



Robustness of public equilibria in repeated games with private monitoring [☆]

Richard McLean ^a, Ichiro Obara ^{b,*}, Andrew Postlewaite ^c

^a Rutgers University, United States

^b University of California, Los Angeles, United States

^c University of Pennsylvania, United States

Received 10 April 2014; accepted 25 June 2014

Available online 8 July 2014

Abstract

A repeated game with private monitoring is “close” to a repeated game with public monitoring (or perfect monitoring) when (i) the expected payoff structures are close and (ii) the informational structures are close in the sense that private signals in the private monitoring game can be aggregated by some public coordination device to generate a public signal whose distribution is close to the distribution of the public signal in the public monitoring game. We provide a sufficient condition for the set of uniformly strict perfect public equilibria for a public monitoring game to be robust in nearby private monitoring games in the sense that they remain equilibria with respect to the public signal that is generated by such public coordination devices with truthful reporting. Our sufficient condition requires that every player is informationally small in a well-defined sense.

© 2014 Elsevier Inc. All rights reserved.

JEL classification: C72; C73; D82

Keywords: Communication; Informational size; Private monitoring; Public monitoring; Repeated games; Robustness

[☆] We are grateful to the associate editor and three anonymous referees for extensive comments. We thank George Mailath for very helpful discussions and the audience in numerous conferences and seminars for helpful comments. This paper was originally a part of our paper “Informational Smallness and Private Monitoring in Repeated Games” [25]. Postlewaite gratefully acknowledges support from National Science Foundation grant SBR-9810693.

* Corresponding author.

1. Introduction

Cooperation within groups is an important and commonly observed social phenomenon, but the way in which cooperation arises is one of the least understood questions in economics. The theory of repeated games has improved our understanding by showing how coordinated threats to punish can prevent deviations from cooperative behavior, but much of the work in repeated games rests on very restrictive assumptions that all players share the same public information either perfectly or imperfectly. For the case in which each player can observe all other players' actions directly (perfect monitoring), Aumann and Shapley [5] and Rubinstein [34] proved a folk theorem without discounting, and Fudenberg and Maskin [13] proved a folk theorem with discounting. For the case in which each player observes a noisy public signal (imperfect public monitoring), Abreu, Pearce and Stacchetti [1] characterized the set of pure strategy sequential equilibrium payoffs and Fudenberg, Levine and Maskin [14] proved a folk theorem.

But a theory that rests on the assumption that there is common knowledge of a sufficient statistic about all past behavior is, at best, incomplete. Such a theory is of little help in understanding behavior in groups in which there are idiosyncratic errors in individuals' observations of outcomes.¹ For many problems, it is more realistic to consider players who possess only partial information about the environment and, most importantly, players may not know the information possessed by other players. In such problems, players may communicate their partial information to other players in order to build a "consensus" regarding the current situation, which can be used to coordinate their future behavior. In this view, repeated games with public information can be thought of as a reduced form of a more complex interaction involving information sharing.

This point of view leads us to examine the *robustness* of equilibria with public monitoring when monitoring is private, but "close" to public monitoring. For example, one can think of a situation in which information contained in the public signal is dispersed among the players in the form of noisy private signals. If the amount of information contained in each player's private signal is negligible, then it is tempting to consider the game with such private signals and the underlying game with public signals as being "close." In this paper, we examine whether an equilibrium with public monitoring remains an equilibrium with respect to a public signal generated from private monitoring and communication, and whether (and how) players can be induced to reveal their private information.

To make these ideas precise, consider a public monitoring game (G, π) and a private monitoring game (G', p) , where G and G' are normal form games with public monitoring and private monitoring respectively. In (G, π) , each action profile a generates a public signal y from a set Y with probability $\pi(y|a)$. In (G', p) , each action profile a generates a private signal profile $s = (s_1, \dots, s_n)$ with probability $p(s|a)$. In our analysis of the private monitoring game (G', p) , we will augment the model with a "public coordination device" ϕ that chooses a public coordinating signal (possibly randomly) from Y based on the reported profile of private signals. In this expanded game, players choose an action profile a , observe their private signals (s_1, \dots, s_n) , and publicly announce the (not necessarily honest) profile (s'_1, \dots, s'_n) . A public coordinating signal $y \in Y$ is then selected with probability $\phi(y|s'_1, \dots, s'_n)$. If the players report their private signals truthfully, then the probability that the realized public coordinating signal is y given a and ϕ is equal to $p^\phi(y|a) = \sum_{s \in S} \phi(y|s) p(s|a)$. We say that (G, π) and (G', p) are close when G and G' are close in terms of expected payoffs and there exists a public coordinating device ϕ such

¹ For example, team production in which each individual observes the outcome with error lies outside this framework.

that $\pi(y|a) \approx p^\phi(y|a)$. Consider the private monitoring repeated game associated with (G', p) and augment it with such a public coordination device ϕ_{h_t} given each history of publicly coordinating signals h_t . We call this augmented game a communication extension of the repeated game associated with (G', p) . We then ask the following question: for a given perfect public equilibrium α^* of the repeated game associated with (G, π) and given the repeated game associated with any “close” private monitoring game (G', p) , can we find a communication extension and an equilibrium in which players (i) truthfully reveal their signals along the equilibrium path and (ii) choose their actions as a function of the history of public coordinating signals according to α^* ?

If revelation constraints can be ignored and players are assumed to announce their private signals truthfully, then the private monitoring game $(G', p^{\phi_{h_t}})$ is really a new public monitoring game for each h_t . In this case, it is straightforward to show that α^* is an equilibrium in any communication extension when G' is close to G , $p^{\phi_{h_t}}$ is close to π at each h_t , and α^* is a uniformly strict equilibrium of the original repeated game (G, π) . Hence our analysis is mainly concerned with revelation constraints. The revelation of private information can be problematic, as can be seen in a simple trigger strategy equilibrium to support collusion. For a trigger strategy equilibrium to work, it is essential that every player reports “bad” outcomes honestly. However it is clear that players will not want to reveal any private information that may trigger mutual punishment.

We find that the following two concepts are the key to deal with the revelation constraints: *informational size* and *distributional variability*. Roughly speaking, player i is informationally small if for each action profile a , her private information is unlikely to have a large effect on the distribution of the public coordinating signal $p^\phi(\cdot|a)$. Consequently, small informational size will imply that she will have little incentive to misreport her private signal in order to manipulate the other players' behavior to her advantage. Players are naturally informationally small in numerous settings. Suppose, for example, that there are many players whose signals are noisy observations of an underlying (but unobserved) common signal, and that these noisy observations are conditionally i.i.d. If ϕ maps each signal profile into the posterior distribution of the unobserved signal, then each player is informationally small by the law of large numbers. Alternatively, with the same function ϕ , agents receiving conditionally i.i.d. signals of the unobserved signal would be informationally small if their signals are very precise, even if the number of players is small (but at least three). Distributional variability is an index that measures the correlation between a player's private signal and the public coordinating signal which she would expect when she reports her signal truthfully. If this index is large, that means that a player's conditional belief about the public coordinating signal varies widely with respect to her private information. The larger this index is, the easier it is to detect and punish a lie. With these concepts, our result can be stated as follows: a uniformly strict equilibrium is robust when, for some public coordination device, (1) (G', p) is close to (G, π) in the sense described above and (2) each player's informational size is small relative to her distributional variability.

