a positive probability to being idle. Intuitively, if under a PPE σ in $\mathscr{E}_*(G_T^{\text{pub}})$, after a given history an agent, say, Agent 1, adheres and the probability that the inspector remains idle is positive, then by changing the inspector's strategy to inspect Agent 1 whenever she is supposed to be idle, we do not alter the incentives of the agents and this new PPE yields the agents the same payoff as σ.

Throughout the proof we focus on PPEs that satisfy propositions 7 and 8. Recall that $g_T^{\text{pub}}(x)$ is the minimum payoff of Agent 1, where the minimum is taken over all PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ that yield Agent 2 the payoff x.

Proposition 9. *The function $g_T^{\text{pub}}(x) : [0, \Delta_T] \to [0, \Delta_T]$ is well defined. Moreover, (i) the function g_T^{pub} is non-negative and non-increasing, and strictly decreasing on $[0, \gamma_T]$; (ii) the function g_T^{pub} is convex; (iii) $g_T^{\text{pub}}(\gamma_T) = 0$; (iv) the function g_T^{pub} is continuous on $[0, \gamma_T]$; (v) for every $x \in (0, \gamma_T)$, we have $g_T^{\text{pub}}\big(g_T^{\text{pub}}(x)\big) = x$.*

Proposition 9 is analogous to Proposition 8 in Solan and Zhao (2018), and its proof is omitted. See Figure 21 for an illustration of the function g_T^{pub}.

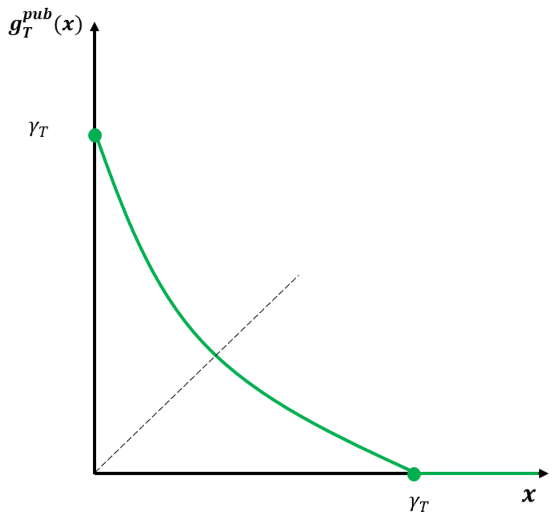

Figure 21: An illustration of the function g_T^{pub}.

Proposition 10. *Suppose $\gamma_{T-1} > \frac{1-c}{\delta}$. For every $x \in (0, \gamma_T)$, there exists a PPE $\widehat{\sigma} \in \mathscr{E}_*(G_T^{\text{pub}})$ that implements the payoff vector $\big(v_1 = x, v_2 = g_T^{\text{pub}}(x)\big)$ with the following structure (see Figure 23): the continuation payoff from the second period and on is $\big(W, g_{T-1}^{\text{pub}}(W)\big)$ (resp., $\big(g_{T-1}^{\text{pub}}(Z), Z\big)$) if Agent 1 (resp., Agent 2) is inspected in the first period. Moreover, $W \in (0, \gamma_{T-1})$ or $Z \in (0, \gamma_{T-1})$.*

Proposition 10 asserts that, to implement the payoff vector $\big(x, g_T^{\text{pub}}(x)\big)$, regardless of which agent is inspected in the current period, the continuation payoff vector lies on the boundary of the PPE payoff set of G_{T-1}^{pub}. That is, the continuation payoff of the agent who is *not inspected* is set to be as low as possible.

Proof. Consider the game G_T^{pub}. Suppose that the strategy profile $\widehat{\sigma} \in \mathscr{E}_*(G_T^{\text{pub}})$ implements the payoffs $\big(v_1 = x, v_2 = g_T^{\text{pub}}(x)\big)$, where $x \in (0, \gamma_T)$. By Proposition 7, we can assume that in the first period of $\widehat{\sigma}$ the correlation device is not used and both agents adhere. Since $x \in (0, \gamma_T)$, the monotonicity of g_T^{pub} implies that $g_T^{\text{pub}}(x) \in (0, \gamma_T)$. The first period under $\widehat{\sigma}$ is shown in Figure 22, where $p_1 := \widehat{\sigma}_0(I_1)$, $W := v_1(\widehat{\sigma}|_{I_1})$, $X := v_2(\widehat{\sigma}|_{I_1})$, $Y := v_1(\widehat{\sigma}|_{I_2})$, and

$Z := v_2(\widehat{\sigma}|_{I_2})$. Since we focus on PPEs that satisfy Proposition 7, both agents adhere in the first period and hence $W \geq f(p_1)$ and $Z \geq f(1-p_1)$.

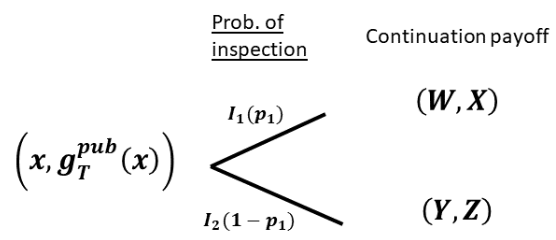

Figure 22: The first period under $\widehat{\sigma}$.

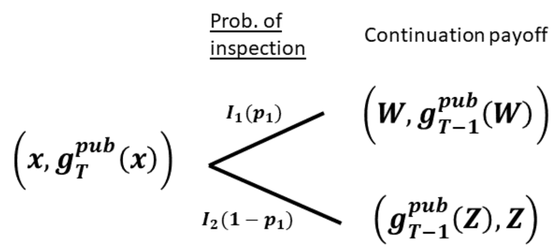

Figure 23: The first period under $\widehat{\sigma}$ (updated).

To make both agents' expected discounted payoff agree with their continuation payoffs, it is necessary that

$$x = 0 + \delta \cdot \big(p_1 \cdot W + (1-p_1) \cdot Y\big), \quad g_T^{\text{pub}}(x) = 0 + \delta \cdot \big(p_1 \cdot X + (1-p_1) \cdot Z\big). \tag{28}$$

We argue that $X = g_{T-1}^{\text{pub}}(W)$. Indeed, since we study the game with public monitoring, the continuation play is a PPE, hence $X \geq g_{T-1}^{\text{pub}}(W)$. Moreover, to minimize Agent 2's payoff at the outset of the game, one needs to lower Agent 2's continuation payoff, hence $X = g_{T-1}^{\text{pub}}(W)$. Since $x = g_T^{\text{pub}}\big(g_T^{\text{pub}}(x)\big)$ (see Proposition 9(v)), given that Agent 2 obtains a payoff $g_T^{\text{pub}}(x)$, the value x is the minimal equilibrium payoff for Agent 1. Therefore, a similar argument applies and $Y = g_{T-1}^{\text{pub}}(Z)$. We summarize the updated continuation payoffs in Figure 23.