The way to induce honest reporting is roughly as follows. If ϕ is employed in every period, then $p^\phi(y|a)$ is always close to $\pi(y|a)$, but players may have an incentive to send false reports. To address this, we employ different public coordinating devices at different public histories $\{\phi_{h^t} | h^t \in H\}$, where each ϕ_{h^t} is a perturbation of ϕ . When every player's informational size is small relative to her distributional variability, we can construct the collection $\{\phi_{h^t} | h^t \in H\}$ so that every revelation constraint is satisfied on the equilibrium path (i.e. after the equilibrium action $\alpha^*(h^t)$ is played in the same period), while keeping each perturbation small enough so that the incentive to play the equilibrium action specified by α^* is not altered.

The model is described in Section 2. The notions of informational size and distributional variability are introduced in Section 3. Section 4 proves our robustness result. Section 5 discusses the related literature. All the proofs are relegated to Appendix A.

2. Preliminaries

2.1. Repeated games with public monitoring

The set of players is $N = \{1, \dots, n\}$. Player i chooses an action from a finite set A_i . An action profile is denoted by $a = (a_1, \dots, a_n) \in \times_i A_i := A$. Actions are not publicly observable, but the players observe a public signal from a finite set Y . The probability that $y \in Y$ is realized given $a \in A$ is denoted $\pi(y|a)$. We do not assume full support. That is, the set $\{y \in Y | \pi(y|a) > 0\}$ can depend on $a \in A$. This allows for perfect monitoring ($Y = A$ and $\pi(y|a) = 1$ if $y = a$) as a special case. Player i 's stage game payoff is $u_i(a_i, y)$ and player i 's expected stage game payoff is $g_i(a) = \sum_y u_i(a_i, y)\pi(y|a)$. Consequently, players do not obtain any additional information regarding the actions of other players from realized payoffs. This stage game is denoted by (G, π) , where $G = (N, A, g)$.

We now construct the repeated game associated with (G, π) . A private history for player i at stage t is denoted $h_i^t = (a_i^0, \dots, a_i^{t-1}) \in H_i^t = A_i^t$ while a public history is denoted $h^t = (y^0, \dots, y^{t-1}) \in H^t = Y^t$ with $H_i^0 = H^0 := \{\emptyset\}$. A pure strategy for player i is a sequence $\alpha_i = \{\alpha_i^t\}_{t=0}^\infty$, where α_i^t is a mapping from $H_i^t \times H^t$ to A_i . The set of pure strategies for player i is denoted Σ_i . We restrict ourselves to pure strategies. A strategy profile is denoted $\alpha = \{\alpha_i\}_{i \in N} \in \Sigma := \times_i \Sigma_i$. A strategy profile induces a probability measure on A^∞ . Player i 's discounted average payoff given α and $\delta \in (0, 1)$ is $w_i^{\alpha, \delta} = (1 - \delta) \sum_{t=0}^\infty \delta^t E[g_i(\bar{a}^t) | \alpha]$.² We denote this repeated game associated with (G, π) by $G_\pi^\infty(\delta)$.

A strategy profile is *public* if it only depends on H^t . A profile of public strategies is a *perfect public equilibrium* (PPE) if, after every public history, the continuation (public) strategy profile constitutes a Nash equilibrium (Fudenberg, Levine, and Maskin [14]). Note that a perfect public equilibrium is a subgame perfect equilibrium when the stage game is one of perfect monitoring. Since we focus on perfect public equilibrium, we will omit the dependence of strategies on private histories and write $\alpha_i^t(h^t)$ instead of $\alpha_i^t(h_i^t, h^t)$.³

Given α, δ and a public history $h^{t+1} = (h^t, y) \in H^{t+1} = H^t \times Y$, let $w_i^{\alpha, \delta}(h^t, y)$ denote player i 's continuation payoff from period $t + 1$. We define η -uniformly strict perfect public equilibrium (η -USPPE) as follows.

Definition 1. A pure strategy perfect public equilibrium $\alpha \in \Sigma$ for $G_\pi^\infty(\delta)$ is η -uniformly strict if

$$(1 - \delta)g_i(\alpha^t(h^t)) + \delta \sum_{y \in Y} \pi(y | \alpha^t(h^t))w_i^{\alpha, \delta}(h^t, y) - \eta$$

$$\geq (1 - \delta)g_i(a'_i, \alpha_{-i}^t(h^t)) + \delta \sum_{y \in Y} \pi(y | a'_i, \alpha_{-i}^t(h^t))w_i^{\alpha, \delta}(h^t, y)$$

for all $h^t \in H^t, t \geq 0, a'_i \neq \alpha_i^t(h^t)$, and $i \in N$.

² When x is a generic outcome of some random variable, we often use \tilde{x} to denote this random variable.

³ Thus we ignore *private strategies* (Kandori and Obara [18], Mailath, Matthews and Sekiguchi [20]).

This means that, at any public history, any player would lose at least η by any unilateral one-shot deviation. This is stronger than requiring all incentive constraints to hold strictly.

2.2. The repeated game with private monitoring and its public communication extension

Fix a stage game (G, π) with public monitoring as defined in the previous section. Consider a *private monitoring game* with the same set of players and the same action sets as those of G . Player i observes a private signal s_i from a finite set S_i instead of the public signal. A private signal profile is denoted $s = (s_1, \dots, s_n) \in \times_i S_i := S$. Player i 's stage game payoff is $v_i(a_i, s_i)$ and player i 's expected stage game payoff is $g'_i(a) = \sum_s v_i(a_i, s_i)p(s|a)$ where the conditional distribution on S given a is denoted $p(\cdot|a)$. We assume that the marginal distributions have full support, that is, $p(s_i|a) := \sum_{s_{-i}} p(s_i, s_{-i}|a) > 0$ for all $s_i \in S_i, a \in A$ and $i \in N$. Let $p(s_{-i}|a, s_i) := \frac{p(s_i, s_{-i}|a)}{p(s_i|a)}$ denote the conditional probability of $s_{-i} \in S_{-i}$ given (a, s_i) . We denote this private monitoring stage game by (G', p) , where $G' = (N, A, g')$. Discounted average payoffs are defined as in the public monitoring case. Let $G'^\infty(\delta)$ be the corresponding repeated game with private monitoring given $\delta \in (0, 1)$.

Players communicate directly each period.⁴ At the end of each period, players publicly announce a profile $s \in S$.⁵ Then, a public coordinating device $\phi : S \rightarrow \Delta(Y)$ generates a public signal $y \in Y$ with probability $\phi(y|s)$. A convex combination of two public coordination devices ϕ and ϕ' is denoted by $(1 - \lambda)\phi + \lambda\phi'$ and is defined by

$$((1 - \lambda)\phi + \lambda\phi')(y|s) := (1 - \lambda)\phi(y|s) + \lambda\phi'(y|s).$$

The distribution of the public signal generated by ϕ given a with honest reporting is denoted by

$$p^\phi(y|a) := \sum_{s \in S} \phi(y|s)p(s|a)$$

and the analogous conditional distribution given (a, s_i) is denoted by

$$p^\phi(y|a, s_i) := \sum_{s_{-i} \in S_{-i}} \phi(y|s_i, s_{-i})p(s_{-i}|a, s_i).$$

Next, let R_i be the set of all maps $r_i : S_i \rightarrow S_i$. The distribution of the generated public signal given a and $r_i \in R_i$ for player i and honest reporting for the other players is denoted by

$$p^\phi(y|a, r_i) := \sum_{s \in S} \phi(y|r_i(s_i), s_{-i})p(s|a).$$

We denote expectation with respect to $p^\phi(\cdot|a, r_i)$ by $E^\phi[\cdot|a, r_i]$ and expectation with respect to $p^\phi(y|a)$ simply by $E^\phi[\cdot|a]$.

⁴ This formulation of communication is very special. A more general communication structure would allow for a mediator who receives and sends confidential private information from and to the players (see Forges [11] and Myerson [29]). However, we do not need such a general communication structure for our robustness result. We ask when a perfect public equilibrium for $G^\infty_\pi(\delta)$ remains a perfect public equilibrium when players are engaged in a ‘‘close’’ private monitoring game (G', p) augmented with communication. As part of our notion of ‘‘closeness,’’ we require that private signals in (G', p) can be aggregated so as to generate a public signal whose distribution is close to π . Since public strategies can only depend on this public signal by definition, there is no role for confidential announcements or confidential recommendations.

⁵ Announcements do not have to be public. Our result holds even with private announcements.

In the repeated game (G', p) augmented with communication as described above, a public history in period t consists of a sequence of realized public signals $h^t \in Y^t$ and a sequence of public announcements $h^t_R \in S^t$. We allow different coordinating devices to be employed at different $h^t \in Y^t$. Given a private monitoring game (G', p) , a *public communication device* for (G', p) is a collection $\Phi = \{\phi_{h^t} : h^t \in Y^t, t \geq 0, \}$ where each $\phi_{h^t} : S \rightarrow \Delta(Y)$ is a public coordination device. Given a private monitoring game (G', p) , a discount factor δ , and a public communication device Φ , let $G_p^\infty(\delta, \Phi)$ denote the *public communication extension* of the repeated game with private monitoring $G_p^\infty(\delta)$.

In $G_p^\infty(\delta, \Phi)$, play proceeds in the following way. At the beginning of period t , player i chooses an action contingent on (h_i^t, h^t, h_R^t) , where $h_i^t \in A_i^t \times S_i^t$ is a sequence of her own private actions and private signals. If the resulting action profile is a , then players receive private signals according to $p(\cdot|a)$. Let s denote the realized signal profile. Then player i makes a public announcement s_i^t contingent on $(h_i^t, h^t, h_R^t, a_i, s_i)$. Of course, s_i^t may differ from s_i . Let $s' \in S$ denote the profile of announcements. Then a public signal is chosen according to $\phi_{h^t}(\cdot|s')$.

To describe a strategy in $G_p^\infty(\delta, \Phi)$, let $H^t = Y^t$ denote the set of histories of realized public signals in period t , $H_R^t = S^t$ denote the set of public reporting histories, and $H_i^t = A_i^t \times S_i^t$ denote the set of private histories for player i in period t . Player i 's (pure) strategy consists of two components, an “action strategy” $\alpha_i^t : H_i^t \times H^t \times H_R^t \rightarrow A_i$ and a “reporting strategy” $\rho_i^t : H_i^t \times H^t \times H_R^t \times A_i \rightarrow R_i$. Let $\alpha_i = (\alpha_i^0, \alpha_i^1, \dots)$, $\rho_i = (\rho_i^0, \rho_i^1, \dots)$, $\alpha = \{\alpha_i\}_{i \in N}$, $\rho = \{\rho_i\}_{i \in N}$ and let $\sigma = (\alpha, \rho)$. As in the repeated game without communication, pure strategies induce probability measures on A^∞ . Player i 's discounted expected payoff in $G_p^\infty(\delta, \Phi)$ is $w_i^{\sigma, \delta}(\Phi) = (1 - \delta) \sum_{t=0}^\infty \delta^t E[g_i^t(\tilde{a}^t)|\sigma, \Phi]$. We usually drop Φ when the public communication device being used is clear from the context.

A strategy $\sigma_i = (\alpha_i, \rho_i)$ for player i is *truthful* if player i reports her private signal truthfully whenever she played according to α_i in the same period, i.e., $\rho_i^t(s_i|h_i^t, h^t, h_R^t, \alpha_i^t(h_i^t, h^t, h_R^t)) = s_i$ for every (h_i^t, h^t, h_R^t) and s_i . Note that we allow players to lie immediately after a deviation in action. A strategy $\sigma_i = (\alpha_i, \rho_i)$ is *public* if α_i^t only depends on $h^t = (y^0, \dots, y^{t-1}) \in H^t$ and ρ_i^t depends only on $(h^t, a_i) \in H^t \times A_i$. Since we focus on this class of strategies in the public communication extension, we will write $\alpha_i^t(h^t) \in A_i$ instead of $\alpha_i^t(h_i^t, h^t, h_R^t)$ and $\rho_i^t(\cdot|h^t, \alpha_i^t(h^t)) \in R_i$ instead of $\rho_i^t(\cdot|h_i^t, h^t, h_R^t, \alpha_i^t(h_i^t, h^t, h_R^t))$. Notice that there is a natural one-to-one relationship between public strategies in $G_p^\infty(\delta)$ and the action strategy components of public strategies in $G_p^\infty(\delta, \Phi)$. Note that we need to check incentive constraints only at each h_t when every player uses a public strategy.