We now show that without loss of generality we can assume that $W \in (0, \gamma_{T-1})$ or $Z \in (0, \gamma_{T-1})$. To this end, we rule out all cases where $W \notin (0, \gamma_{T-1})$ and $Z \notin (0, \gamma_{T-1})$.

Case 1: $W = 0$ and $Z \geq \gamma_{T-1}$. In this case $g_{T-1}^{\text{pub}}(Z) = 0$ and hence $x = \delta \cdot \big(p_1 \cdot W + (1 - p_1) \cdot g_{T-1}^{\text{pub}}(Z)\big) = 0$, a contradiction to $x \in (0, \gamma_T)$.

Case 2: $W \geq \gamma_{T-1}$ and $Z = 0$. In this case $g_{T-1}^{\text{pub}}(W) = 0$ and hence $g_T^{\text{pub}}(x) = \delta \cdot \big(p_1 \cdot g_{T-1}^{\text{pub}}(W) + (1-p_1) \cdot Z\big) = 0$. This implies that $x \geq \gamma_T$, a contradiction to $x \in (0, \gamma_T)$.

Case 3: $W = 0$ and $Z = 0$. Since both agents adhere in the first period, $W \geq f(p_1)$ and $Z \geq f(1-p_1)$. This implies that $p_1 \geq \frac{1}{1+c}$ and $1 - p_1 \geq \frac{1}{1+c}$, which, due to Assumption 1, is impossible to attain.

Case 4: $W \geq \gamma_{T-1}$ and $Z \geq \gamma_{T-1}$. This is the most challenging case. We will show that there exists another strategy profile σ' in the game G_T^{pub} with the structure shown in Figure 23 that also implements the payoffs $\big(v_1 = x, v_2 = g_T^{\text{pub}}(x)\big)$, with either $W' := v_1(\sigma'|_{I_1}) \in (0, \gamma_{T-1})$ or $Z' := v_2(\sigma'|_{I_2}) \in (0, \gamma_{T-1})$.

Since $\max(p_1, 1-p_1) \geq \frac{1}{2}$ and $\gamma_{T-1} > \frac{1-c}{\delta} = f(\frac{1}{2})$ (by the assumption of Proposition 10), we have

$$W, Z \geq \gamma_{T-1} > f(\frac{1}{2}) \geq f\big(\max\{p_1, 1-p_1\}\big) = \min\{f(p_1), f(1-p_1)\}. \tag{29}$$

This observation, together with the inequalities $W \geq f(p_1)$ and $Z \geq f(1-p_1)$, imply that either $W > f(p_1)$, or $Z > f(1-p_1)$. Suppose without loss of generality that $Z > f(1-p_1)$. Set $p_2' := f^{-1}(Z)$. Since $Z > f(1-p_1)$, we have $p_2' < 1-p_1$. Consider the following strategy profile σ' in the game G_T^{pub} (see Figure 24).

- In the first period, the principal inspects Agent 1 with probability $1-p_2'$, and Agent 2 with probability p_2'. Both agents adhere.
- Let $W' := \frac{p_1}{1-p_2'} \cdot W + \frac{1-p_1-p_2'}{1-p_2'} \cdot g_{T-1}^{\text{pub}}(Z) > 0$. If Agent 1 is inspected and found adhering, from the second period on the players implement a strategy profile that yields the payoffs $\big(W', g_{T-1}^{\text{pub}}(W')\big)$.
- If Agent 2 is inspected and found adhering, from the second period on the players implement a strategy profile that yields the payoffs $\big(v_1 = g_{T-1}^{\text{pub}}(Z), v_2 = Z\big)$.
- Punishment: If in the first period an agent is found violating, then in all future periods that agent is inspected with probability 1, he adheres, and the other agent violates.

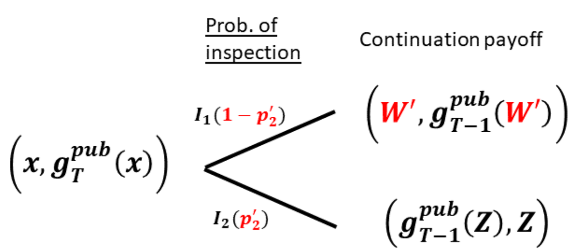

Figure 24: The first period under σ'.

Since $W'(1-p_2') = p_1 W + (1-p_1-p_2')g_{T-1}^{\text{pub}}(Z) \geq p_1 f(p_1) > (1-p_2')f(1-p_2')$, we have $Z = f(p_2')$ and $W' \geq f(1-p_2')$, and hence the strategy profile σ' is a PPE in $\mathscr{E}_*(G_T^{\text{pub}})$. By the convexity of g_{T-1}^{pub}, and since $g_{T-1}^{\text{pub}}\big(g_{T-1}^{\text{pub}}(Z)\big) = Z$,

$$g_{T-1}^{\text{pub}}(W') \leq \frac{p_1}{1-p_2'} \cdot g_{T-1}^{\text{pub}}(W) + \frac{1-p_1-p_2'}{1-p_2'} \cdot Z. \tag{30}$$

In fact, Eq. (30) holds with equality, since otherwise the strategy profile σ' constitutes a PPE in G_T^{pub} that yields the outcome $(v_1 = x, v_2 = y)$, where $y < g_T^{\text{pub}}(x)$, a contradiction to the definition of g_T^{pub}. Consequently, σ' also yields the agents the payoffs $\big(v_1 = x, v_2 = g_T^{\text{pub}}(x)\big)$. As can be seen in Figure 23, since $g_T^{\text{pub}}(x) > 0$, the right-hand side of Eq. (30) is positive, and hence $g_{T-1}^{\text{pub}}(W') > 0$. This observation, together with $W' > 0$ imply that $W' \in (0, \gamma_{T-1})$, as desired. □

We are now ready to compare public and private monitoring. To this end, it is convenient to restrict attention to a family of SEs in G_T^{pri}, which is analogous to the family $\mathscr{E}_*(G_T^{\text{pub}})$. Specifically, let $\mathscr{E}_*(G_T^{\text{pri}})$ be the set of SEs under which (i) the agents play pure strategies, and whenever an agent is indifferent between adhering and violating, he adheres; (ii) the loss of the principal equals the sum of payoffs of the agents; and (iii) on the equilibrium path, an agent who is inspected and found violating will be punished by being inspected with probability 1 in all future periods. For every x, let $g_T^{\text{pri}}(x)$ be the minimal payoff of Agent 1 among all SEs in $\mathscr{E}_*(G_T^{\text{pri}})$ in which Agent 2's payoff is x. Since we allow the use of the randomization device, the function g_T^{pri} is convex.