We extend the standard definition of perfect public equilibrium to the public communication extension as follows: a strategy profile σ for the public communication extension is a *perfect public equilibrium with communication* (which we will refer to as PPE from now on) if σ is a profile of truthful public strategies and the continuation strategy profile in the beginning of any period, which is a truthful public strategy by definition, constitutes a Nash equilibrium in the continuation game. A strategy profile σ is an η -uniformly strict perfect public equilibrium with communication if σ is a perfect public equilibrium with communication and any player would lose at least η in terms of discounted average payoff at any public history h^t when she deviates from her *equilibrium action*.

3. Informational size and incentive compatibility

3.1. Distance between stage games

We focus on public monitoring games (G, π) and private monitoring games (G', p) that are close in two respects. First, g and g' are close. Second, there exists a public coordination device ϕ that can generate a public signal distribution close to π .

Definition 2. Let (G', p) be a private monitoring game and (G, π) be a public monitoring game. Given any public coordinating device ϕ , we say that p^ϕ is an ε -approximation of π if $\max_a \|\pi(\cdot|a) - p^\phi(\cdot|a)\|_2 \leq \varepsilon$.⁶ (G', p, ϕ) is an ε -approximation of (G, π) if p^ϕ is an ε -approximation of π and $\max_{i,a} |g_i(a) - g'_i(a)| \leq \varepsilon$.

The following is a canonical example of ε -approximation.

Example 1 (Conditionally Independent Signals). Let \tilde{y} denote a random variable that can take the values 0 or 1 with equal probability. There are n players (n odd). They do not observe the realization of \tilde{y} but each observes a noisy private signal correlated with \tilde{y} . Specifically, if $\tilde{y} = y$, then player i observes a private signal $s_i \in S_i = \{0, 1\}$, which agrees with y with probability $\beta \in (\frac{1}{2}, 1]$ and differs from y with probability $1 - \beta$. These private signals are conditionally independent. Suppose that all players report their private signals (s_1, \dots, s_n) simultaneously and truthfully and $\phi(1|s_1, \dots, s_n) = 1$ if the majority of the players announce 1 while $\phi(0|s_1, \dots, s_n) = 0$ if the majority of the players announce 0.⁷ Clearly the distribution on $\{0, 1\}$ generated this way is a good approximation of the original distribution of \tilde{y} when β is sufficiently close to 1 or when the number of players is sufficiently large.

Since our notion of approximation is relatively weak, it may happen that two seemingly different monitoring structures are close to each other. For example, Y and S_i can be different sets. As a second example, consider the following monitoring structure.

Example 2 (Perfectly Complementary Information). Suppose that $Y = \{0, 1\}$ and $S_i = \{0, 1\}$ and let π be the distribution of a hidden signal \tilde{y} on Y . There are six players. The distribution p of private signals on S satisfies the following. When \tilde{y} is 1, the private signal profile is such that three players receive signal 0 and three players receive signal 1, and each such profile of signals is equally likely. When \tilde{y} is 0, the private signal profile is either $(1, 1, 1, 1, 1, 1)$ or $(0, 0, 0, 0, 0, 0)$, each with probability $\frac{1}{2}$. Consider a public coordination mechanism ϕ such that $\phi(0|s) = 0$ if at least five players announce the same signal and $\phi(0|s) = 1$ otherwise. Then p^ϕ is a 0-approximation of π .⁸

⁶ Throughout the text, $\|\cdot\|_1$ will denote the ℓ_1 -norm and $\|\cdot\|_2$ will denote the ℓ_2 -norm.

⁷ The generated public signal is the maximum likelihood estimate of the true realization of \tilde{y} .

⁸ In this example, we use six players to make sure that no player's report affects the realization of \tilde{y} . This is not the case with five players.

Suppose that, with five players, either $(1, 1, 1, 1, 1)$ or $(0, 0, 0, 0, 0)$ is observed with probability $\frac{1}{2}$ given $\tilde{y} = 0$, signal profiles with at most one 0 or 1 is never observed and every other profile is observed with equal probability ($= \frac{1}{20}$) given $\tilde{y} = 1$. If we use a similar public coordination device, then every player is pivotal given a realization of signal profile such as $(1, 1, 1, 1, 0)$.

However note that (p, ϕ) is still nonexclusive because no player's reporting strategy can change the distribution of \tilde{y} .

In this example, each player’s prior beliefs and posterior beliefs on Y are the same. Hence, a player’s signal alone provides no information about the true value of the underlying signal, yet the aggregated private signals completely reveal the true underlying signal.

In this example and [Example 1](#) with $\beta = 1$ and $n \geq 3$, one player’s report does not affect the generated public signal at all. Following Postlewaite and Schmeidler [33], we say that a pair (p, ϕ) is nonexclusive when $p^\phi(y|a, r_i) = p^\phi(y|a)$ for any $a, y, r_i \in R_i$ and i .

3.2. Informational size, distributional variability, and one-shot revelation game

We turn to the issue of truthful revelation of private information in this subsection. Although our main interest is in repeated games, it is useful to consider the following simple *one-shot information revelation game* ([26]) first. Fix any private monitoring game (G', p) . For any public coordination device ϕ , any profile of payoff functions $w : Y \rightarrow \mathbb{R}^n$, and any $a \in A$, the one-shot information revelation game (G', p, ϕ, w, a) is defined as follows. Player i observes a private signal s_i , which is distributed according to $p(\cdot|a)$. Players report s' , then a public coordinating signal y is generated with probability $\phi(y|s')$. Finally, player i receives payoff $w_i(y)$ if the realized value of the public signal is y . In the context of repeated games, this payoff would be player i ’s continuation payoff. (G', p, ϕ, w, a) defines a game of incomplete information. Truthful reporting is an equilibrium of this game if for each i ,

$$E^\phi[w_i(\cdot)|a] \geq E^\phi[w_i(\cdot)|a, r_i]$$

for any $r_i \in R_i$.

When do players have an incentive to report their private signals truthfully in this game? To fix ideas, consider the extreme case in which (p, ϕ) is nonexclusive. Then no player has an incentive to lie because what she reports is irrelevant in the presence of the reports of the other players and has no effect on the distribution of the generated public signal. Hence honest reporting can be implemented in a one-shot revelation game given any (w, a) when (p, ϕ) is nonexclusive. More generally, it should be “easier” to induce honest reporting as each player’s influence on the public coordinating signal becomes “smaller.” The following index measures the size of this influence for each player.

Definition 3 (Informational Size).

Player i ’s informational size $v_i^\phi(s_i, s'_i, a)$ given ϕ and $(s_i, s'_i, a) \in S_i \times S_i \times A$ is the smallest ϵ satisfying

$$\Pr(\|\phi(\cdot|s_i, \tilde{s}_{-i}) - \phi(\cdot|s'_i, \tilde{s}_{-i})\|_1 > \epsilon | s_i, a) \leq \epsilon.$$

This means that, conditional on (s_i, a) , player i ’s posterior probability that she can manipulate the public signal distribution by more than $v_i^\phi(s_i, s'_i, a)$ by announcing $s'_i \neq s_i$ is at most $v_i^\phi(s_i, s'_i, a)$. For example, each player’s informational size is 0 given any (s_i, s'_i) in [Example 2](#).