Remark 5. While in the game with public monitoring G_T^{pub} the restriction to $\mathscr{E}_*(G_T^{\text{pub}})$ is without loss of generality, this may not be the case in the game G_T^{pri} with private monitoring. However, as we will see, restricting attention to $\mathscr{E}_*(G_T^{\text{pri}})$ is sufficient for determining the relationship between the two monitoring structures.

The next proposition is analogous to Proposition 1, and states that the reward an agent should receive for adhering (if his behavior is observed by the principal) is given by the function f.

Proposition 11. *Suppose that σ is an SE in $\mathscr{E}_*(G_T^{\text{pri}})$. Then for every private history h_i^k, $i = 1, 2$, that occurs with positive probability under σ, $\sigma_i(A|h_i^k) = 1$ if and only if $v_i(\sigma|_{(h_i^k, A_i)}) - v_i(\sigma|_{(h_i^k, V_i)}) \geq f\big(E[\sigma_0(I_i|h_i^k)]\big)$, where $E[\sigma_0(I_i|h_i^k)]$ is the probability assigned by agent i after the private history h_i^k to the event that he is inspected.*

To show the superiority of private monitoring, we divide the argument into several steps. We first show that if $\gamma_{T-1} > \frac{1-c}{\delta}$ and $g_{T-1}^{\text{pri}}(x) < g_{T-1}^{\text{pub}}(x)$ for every $x \in (0, \gamma_{T-1})$, then $g_T^{\text{pri}}(x) < g_T^{\text{pub}}(x)$ for every $x \in (0, \gamma_T)$ (Step 1). We next show that $g_2^{\text{pri}}(x) < g_2^{\text{pub}}(x)$ for every $x \in (0, \gamma_2)$ (Step 2). By induction, we show that if $\gamma_{T-1} > \frac{1-c}{\delta}$, then $g_T^{\text{pri}}(x) < g_T^{\text{pub}}(x)$ for every $x \in (0, \gamma_T)$ (Step 3). Since the principal minimizes the sum of the agents' gain (that is, $x + g_T^{\text{pub}}(x)$ under public monitoring and $x + g_T^{\text{pri}}(x)$ under private monitoring), the above result implies that private monitoring is better for the principal than public monitoring in the T-period game.

Step 1: If $\gamma_{T-1} > \frac{1-c}{\delta}$ and $g^{\text{pri}}_{T-1}(x) < g^{\text{pub}}_{T-1}(x)$ for every $x \in (0, \gamma_{T-1})$, then $g^{\text{pri}}_{T}(x) < g^{\text{pub}}_{T}(x)$ for every $x \in (0, \gamma_T)$.

Fix $x \in (0, \gamma_{T-1})$, and let $\widehat{\sigma} = (\widehat{\sigma}_0, \widehat{\sigma}_1, \widehat{\sigma}_2)$ be a PPE in G^{pub}_T that implements the payoff $\left(v_1 = x, v_2 = g^{\text{pub}}_T(x)\right)$ and has the structure shown in Figure 23 with $W \in (0, \gamma_{T-1})$ or $Z \in (0, \gamma_{T-1})$. Denote by p_1 the probability under $\widehat{\sigma}_0$ that the inspector inspects Agent 1 in the first period.

Suppose first that $W \in (0, \gamma_{T-1})$. By the assumption in this step, $g^{\text{pri}}_{T-1}(W) < g^{\text{pub}}_{T-1}(W)$. We argue that in the game G^{pri}_T the principal can achieve an SE outcome $(v_1 = x, v_2 = y)$ for some $y < g^{\text{pub}}_T(x)$. This will prove that when $W \in (0, \gamma_{T-1})$ we have $g^{\text{pri}}_T(x) < g^{\text{pub}}_T(x)$.

Let (σ^W, μ^W), where $\sigma^W = (\sigma^W_0, \sigma^W_1, \sigma^W_2)$, be the SE in G^{pri}_{T-1} that implements the payoff vector $\left(v_1 = W, v_2 = g^{\text{pri}}_{T-1}(W)\right)$, and let (σ^Z, μ^Z), where $\sigma^Z = (\sigma^Z_0, \sigma^Z_1, \sigma^Z_2)$, be the SE in G^{pri}_{T-1} that implements the payoff vector $\left(v_1 = g^{\text{pub}}_{T-1}(Z), v_2 = Z\right)$. Consider the following strategy profile $\sigma^{\text{pri}} = (\sigma^{\text{pri}}_0, \sigma^{\text{pri}}_1, \sigma^{\text{pri}}_2)$ in G^{pri}_T (see Figure 25):

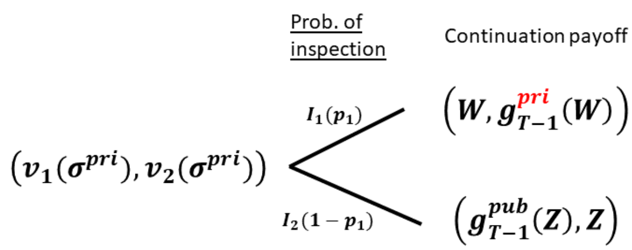

Figure 25: The first period under σ^{pri}.

- Strategy of the principal:
  - In the first period, inspect Agent 1 and Agent 2 with probabilities p_1 and $1 - p_1$, respectively.
  - If Agent 1 (resp., Agent 2) is inspected in the first period and found adhering, then from the second period on the inspector follows σ^W_0 (resp., σ^Z_0).
  - Punishment: if an agent is found violating, he is inspected with probability 1 in all future periods.
- Strategy of each agent i:
  - agent i adheres in the first period.
  - If agent i is inspected in the first period and found adhering, then from the second period on he implements σ^W_1 (if $i = 1$) or σ^Z_2 (if $i = 2$), ignoring the play in the first period.

– If agent i is not inspected in the first period, then from the second period on he implements σ_1^Z (if $i = 1$) or σ_2^W (if $i = 2$), ignoring the play in the first period.

– Punishment: if in the first period agent i is inspected and found violating, he adheres in all future periods.