Note that there may be a trade-off between keeping each player’s informational size small and approximating a particular public signal distribution. Suppose that only player 1 is perfectly informed regarding the realization of \tilde{y} in [Example 1](#), i.e., $\beta_1 = 1$ and $\beta_i = \beta \in (1/2, 1)$ for all $i \neq 1$. If the goal is to replicate the true value of \tilde{y} from private signals, then one should choose a device $\hat{\phi}$ that uses player 1’s signal exclusively (i.e. $\hat{\phi}(y|s) = 1$ for $y = s_1$). However, this makes player 1 informationally large. In such a case, although informationally inferior to $\hat{\phi}$, the

majority rule device described in Example 1 may be more desirable as it makes every player informationally small but still keeps $p^\phi(\cdot|a)$ close to $\pi(\cdot|a)$ when n is large.

Of course, a player’s informational size alone is not enough to induce honest reporting. Since players may still have an incentive to misreport their signals, we need to introduce some scheme to punish dishonest reporting. So we consider the following mechanism design problem: given that $a \in A$ is played, find a public coordination device ϕ' that generates approximately the same distribution as $p^\phi(\cdot|a)$ and makes truthful reporting a Bayesian Nash equilibrium for the one-shot revelation game (G', p, ϕ', w, a) . For this purpose, we construct a certain scoring rule that relies on a player’s conditional beliefs.

Definition 4. (Distributional Variability)

$$\Lambda_i^\phi(s_i, s'_i, a) = \left\| \frac{p^\phi(\cdot|a, s_i)}{\|p^\phi(\cdot|a, s_i)\|_2} - \frac{p^\phi(\cdot|a, s'_i)}{\|p^\phi(\cdot|a, s'_i)\|_2} \right\|_2^2$$

The number $\Lambda_i^\phi(s_i, s'_i, a) \in [0, 2]$ measures the extent to which player i ’s conditional (normalized) beliefs regarding the public coordinating signal are different given s_i and s'_i (assuming honest reporting by others).⁹ This is always close to 2 in Example 1 if β is large and is always equal to 0 in Example 2. We use this variation of player i ’s beliefs to induce her to report her private signals truthfully.¹⁰

Intuitively, it will be easier to induce players to report their private signals truthfully when the first indices are smaller and the second indices are larger. Indeed, it is the ratio of these two indexes at each (s_i, s'_i, a) that is significant for truthful revelation.

Definition 5. The probability measure p^ϕ is γ -regular for ϕ if $v_i^\phi(s_i, s'_i, a) \leq \gamma \Lambda_i^\phi(s_i, s'_i, a)$ for all $s_i \in S_i, s'_i \in S_i, a \in A$ and $i \in N$.

For example, p^ϕ is 0-regular if (p, ϕ) is nonexclusive. When we say a player is *informationally small*, we mean that the ratio of her informational size to her distributional variability given every (s_i, s'_i, a) is small. We now state a result proved in [26] that provides a sufficient condition for truthful reporting to be an equilibrium in a one shot information revelation game. Suppose that $a \in A$ and $w : Y \rightarrow \mathbb{R}^n$ is a function. Let $\bar{\mu}_i(y)$ and $\underline{\mu}_i(y)$ be probability distributions on Y whose supports are contained in $\arg \max_{y \in Y} w_i(y)$ and $\arg \min_{y \in Y} w_i(y)$ respectively.

Let $\psi_i(a, s) := \sum_{y \in Y} \frac{p^\phi(y|a, s_i)}{\|p^\phi(\cdot|a, s_i)\|_2} \cdot \phi(y|s)$ and define $\phi'_{a,w}$ as follows:

$$\phi'_{a,w}(y|s) := \frac{\sum_{i=1}^n \bar{\mu}_i(y)\psi_i(a, s) + \underline{\mu}_i(y)(1 - \psi_i(a, s))}{n} \tag{1}$$

Intuitively, $\bar{\mu}_i(y)$ is a reward and $\underline{\mu}_i(y)$ is a punishment and $\psi_i(a, s)$ is the probability of receiving a reward defined so that it is maximized given truthful reporting.

⁹ It is bounded above by 2 because $\Lambda_i^\phi(s_i, s'_i, a) = 2 \left(1 - \frac{p^\phi(\cdot|a, s_i) \cdot p^\phi(\cdot|a, s'_i)}{\|p^\phi(\cdot|a, s_i)\|_2 \|p^\phi(\cdot|a, s'_i)\|_2} \right)$.

¹⁰ Our distributional variability is similar to, but different from, the condition with the same name in McLean and Postlewaite [27]. Our condition measures the distance between a player’s conditional belief regarding the aggregated public signal, whereas their condition measures the distance between a player’s conditional belief about the other players’s private signals, i.e. $\left\| \frac{p(\cdot|s_i)}{\|p(\cdot|s_i)\|_2} - \frac{p(\cdot|s'_i)}{\|p(\cdot|s'_i)\|_2} \right\|_2^2$ (there is no action in [27]).

Proposition 1. For any private monitoring game (G', p) and any $\lambda \in (0, 1)$, if there exists a public coordination device ϕ such that p^ϕ is γ -regular for some $\gamma \in (0, \frac{\lambda}{6\sqrt{|Y|n}})$, then for any $a \in A$ and any payoff function $w : Y \rightarrow \mathbb{R}^n$, truthful reporting is a Bayesian Nash equilibrium for the one-shot information revelation game $(G', p, (1 - \lambda)\phi + \lambda\phi'_{a,w}, w, a)$ where $\phi'_{a,w}$ is a public coordination device as defined in (1) above.

Proposition 1 shows that honest reporting can be induced for any one-shot revelation game by perturbing ϕ slightly. A smaller λ would require a smaller γ but note that γ does not depend on the underlying action and the payoff function. It is easy to see that γ is independent of a , since γ -regularity requires a certain property across all actions. In addition, γ is independent of w since both gains and losses from misreporting come from the same w implying that the size of w does not matter. This independence of γ from a and w will be crucial for our robustness result.

This proposition is similar to Theorem 1 in McLean and Postlewaite [27]. However, there is a difference between our result and the result in [27]. In [27], each player’s preference is given by $u_i(x, \theta)$, where θ is an unobservable payoff relevant parameter and x is an allocation, so it is important that the true θ is recovered from private signals almost surely (“negligible aggregate uncertainty”) to implement an ex post efficient allocation $x(\theta)$. In our paper, we only require that the distribution of the public signal generated by ϕ be close to the distribution of the true public signal conditional on any action profile, because no signal has any direct effect on the players’ payoffs.