We next specify a proper belief system for σ^{pri}. Since (σ^W, μ^W) is an SE, there exists a sequence of strategies where agents completely mix their actions $\{\sigma^{W,k}\}_{k\in\mathbb{N}}$ that converges to σ^W, with a corresponding sequence of beliefs $\{\mu^{W,k}\}_{k\in\mathbb{N}}$ that converges to μ^W. An analogous statement applies to (σ^Z, μ^Z). For every $k \in \mathbb{N}$, consider those plays that occur with positive probability[52] under $\sigma^{W,k}$ and $\sigma^{Z,k}$, and denote by $\epsilon_k > 0$ the minimum among all these probabilities. Since agents play pure strategies in σ^W and σ^Z, the sequence ϵ_k converges to zero as $k \to \infty$.

- Consider the following strategy profile $\sigma_i^{pri,k}$ of agent i that is completely mixed:
  – In the first period agent i adheres with probability $1 - (\epsilon_k)^2$ and violates with probability $(\epsilon_k)^2$.
  – If agent i is inspected in the first period and found adhering, then from the second period on he implements $\sigma_1^{W,k}$ (if $i = 1$) or $\sigma_2^{Z,k}$ (if $i = 2$).
  – If agent i is not inspected in the first period, then from the second period on he implements $\sigma_1^{Z,k}$ (if $i = 1$) or $\sigma_2^{W,k}$ (if $i = 2$).
  – Punishment: if in the first period agent i is inspected and found violating, he adheres with probability $\frac{k-1}{k}$ in all future periods.

By definition, $\lim_{k\to\infty} \sigma_i^{pri,k} = \sigma_i^{\text{pri}}$. The reader can verify that the sequence of beliefs $(\mu^{pri,k})_{k\in\mathbb{N}}$ induced by $(\sigma^{pri,k})_{k\in\mathbb{N}}$ converges. Denote by μ^{pri} the limit belief system. In Appendix B.2 we show that under μ^{pri}, if Agent 1 (resp., Agent 2) is inspected in the first period, then the agents' continuation beliefs under μ^{pri} coincide with μ^W (resp., μ^Z). The difficulty lies in showing that if an agent, say, Agent 1, is not inspected in the first period, then his continuation belief under μ^{pri} coincides with μ^Z.

We now argue that $(\sigma^{\text{pri}}, \mu^{\text{pri}})$ is an SE. When players follow σ^{pri}, even though the agent who is not inspected in the first period cannot observe the first-period actions of the inspector and of the other agent, he forms a correct conjecture about it: that is, the other agent is inspected and found adhering. Therefore, if Agent 1 (resp., Agent 2) is inspected in the first

[52]The agents' strategies under $\sigma^{W,k}$ and $\sigma^{Z,k}$ are completely mixed. Nevertheless, the principal may use pure actions at some histories. Therefore, under $\sigma^{W,k}$ and $\sigma^{Z,k}$, there may still exist events of probability zero.

period, both agents believe that they are following the continuation strategy σ^W (resp., σ^Z). Since (σ^W, μ^W) and (σ^Z, μ^W) are SEs from the second period and on, it is sufficient to verify that both agents are better off adhering in the first period. Indeed, since $\widehat{\sigma}$ is a PPE that satisfies Proposition 7, we have $W \geq f(p_1)$ and $Z \geq f(1-p_1)$, and hence both agents adhere in the first period of σ^{pri}.

By the definition of σ^{pri}, we have (see Figure 25)

$$v_1(\sigma^{\text{pri}}) = \delta \cdot \big(p_1 \cdot W + (1-p_1) \cdot g^{\text{pub}}_{T-1}(Z)\big) = v_1(\widehat{\sigma}) = x.$$

Since by assumption $g^{\text{pri}}_{T-1}(W) < g^{\text{pub}}_{T-1}(W)$,

$$v_2(\sigma^{\text{pri}}) = \delta \cdot \big(p_1 \cdot g^{\text{pri}}_{T-1}(W) + (1-p_1) \cdot Z\big) < \delta \cdot \big(p_1 \cdot g^{\text{pub}}_{T-1}(W) + (1-p_1) \cdot Z\big) = v_2(\widehat{\sigma}) = g^{\text{pub}}_T(x).$$

Thus, $(\sigma^{\text{pri}}, \mu^{\text{pri}})$ is an SE in G^{pri}_T that yields the payoffs $\big(v_1 = x, v_2 < g^{\text{pub}}_T(x)\big)$, as desired. This completes the proof for the case $W \in (0, \gamma_{T-1})$.

To prove the case $Z \in (0, \gamma_{T-1})$, we first argue that under private monitoring, g^{pri}_T is strictly decreasing on $(0, g^{\text{pri}}_T(0))$. Since there is no violation-free mechanism, $g^{\text{pri}}_T(0) > 0$. By the definition of g^{pri}_T, we have $g^{\text{pri}}_T\big(g^{\text{pri}}_T(0)\big) = 0$. Due to the convexity and non-negativity of g^{pri}_T, the function g^{pri}_T is strictly decreasing on $[0, \gamma^{\text{pri}}_T]$. Note that by Theorem 1, $g^{\text{pri}}_T(\gamma_T) \leq g^{\text{pub}}_T(\gamma_T) = 0$, and hence $0 < g^{\text{pri}}_T(0) \leq \gamma_T$.

Suppose that W is not in $(0, \gamma_{T-1})$ and $Z \in (0, \gamma_{T-1})$. By an argument symmetric to the previous case, under G^{pri}_T the principal can attain an SE outcome $(v_1 = x', v_2 = g^{\text{pub}}_T(x))$, where $x' < x = v_1(\widehat{\sigma})$ (see Figure 23). Since $Z \in (0, \gamma_{T-1})$ and $p_1 \in (0,1)$ (recall that both agents adhere in the first period under $\widehat{\sigma}$), we have $g^{\text{pub}}_T(x) > 0$ and $x > 0$. If $x' < g^{\text{pri}}_T(0)$, then since the function g^{pri}_T is strictly decreasing, the payoff vector $(v_1 = x, v_2 = y)$ for some $y < g^{\text{pub}}_T(x)$ can be supported as an SE outcome in G^{pri}_T, as desired.

Consider now the case $x' \geq g^{\text{pri}}_T(0)$. Since g^{pri}_T is non-increasing, $g^{\text{pri}}_T(x') \leq g^{\text{pri}}_T\big(g^{\text{pri}}_T(0)\big) = 0$. Since g^{pri}_T is non-negative, $g^{\text{pri}}_T(x') = 0$. Since $x > x'$, we have $g^{\text{pri}}_T(x) = 0$. This implies that the payoff vector $(v_1 = x, v_2 = 0)$ can be supported as an SE outcome in G^{pri}_T. Since $0 < g^{\text{pub}}_T(x)$, private monitoring improves upon public monitoring, as claimed.