When is p^ϕ likely to be γ -regular? Consider a more general version of Example 1 with conditionally independent signals.

Definition 6. A private monitoring game (G', p) is called a β -perturbation of the public monitoring game (G, π) if $S_i = Y$ for all i , $v_i(a_i, y) = u_i(a_i, y)$ for all (a, y) and i , and there exists $q_i(\cdot|y) \in \Delta(S_i)$ for all y and i such that $p(s|a) = \sum_{y \in Y} \prod_i q_i(s_i|y)\pi(y|a)$ and $q_i(y|y) \geq \beta$ for any y and i .

Note that, if (G', p) is a β -perturbation of (G, π) , then it is a β' -perturbation of (G, π) for any $\beta' \in (0, \beta)$. Suppose that (G', p) is a β -perturbation of (G, π) . Let ϕ_M be the “majority rule”, which is a public coordination device that chooses $y \in Y$ reported by the largest number of players (with some tie-breaking rule). Let $p^{\phi_M}(y|a)$ be the probability that y is generated given truthful reporting and action profile $a \in A$. Then it is easy to show that p^{ϕ_M} converges to π as $\beta \rightarrow 1$. Hence (G', p) is an ε -approximation of (G, π) for any given ε if β is close enough to one. Furthermore, player i ’s informational size $v_i^{\phi_M}(s_i, s'_i, a)$ converges to 0 given any $s_i \neq s'_i$ and a as long as $n \geq 3$ and, when π satisfies full support (i.e. $\forall y \in Y, \forall a \in A, \pi(y|a) > 0$), player i ’s distributional variability $\Lambda_i^{\phi_M}(s_i, s'_i, a)$ converges to its maximum possible value 2 given any $s_i \neq s'_i$ and a as $\beta \rightarrow 1$ for every $i \in N$. The following proposition summarizes this observation.

Proposition 2. Fix a public monitoring game (G, π) with $n \geq 3$, where π satisfies full support. Let ϕ_M be any public coordination device for β -perturbation games that selects a majority opinion. Then, for any $\varepsilon, \gamma > 0$, there exists $\beta \in (0, 1)$ such that, for every β -perturbation (G', p) , (G', p, ϕ_M) is an ε -approximation of (G, π) and p^{ϕ_M} is γ -regular.

4. Robustness of PPE

Consider any (G, π) . In this section, we ask the following question: if α^* is a PPE of $G_\pi^\infty(\delta)$, when can we find a reporting strategy ρ^* such that (α^*, ρ^*) is truthful and a PPE of $G_p^{\prime\infty}(\delta, \Phi)$ with the help of some public communication device Φ ?

The answer is simple in some extreme cases. Let ϕ be a public coordination device for which p^ϕ is an ε -approximation of π , and suppose that (p, ϕ) is nonexclusive. Let $\Phi = \{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ be the public communication device with $\phi_{h^t} = \phi$ for all t and all h^t . In this case, incentive compatibility for truthful reporting is an immediate consequence of nonexclusiveness. Essentially, we simply have another repeated game with public monitoring where the stage game payoffs and the public signal distribution are slightly perturbed. Therefore every η -uniformly strict PPE of (G, π) is an η' -uniformly strict PPE of $G_p^{\prime\infty}(\delta, \Phi)$ for some positive η' as long as ε is small enough because continuation payoffs in $G_p^{\prime\infty}(\delta, \Phi)$ converge to continuation payoffs in $G_\pi^\infty(\delta)$ uniformly across all public histories and all public strategies as $\varepsilon \rightarrow 0$ (as formally proved for [Theorem 1](#)).

Our robustness result generalizes this observation by relaxing nonexclusiveness. It says that an η -uniformly strict PPE of $G_\pi^\infty(\delta)$ is robust if, for some small enough ε and γ , there exists some public coordinating device ϕ such that (G', p, ϕ) is an ε -approximation of (G, π) and p^ϕ is γ -regular. Note that the choices of γ and ε depend on η , but do not depend on the particular selection of an η -uniformly strict PPE.

Definition 7. An η -uniformly strict PPE α^* of $G_\pi^\infty(\delta)$ is strictly robust with respect to (G', p) if there exists a public communication device Φ^* and a reporting strategy ρ^* such that $\sigma^* = (\alpha^*, \rho^*)$ is an η' -uniformly strict truthful PPE of $G_p^{\prime\infty}(\delta, \Phi^*)$ for some $\eta' > 0$.

Theorem 1. Fix $\delta \in (0, 1)$ and a public monitoring game (G, π) . For any $\eta > 0$, there exist $\gamma, \varepsilon > 0$ such that every η -uniformly strict PPE of $G_\pi^\infty(\delta)$ is strictly robust with respect to any (G', p) for which there exists ϕ such that (G', p, ϕ) is an ε -approximation of (G, π) and p^ϕ is γ -regular.

Proof. See [Appendix A](#). \square

To understand the structure of our proof, let α^* be any η -uniformly strict perfect public equilibrium in $G_\pi^\infty(\delta)$. Given $\eta > 0$, we can choose small enough $\varepsilon > 0$ and $\lambda > 0$ such that, if (G', p, ϕ) is an ε -approximation of (G, π) and players are constrained to reveal their signals truthfully, then it is strictly optimal to follow α^* at any public history given a public communication device $\Phi' = \{(1 - \lambda)\phi + \lambda\phi'_{h^t} : h^t \in Y^t, t \geq 0\}$ for any $\phi'_{h^t}, h^t \in Y^t$. This is because, as we already mentioned, the continuation payoffs in $G_p^{\prime\infty}(\delta, \Phi')$ become close to the continuation payoffs in $G_\pi^\infty(\delta)$ uniformly across all public histories given α^* (or any profile of public strategies) and truthful revelation when (G', p, ϕ) approximates (G, π) .

Since players are not constrained to report truthfully, we must construct a public communication device that induces them to do so. We construct such a public communication device as follows. To begin, fix any such public communication device Φ' . Pick any truthful public strategy $\sigma = (\alpha^*, \rho)$ in $G_p^{\prime\infty}(\delta, \Phi')$ and compute continuation payoffs after every public history. Define $\{w^{\Phi'}(h^t, \cdot) : h^t \in Y^t, t \geq 0\}$ to be the collection of continuation payoff functions. Note that they are independent of the choice of a reporting strategy as long as σ is truthful. Note also that σ need not be an equilibrium in $G_p^{\prime\infty}(\delta, \Phi')$.