Step 2: If $\delta > 1 - c$, then $g^{\text{pri}}_2(x) < g^{\text{pub}}_2(x)$ for every $x \in (0, \gamma_2)$.

As argued before, $\gamma_2 = \delta$ and hence the payoff vectors $(v_1 = 0, v_2 = \delta)$ and $(v_1 = \delta, v_2 = 0)$ can be supported as PPE outcomes in G^{pub}_2. We now argue that $g^{\text{pub}}_2(x) = \delta - x$. First notice that the payoff vectors that can be attained by a linear combination of $(v_1 = 0, v_2 = \delta)$ and $(v_1 = \delta, v_2 = 0)$ are PPE outcomes of G^{pub}_2. It is left to show that no PPE outcome in $\mathscr{E}_*(G_2^{\text{pub}})$ can yield the agents a payoff (v_1, v_2) such that $v_1 + v_2 < \delta$. Indeed, in the second period there must be an agent who is inspected with probability less than $\frac{1}{1+c}$, and consequently this agent violates. Since we consider PPEs in $\mathscr{E}_*(G_2^{\text{pub}})$, the violating agent gains 1 in the second period, and hence $v_1 + v_2 \geq \delta$, as desired.

Recall that $\gamma_2 = \delta$ applies also to private monitoring. Hence the payoff vectors $(v_1 = 0, v_2 = \delta)$ and $(v_1 = \delta, v_2 = 0)$ can also be supported as PPE outcomes in G_2^{pri}. Since g_2^{pri} is a convex function, to show that $g_2^{\text{pri}}(x) < g_2^{\text{pub}}(x)$ for every $x \in (0, \gamma_2)$, it is therefore sufficient to show that there exists an SE outcome $(\widehat{v}_1, \widehat{v}_2)$ in $\mathscr{E}_*(G_2^{\text{pri}})$ such that $\widehat{v}_1 + \widehat{v}_2 < \delta$. Such an SE was constructed in the proof of Theorem 1.

Step 3: For every $T \geq 3$, if $\gamma_{T-1} > \frac{1-c}{\delta}$, then $g_T^{\text{pri}}(x) < g_T^{\text{pub}}(x)$ for every $x \in (0, \gamma_T)$.

The proof is by induction on T. As shown in Step 2, $g_2^{\text{pri}}(x) < g_2^{\text{pub}}(x)$ for every $x \in (0, \gamma_2)$. By the assumption in Proposition 3, $\delta > \frac{1-c^2}{c}$. This implies that $\gamma_1 = 1 > \frac{1-c}{\delta}$. Since the sequence $\{\gamma_k\}_{k=1}^{T-1}$ is decreasing in k and since $\gamma_{T-1} > \frac{1-c}{\delta}$ by assumption, we have $\gamma_k > \frac{1-c}{\delta}$ for every $k \leq T-1$. Using Step 1, it follows by induction that $g_k^{\text{pri}}(x) < g_k^{\text{pub}}(x)$ for every $3 \leq k \leq T$.

B.2 The consistency of beliefs

In Figure 25 in the proof of Proposition 3, we construct a sequence of strategies and belief systems $(\sigma^{pri,k}, \mu^{pri,k})_{k \in \mathbb{N}}$. As $k \to \infty$, the limit strategy and belief system $(\sigma^{\text{pri}}, \mu^{\text{pri}})$ exists and the two components are consistent with each other. In this section we show that under μ^{pri}, if Agent 1 (resp., Agent 2) is inspected in the first period, the agents' continuation beliefs under μ^{pri} coincide with μ^W (resp., μ^Z).

Let h be a private history of, say, Agent 1. In this section we show that $\mu^{\text{pri}}(A_1, h) = \mu^W(h)$, and $\mu^{\text{pri}}(N_1, h) = \mu^Z(h)$, where A_1 represents the history that Agent 1 is inspected and found adhering, and N_1 represents the history that Agent 1 is not inspected. The argument for Agent 2's beliefs is similar.

Fix $k \in \mathbb{N}$ and consider the strategy profile $\sigma^{pri,k}$. Suppose that Agent 1 is inspected and found adhering in the first period. In this case, from the second period on Agent 1 implements the strategy $\sigma_1^{W,k}$, and he assigns probability 1 to the event that Agent 2 implements $\sigma_2^{W,k}$. Agent 1's belief from the second period under $\mu^{pri,k}$ induced by $\sigma^{pri,k}$ is hence the same as the belief $\mu^{W,k}$ induced by $\sigma^{W,k}$, and therefore the limit belief when $k \to \infty$ is the same as well: $\mu^{\text{pri}}(A_1, h) = \mu^W(h)$.

We next consider the case where Agent 1 is not inspected in the first period. Denote by $\mathcal{I}$ the information set of Agent 1 that contains the private history h, and suppose that it can be reached with positive probability under $\sigma^{Z,k}$. Note that, in fact, for all $k \in \mathbb{N}$, the information sets that can be reached with positive probability under $\sigma^{Z,k}$ are the same. Under $(\sigma^{pri,k}, \mu^{pri,k})$, if Agent 1 is not inspected in the first period, the probability he assigns to a play $a \in \mathcal{I}$ is

$$P_{\sigma^{pri,k}}\big((N_1, a)|(N_1, \mathcal{I})\big) = \frac{P(A_2) \cdot P_{\sigma^{Z,k}}(a) + P(V_2) \cdot P_{\sigma^{pun,k}}(a)}{P(A_2) \cdot \sum_{a' \in \mathcal{I}} P_{\sigma^{Z,k}}(a) + P(V_2) \cdot \sum_{a' \in \mathcal{I}} P_{\sigma^{pun,k}}(a)}, \tag{31}$$

where $P(A_2) = 1 - (\epsilon_k)^2$ (resp., $P(V_2) = (\epsilon_k)^2$) is the probability that Agent 2 is inspected and found adhering (resp., violating), $P_{\sigma^{Z,k}}(a)$ is the probability that play a occurs under strategy profile $\sigma^{Z,k}$, and $\sigma^{pun,k}(a)$ is the probability that play a occurs when Agent 2 is being inspected with probability 1 in all periods (*pun* stands for *punishment*).