By Proposition 1, we can pick small enough $\gamma > 0$ such that, if p^ϕ is γ -regular, then we can find $\{\phi_{h^t}^t : h^t \in Y^t, t \geq 0\}$ such that the incentive constraints for truth-telling are always satisfied on the equilibrium path given a public communication device $\Phi'' = \{(1 - \lambda)\phi + \lambda\phi_{h^t}^t : h^t \in Y^t, t \geq 0\}$ and the continuation payoffs specified by $w^{\Phi'}(h^t, \cdot)$. What is important here is that honest reporting can be induced given any equilibrium action and any equilibrium continuation payoff function for the same number γ as Proposition 1 states. Note that this does not mean that incentive constraints for honest reporting are satisfied off the equilibrium path, namely when a player deviates from the equilibrium action within the same period. In fact, the optimal one-period deviation may be a joint deviation in action and in report. However it can be easily shown that such joint deviation is not profitable when each player is informationally small, because of the strict incentive to play α^* at every public history.

A somewhat subtle issue arises here. Honest reporting may not be incentive compatible with respect to the public communication device Φ'' because the system of continuation payoff functions $\{w^{\Phi''}(h^t, \cdot) : h^t \in Y^t, t \geq 0\}$ induced by Φ'' is different from the system of continuation payoff function $\{w^{\Phi'}(h^t, \cdot) : h^t \in Y^t, t \geq 0\}$ induced by Φ' , with respect to which the incentive constraints for truth-telling are checked. Let $F(\Phi') = \Phi''$ be a mapping that associates Φ'' for each Φ' . If $F(\Phi^*) = \Phi^*$ for some public communication device Φ^* , then the above argument suggests that we can find a truthful perfect public equilibrium with communication $\sigma^* = (\alpha^*, \rho^*)$ in $G_p^\infty(\delta, \Phi^*)$. So the problem is reduced to finding a fixed point of the mapping F from the space of profiles of countable public coordination devices to itself. In the proof, we use a certain subset of this space that is mapped to itself by F and show that F satisfies all the properties required for applying the Fan–Glicksberg fixed point theorem.

When does our robustness result apply? The class of β -perturbations of the public monitoring game (G, π) provides a natural framework in which to pose the robustness problem. A combination of Proposition 2 and Theorem 1 immediately implies that the set of uniformly strict perfect PPE for a repeated game with public monitoring is strictly robust with respect to any β -perturbation of the stage game (G, π) if β is large enough.

Corollary 1. Fix a public monitoring game (G, π) with $n \geq 3$ and $\delta \in (0, 1)$, where π satisfies full support. For every $\eta > 0$, there exists $\beta > 0$ such that every η -uniformly strict PPE of $G_\pi^\infty(\delta)$ is strictly robust with respect to any β -perturbation of (G, π) .

We can also show that our robustness result is robust with respect to a small perturbation of p in $\Delta(S)$ in the following sense. Suppose that, for given $\eta > 0$, there exists a public coordination device ϕ such that (G, p, ϕ) is a strict ε -approximation of (G, π) (i.e. $\max_{i,a} |g_i(a) - g'_i(a)| < \varepsilon$ and $\max_a \|\pi(\cdot|a) - p^\phi(\cdot|a)\|_2 < \varepsilon$) and p^ϕ satisfies γ -regularity strictly (i.e. $\max_{i,a,s_i,s'_i} \frac{v_i^\phi(s_i,s'_i,a)}{\Lambda_i^\phi(s_i,s'_i,a)} < \gamma$) for some $\varepsilon > 0$ and $\gamma > 0$ in the statement of Theorem 1. Theorem 1 says that every η -uniformly strict PPE is strictly robust with respect to (G', p') . Now perturb p slightly in $\Delta(S)$. Note that the left hand sides of these strict inequalities are continuous in p . So all these inequalities continue to hold for any p' that is close to p in $\Delta(S)$. Therefore every η -uniformly strict PPE remains strictly robust with respect to (G', p') for any p' in some neighborhood of p in $\Delta(S)$. Note that the above corollary is a special case of this observation

where p is the (degenerate) private signal distribution with $\beta = 1$ and p' is any β -perturbation of p with large enough β .¹¹

However our result does not satisfy the following stronger notion of robustness. Let (G, π) be a game with public monitoring and define the ε -perturbation class $\Delta_\varepsilon(G, \pi)$ of (G, π) as follows:

$$\Delta_\varepsilon(G, \pi) = \{(G', p) \mid \exists \phi \text{ s.t. } (G', p, \phi) \text{ is an } \varepsilon\text{-approximation of } (G, \pi)\}.$$

As a consequence of Proposition 2, it follows that $\Delta_\varepsilon(G, \pi) \neq \emptyset$ for all ε if $n \geq 3$. We will say that a PPE α^* of $G_\pi^\infty(\delta)$ is ε -robust if for every $(G', p, \phi) \in \Delta_\varepsilon(G, \pi)$, there exists a public communication device Φ^* and a reporting strategy ρ^* such that $\sigma^* = (\alpha^*, \rho^*)$ is a truthful PPE of $G_p'^\infty(\delta, \Phi^*)$. In terms of ε -robustness, a stronger continuity result would be stated as follows: for every $\eta > 0$, there exists an $\varepsilon > 0$ such that every η -uniformly strict PPE of $G_\pi^\infty(\delta)$ is ε -robust. The statement of our Theorem is close to, but not quite the same as, this stronger notion of robustness.

5. Related literature and discussion

There is a large literature on repeated games with private monitoring and communication. Most papers in the literature focus on a folk theorem rather than robustness (Aoyagi [4], Ben-Porath and Kahneman [6], Compte [8], Fudenberg and Levine [12], Kandori and Matsushima [17], Obara [30], Tomala [37]).

Mailath and Morris [21] study robustness of PPE when a public monitoring structure is perturbed slightly, but without any communication. One of the assumptions they need for robustness is that private monitoring is *almost public*. Private monitoring is almost public if $S_i = Y$ for all i and $|\Pr(s = (y, \dots, y) \mid a) - \pi(y \mid a)|$ is small for all a and y . In a subsequent work, Mailath and Morris [22] introduce a weaker notion of approximation called ε -closeness, which does not assume $S_i = Y$. Their definition of ε -closeness (see Definition 2 in [22]) is $\max_{a,y} |\Pr(\forall i, f_i(s_i) = y \mid a) - \pi(y \mid a)| \leq \varepsilon$ for all a and y given some mapping $f_i : S_i \rightarrow Y \times \{\emptyset\}$, $i \in N$. While we allow communication for our robustness result, our notion of closeness is weaker. It is easy to show that, when p is ε -close to π in their sense, there exists a public coordination device ϕ such that p^ϕ is $\varepsilon(|Y| + 1)$ -approximation of π . In Example 2, p^ϕ is a 0-approximation of π and is 0-regular, but it is not close to π in their sense. They show that some PPE without bounded recall are not robust to small perturbations of public monitoring. Since there exist uniformly strict PPE within this class that are robust in our sense with respect to more general perturbations, our result suggests that communication is essential for the robustness of a certain class of PPE.¹²

Another related paper is Anderlini and Lagunoff [3], which considers dynastic repeated games with communication where short-lived players care about their offspring. As in our paper, players may have an incentive to conceal bad information so that future generations do not suffer from mutual punishments. Their model is based on perfect monitoring and their focus is on characterizing the equilibrium payoff set rather than establishing the robustness of equilibria.