As a comparison, under $\sigma^{Z,k}$, given that Agent 1 is at information set $\mathcal{I}$, the probability he assigns to play a is

$$P_{\sigma^{Z,k}}(a, \mathcal{I}) = \frac{P_{\sigma^{Z,k}}(a)}{\sum_{a' \in \mathcal{I}} P_{\sigma^{Z,k}}(a)}. \tag{32}$$

As argued before, $\epsilon_k \to 0$ as $k \to \infty$. Therefore, in the right-hand-side of Eq. (31), the term $P(V_2)$ goes to zero as $k \to \infty$. Nevertheless, this is not enough to show that (31) converges to (32), since the term $P_{\sigma^{Z,k}}$ may also go to zero. However, recall that the magnitude of $P_{\sigma^{Z,k}}(a)$ is at least ϵ_k, whereas we choose $P(V_2)$ to be $(\epsilon_k)^2$. Consequently, the right-hand side of Eq. (31) converges to the right-hand side of Eq. (32), and hence in the limit, $\mu^{\text{pri}}(N_1, h) = \mu^Z(h)$, as desired.

Note that in this proof we only deal with information sets that can be reached with positive probability under $\mu^{W,k}$ or $\mu^{Z,k}$. For events in the information sets that occur with probability zero under, say, $\mu^{W,k}$ (this happens because the principal may use pure strategies under some histories), the agents' beliefs at these events under μ^W are irrelevant and hence can be selected arbitrarily. In particular, they can be made consistent with the beliefs under μ^{pri}.

B.3 Proof of Proposition 7

We first define a function $\widehat{g}_T^{\text{pub}}$ that is similar to g_T^{pub}, but excludes the use of the correlation device in the first period. Recall that $\Delta_T = 1 + \delta + \delta^2 + \cdots + \delta^{T-1}$ is an agent's maximal equilibrium payoff in the T-period game.

Definition 2. For every $x \in [0, \Delta_T]$, let $\widehat{g}_T^{\text{pub}}(x)$ denote the minimum payoff of Agent 1 over all PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ that yield Agent 2 the payoff x and do not use the correlation device in the first period.

The function g_T^{pub} is the largest convex function that is smaller than $\widehat{g}_T^{\text{pub}}$. Therefore, to prove Proposition 7, it is sufficient to prove that $\widehat{g}_T^{\text{pub}}$ agrees with g_T^{pub}, and that under $\widehat{g}_T^{\text{pub}}$ both agents adhere in the first period. To this end we will study some properties of the function $\widehat{g}_T^{\text{pub}}$. Obviously, $\widehat{g}_T^{\text{pub}}$ is non-negative, since if an agent adheres in every period, he guarantees a payoff 0.

Denote by $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ the set of all PPEs in $\mathscr{E}_*(G_T^{\text{pub}})$ where no correlation device is used in the first period.

Step 1: Let $x \in (0, \gamma_T)$. To implement the outcome $\big(v_1 = x, v_2 = \widehat{g}_T^{\text{pub}}(x)\big)$ by a PPE in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$, both agents adhere in the first period.

We first argue that $x < 1$. Indeed, since $T \geq 2$, we have $\gamma_T \leq \delta$ (see Proposition 6), hence, $x < \gamma_T \leq \delta < 1$. We now argue that $\widehat{g}_T^{\text{pub}}(x) < 1$. To this end we construct a PPE σ in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ that yields the payoffs $(v_1 = x, v_2 = \gamma_T)$. This will imply that $\widehat{g}_T^{\text{pub}}(x) \leq \gamma_T \leq \delta < 1$.

Step 1 in the proof of Proposition 6 shows that $\widehat{g}_T^{\text{pub}}(0) = \gamma_T$. Therefore, there exists a PPE $\widetilde{\sigma}$ in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ that implements the payoffs $(v_1 = 0, v_2 = \gamma_T)$. Under $\widetilde{\sigma}$, Agent 1 adheres in all periods. By an argument similar to that in the proof of Proposition 5, we can modify histories after the first period and construct another equilibrium that yields Agent 1 a payoff x, $0 < x \leq \gamma_T \leq \delta$. Under this new equilibrium, the correlation device in not used in the first period, and the payoff of Agent 2 remains unchanged. This observation implies that for every $x \in (0, \gamma_T)$, the payoffs $(v_1 = x, v_2 = \gamma_T)$ can be implemented by a PPE in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$. By the definition of $\widehat{g}_T^{\text{pub}}(x)$, we have $\widehat{g}_T^{\text{pub}}(x) \leq \gamma_T \leq \delta < 1$, as desired.

Since under PPEs in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$, a violation on the equilibrium path yields an agent a payoff 1, because no correlation device is used, and since $\widehat{g}_T^{\text{pub}}(x) < 1$, in any PPE that implements $\big(v_1 = x, v_2 = \widehat{g}_T^{\text{pub}}(x)\big)$, both agents must adhere in the first period.

Step 2: The function $\widehat{g}_T^{\text{pub}}(x)$ is non-increasing.

By Step 1, to implement the outcome $\big(v_1 = x, v_2 = \widehat{g}_T^{\text{pub}}(x)\big)$ by a PPE in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$, both agents adhere in the first period. Therefore, by an argument similar to that in the proof of Proposition 5, we can modify histories after the first period and construct another equilibrium that yields Agent 1 a payoff x', for every $x < x' \leq \gamma_T \leq \delta$, and keep Agent 2's payoff $\widehat{g}_T^{\text{pub}}(x)$. Under this new equilibrium no correlation device is used in the first period, and hence $\widehat{g}_T^{\text{pub}}(x') \leq \widehat{g}_T^{\text{pub}}(x)$. This completes the proof that the function $\widehat{g}_T^{\text{pub}}(x)$ is non-increasing.

We have shown that the function $\widehat{g}_T^{\text{pub}}$ is non-negative and non-increasing on $[0, \gamma_T]$,

$\widehat{g}_T^{\text{pub}}(0) = \gamma_T$, and $\widehat{g}_T^{\text{pub}}(\gamma_T) = 0$. We next prove that $\widehat{g}_T^{\text{pub}}$ is convex on $(0, \gamma_T)$, and this will imply that $\widehat{g}_T^{\text{pub}}$ agrees with g_T^{pub}.

Step 3: The function $\widehat{g}_T^{\text{pub}}$ is convex on $(0, \gamma_T)$.

To show this, we will prove that

$$\widehat{g}_T^{\text{pub}}\big(q \cdot x + (1-q) \cdot y\big) \leq q \cdot \widehat{g}_T^{\text{pub}}(x) + (1-q) \cdot \widehat{g}_T^{\text{pub}}(y),$$

for every $x, y \in (0, \gamma_T)$ and every $0 < q < 1$.

Let σ be a PPE in $\mathscr{E}_*(G_T^{\text{pub}})$ that uses the correlation device in the first period and plays as follows: with probability q the players implement a PPE in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ that yields the payoffs $\big(v_1 = x, v_2 = \widehat{g}_T^{\text{pub}}(x)\big)$, and with probability $1-q$ they implement a PPE in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ that yields the payoffs $\big(v_1 = y, v_2 = \widehat{g}_T^{\text{pub}}(y)\big)$. Then, $v_1(\sigma) = qx + (1-q)y$ and $v_2(\sigma) = q \cdot \widehat{g}_T^{\text{pub}}(x) + (1-q) \cdot \widehat{g}_T^{\text{pub}}(y)$.

We will show that there exists a PPE $\overline{\sigma}$ in $\mathscr{E}_*^{\text{nc}}(G_T^{\text{pub}})$ such that $v_1(\overline{\sigma}) \leq v_1(\sigma)$ and $v_2(\overline{\sigma}) \leq v_2(\sigma)$. This observation, together with the fact that $\widehat{g}_T^{\text{pub}}$ is non-increasing (see Step 2), imply that

$$\begin{aligned}\widehat{g}_T^{\text{pub}}\big(\overbrace{q \cdot x + (1-q) \cdot y}^{v_1(\sigma)}\big) &= \widehat{g}_T^{\text{pub}}\big(v_1(\sigma)\big)\\ &\leq \widehat{g}_T^{\text{pub}}\big(v_1(\overline{\sigma})\big) \leq v_2(\overline{\sigma}) \leq v_2(\sigma) = q \cdot \widehat{g}_T^{\text{pub}}(x) + (1-q) \cdot \widehat{g}_T^{\text{pub}}(y),\end{aligned}$$

and hence the function $\widehat{g}_T^{\text{pub}}$ is convex.

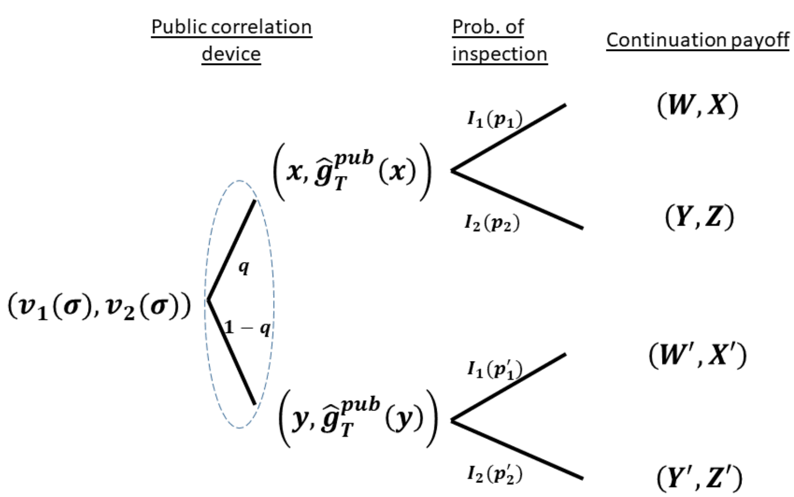

Figure 26: The first period under σ.

Suppose $\sigma \in \mathscr{E}_*(G_T^{\text{pub}})$. The play under σ in the first period is as follows (see Figure 26):

- The PPE that implements $\big(v_1 = x, v_2 = \widehat{g}_T^{\text{pub}}(x)\big)$ has the following structure:
  - In the first period, Agent 1 and Agent 2 are inspected with probabilities p_1 and p_2, respectively. Both agents adhere.
  - If Agent 1 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G_{T-1}^{pub} that yields the payoffs $(v_1 = W, v_2 = X)$.
  - If Agent 2 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G_{T-1}^{pub} that yields the payoffs $(v_1 = Y, v_2 = Z)$.
  - Punishment: If in the first period an agent is found violating, then in all future periods he is inspected with probability 1, he adheres, and the other agent violates.
- The PPE that implements $\big(v_1 = y, v_2 = \widehat{g}_T^{\text{pub}}(y)\big)$ has the following structure:
  - In the first period, inspect Agent 1 and Agent 2 with probabilities p_1' and p_2', respectively. Both agents adhere.
  - If Agent 1 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G_{T-1}^{pub} that yields the payoffs $(v_1 = W', v_2 = X')$.
  - If Agent 2 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G_{T-1}^{pub} that yields the payoffs $(v_1 = Y', v_2 = Z')$.
  - Punishment: If in the first period an agent is found violating, then in all future periods he is inspected with probability 1, he adheres, and the other agent violates.

By Step 1, both agents adhere in the first period, and hence $W \geq f(p_1)$, $W' \geq f(p_1')$, $Z \geq f(p_2)$, and $Z' \geq f(p_2')$. To make the agents' expected payoffs agree with their continuation payoffs, we must have

$$x = \delta \cdot (p_1 W + p_2 Y), \quad \widehat{g}_T^{\text{pub}}(x) = \delta \cdot (p_1 X + p_2 Z), \tag{33}$$

and

$$y = \delta \cdot (p_1' W' + p_2' Y'), \quad \widehat{g}_T^{\text{pub}}(y) = \delta \cdot (p_1' X' + p_2' Z'). \tag{34}$$

Set $\overline{p}_1 := q \cdot p_1 + (1-q) \cdot p_1'$ and $\overline{p}_2 := q \cdot p_2 + (1-q) \cdot p_2'$. Note that $\overline{p}_1 + \overline{p}_2 = 1$. Consider the following strategy profile $\overline{\sigma}$ (see Figure 27):

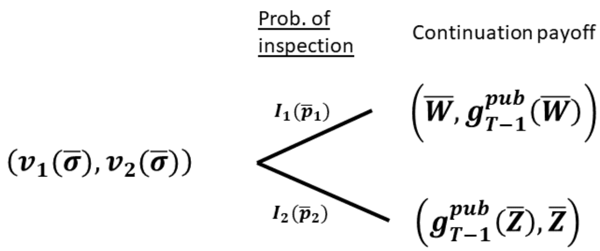

Figure 27: The first period under $\overline{\sigma}$.

- In the first period, Agent 1 and Agent 2 are inspected with probabilities $\overline{p}_1$ and $\overline{p}_2$, respectively. Both agents adhere.
- If Agent 1 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G^{pub}_{T-1} that yields the payoffs $(v_1 = \overline{W}, v_2 = g^{\text{pub}}_{T-1}(\overline{W}))$, where $\overline{W} := \frac{q_1 p_1}{q_1 p_1 + q_2 p'_1} W + \frac{q_2 p'_1}{q_1 p_1 + q_2 p'_1} W'$.
- If Agent 2 is inspected in the first period and found adhering, then from the second period on the players implement a PPE profile in G^{pub}_{T-1} that yields the payoffs $(v_1 = g^{\text{pub}}_{T-1}(\overline{Z}), v_2 = \overline{Z})$, where $\overline{Z} := \frac{q_1 p_2}{q_1 p_2 + q_2 p'_2} Z + \frac{q_2 p'_2}{q_1 p_2 + q_2 p'_2} Z'$.
- Punishment: If in the first period an agent is found violating, then in all future periods he is inspected with probability 1, he adheres, and the other agent violates.

Let us verify that $\overline{\sigma}$ is a PPE. The crucial part is to show that both agents are better off adhering in the first period of $\overline{\sigma}$. And indeed, since $W \geq f(p_1)$, $W' \geq f(p'_1)$, and since the function $p \mapsto pf(p)$ is linear,

$$\begin{aligned}(q_1 p_1 + q_2 p'_1)\overline{W} &= q_1 p_1 W + q_2 p'_1 W' \\ &\geq q_1 p_1 f(p_1) + q_2 p'_1 f(p'_1) \\ &= (q_1 p_1 + q_2 p'_1) \cdot f(q_1 p_1 + q_2 p'_1),\end{aligned} \tag{35}$$

and hence $\overline{W} \geq f(q_1 p_1 + q_2 p'_1) = f(\overline{p}_1)$. By Proposition 1(i), Agent 1 is better off adhering in the first period of $\overline{\sigma}$, as desired. A similar argument shows that $\overline{Z} \geq f(\overline{p}_2)$, and Agent 2 is better off adhering in the first period of $\overline{\sigma}$ as well.

We next show that $v_1(\overline{\sigma}) \leq v_1(\sigma)$ and $v_2(\overline{\sigma}) \leq v_2(\sigma)$. Because of the convexity of g^{pub}_{T-1}, and since $X \geq g^{\text{pub}}_{T-1}(W)$ and $X' \geq g^{\text{pub}}_{T-1}(W')$,

$$\begin{aligned} g^{\text{pub}}_{T-1}(\overline{W}) &\leq \frac{q_1 p_1}{q_1 p_1 + q_2 p'_1} \cdot g^{\text{pub}}_{T-1}(W) + \frac{q_2 p'_1}{q_1 p_1 + q_2 p'_1} \cdot g^{\text{pub}}_{T-1}(W') \\ &\leq \frac{q_1 p_1}{q_1 p_1 + q_2 p'_1} \cdot X + \frac{q_2 p'_1}{q_1 p_1 + q_2 p'_1} \cdot X'. \end{aligned} \tag{36}$$

Similarly,

$$g^{\text{pub}}_{T-1}(\overline{Z}) \leq \frac{q_1 p_2}{q_1 p_2 + q_2 p'_2} \cdot Y + \frac{q_2 p'_2}{q_1 p_2 + q_2 p'_2} \cdot Y'. \tag{37}$$

By Eqs. (33) and (34),

$$\begin{aligned} v_1(\sigma) &= q_1 \cdot x + q_2 \cdot y \\ &= q_1 \cdot \delta(p_1 W + p_2 Y) + q_2 \cdot \delta(p'_1 W' + p'_2 Y') \\ &= \delta \cdot \big(q_1 p_1 W + q_1 p_2 Y + q_2 p'_1 W' + q_2 p'_2 Y'\big). \end{aligned} \tag{38}$$

Next, by the definitions of $\overline{W}$ and $\overline{Z}$, the convexity of g^{pub}_{T-1}, and Eq. (37),

$$\begin{aligned} v_1(\overline{\sigma}) =& \,\delta \cdot \big(\overline{p}_1 \cdot \overline{W} + \overline{p}_2 \cdot g^{\text{pub}}_{T-1}(\overline{Z})\big) \\ =& \,\delta \cdot (q_1 p_1 + q_2 p'_1) \cdot \left(\frac{q_1 p_1}{q_1 p_1 + q_2 p'_1} \cdot W + \frac{q_2 p'_1}{q_1 p_1 + q_2 p'_1} \cdot W'\right) \\ &+ \delta \cdot (q_1 p_2 + q_2 p'_2) \cdot g^{\text{pub}}_{T-1}\left(\overline{Z}\right) \\ \leq& \,\delta \cdot (q_1 p_1 \cdot W + q_2 p'_1 \cdot W') \\ &+ \delta \cdot (q_1 p_2 + q_2 p'_2) \cdot \left(\frac{q_1 p_2}{q_1 p_2 + q_2 p'_2} \cdot Y + \frac{q_2 p'_2}{q_1 p_2 + q_2 p'_2} \cdot Y'\right) \\ =& \,\delta \cdot \big(q_1 p_1 W + q_1 p_2 Y + q_2 p'_1 W' + q_2 p'_2 Y'\big) \\ =& \,v_1(\sigma). \end{aligned} \tag{39}$$

By a similar argument, $v_2(\overline{\sigma}) \leq v_2(\sigma)$.

Since under $\overline{\sigma}$ no correlation device is used in the first period, $\big(v_1(\overline{\sigma}), v_2(\overline{\sigma})\big)$ is a PPE outcome in $\mathscr{E}^{\text{nc}}_*(G^{\text{pub}}_T)$. Since $v_1(\sigma) \geq v_1(\overline{\sigma})$ and by Step 2, the payoff vector $\big(v_1(\sigma), v_2(\overline{\sigma})\big)$ is a PPE outcome in $\mathscr{E}^{\text{nc}}_*(G^{\text{pub}}_T)$, and hence $\widehat{g}^{\text{pub}}_T\big(v_1(\sigma)\big) \leq v_2(\overline{\sigma}) \leq v_2(\sigma)$. That is, for every $q_1 \in (0,1)$,

$$\widehat{g}^{\text{pub}}_T\big(q_1 \cdot x + (1-q_1) \cdot y\big) \leq q_1 \cdot \widehat{g}^{\text{pub}}_T(x) + (1-q_1) \cdot \widehat{g}^{\text{pub}}_T(y). \tag{40}$$

Consequently, for every $T \geq 2$, the function $\widehat{g}^{\text{pub}}_T$ is convex, as desired.