¹¹ With $\beta = 1$, (G, p, ϕ_M) is a strict ε -approximation of (G, π) for any $\varepsilon > 0$ and p^{ϕ_M} is strictly γ -regular for any $\gamma > 0$.

¹² Mailath and Morris [21] also prove a folk theorem for general repeated games with almost-perfect and almost-public private monitoring without any communication (a gap in the proof of the folk theorem in [21] is reported and corrected in Mailath and Olszewski [23]).

The notion of informational size and distributional variability appear in McLean and Postlewaite [27], although our formulation is slightly different.

There is an extensive literature on repeated games with private monitoring without communication, starting with Sekiguchi [35]. Bhaskar and Obara [7], Ely and Välimäki [9], Piccione [31] prove a folk theorem for the repeated prisoners’ dilemma game with private almost-perfect monitoring. The most successful approach to private monitoring games, which is taken by [9] and [31], is to rely on a class of equilibria called Belief-free equilibrium.^{13,14} BFE is formalized and generalized by Ely, Hörner and Olszewski [10]. A limit characterization of BFE payoffs is first obtained by [15] for the two-player case, then generalized to the n -player case by Yamamoto [38]. In general, BFE does not deliver a folk theorem except for particular games such as the prisoners’ dilemma. Various extensions of BFE have been proposed and successfully applied to prove a folk theorem for more general games (Hörner and Olszewski [15], Kandori [16], Matsushima [24], Sugaya [36], Yamamoto [39]). Miyagawa, Miyahara, and Sekiguchi [28] consider private monitoring games where each player can observe the other players’ actions perfectly at some cost, and proves a folk theorem without any further assumption on imperfect monitoring structure.

Appendix A

A.1. Preliminary lemma

First we derive an upper bound on player i ’s ability to manipulate the distribution of a public coordinating signal in terms of player i ’s informational size. Assuming honest reporting by players $j \neq i$, let $p^\phi(y|a, s_i, s'_i) = \sum_{s_{-i}} \phi(y|s'_i, s_{-i})p(s_{-i}|a, s_i)$ be player i ’s conditional belief regarding the public signal given action profile a , true type s_i and reported type s'_i .

Lemma 1. $\|p^\phi(\cdot|a, s_i, s_i) - p^\phi(\cdot|a, s_i, s'_i)\|_1 \leq 3v_i^\phi(s_i, s'_i, a)$ for all $s'_i, s_i \in S_i, a \in A$ and $i \in N$.

Proof.

$$\begin{aligned} & \|p^\phi(\cdot|a, s_i, s_i) - p^\phi(\cdot|a, s_i, s'_i)\|_1 \\ &= \sum_y \left| \sum_{s_{-i}} \{ \phi(y|s_i, s_{-i}) - \phi(y|s'_i, s_{-i}) \} p(s_{-i}|a, s_i) \right| \\ &\leq \sum_{s_{-i}} \| \phi(\cdot|s_i, s_{-i}) - \phi(\cdot|s'_i, s_{-i}) \|_1 p(s_{-i}|a, s_i) \\ &\leq (1 - v_i^\phi(s_i, s'_i, a))v_i^\phi(s_i, s'_i, a) + v_i^\phi(s_i, s'_i, a) \cdot \max_{c,d \in \Delta(Y)} \|c - d\|_1 \\ &\leq 3v_i^\phi(s_i, s'_i, a) \quad \square \end{aligned}$$

Proof of Theorem 1. Let (G, π) and (G', p) be a public monitoring game and a private monitoring game respectively. Define κ as follows

¹³ However, the *Belief-based approach* in [7] also has been studied and refined in recent papers such as Phelan and Skrzypacz [32] and Kandori and Obara [19].

¹⁴ Kandori and Obara [18] introduces a type of belief-free equilibrium in the context of repeated games with *imperfect public monitoring*.

$$\kappa := \max_{i \in N, a \in A} |g_i(a)|.$$

The discount factor δ is fixed throughout the proof.

We proceed in three steps to prove this theorem. In Step 1, we show that continuation payoffs in a public monitoring game and continuation payoffs in a private monitoring game are uniformly close across all public histories given the same public (action) strategy profile if the stage games are close at every public history given some public communication device and the players announce their private signals truthfully. In Step 2, using Proposition 1, we construct a public communication device such that the stage games of the private monitoring game are close to the stage game of the public monitoring game (as required in Step 1) and truth-telling is incentive compatible at every public history when players are informationally small. Finally, Step 3 verifies that all the (relevant) incentive constraints are satisfied in the private monitoring game for any η -uniformly strict equilibrium of the public monitoring game when players are informationally small, using the results from Step 1 and 2.

Step 1. Continuation payoffs are close with truthful reporting when stage games are uniformly close.

Claim. Suppose that $\alpha \in \Sigma$ is a public strategy in $G_\pi^\infty(\delta)$, $\Phi = \{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ is a public communication device, $\sigma = (\alpha, \rho)$ is a truthful strategy in $G'_p{}^\infty(\delta, \Phi)$, and (G', p, ϕ_{h^t}) is an ε -approximation of (G, π) for every $h^t \in Y^t$. We claim that

$$\sup_{i, h^t} |w_i^\alpha(h^t) - w_i^\sigma(h^t)| \leq \left(1 + \frac{\delta\sqrt{|Y|\kappa}}{1 - \delta}\right)\varepsilon$$

where (abusing notation)

$$w_i^\alpha(h^t) = (1 - \delta)g_i(\alpha(h^t)) + \delta \sum_y w_i^\alpha(h^t, y)\pi(y|\alpha(h^t))$$

denotes player i 's continuation payoffs after h^t in $G_\pi^\infty(\delta)$ given α and

$$\begin{aligned} w_i^\sigma(h^t) &= (1 - \delta)g'_i(\alpha(h^t)) + \delta \sum_y \sum_s w_i^\sigma(h^t, y)\phi_{h^t}(y|s)p(s|\alpha(h^t)) \\ &= (1 - \delta)g'_i(\alpha(h^t)) + \delta \sum_y w_i^\sigma(h^t, y)p^{\phi_{h^t}}(y|\alpha(h^t)) \end{aligned}$$

denotes player i 's continuation payoffs after h^t in $G'_p{}^\infty(\delta, \Phi)$ given σ .

Proof of Step 1. Suppose that $\alpha \in \Sigma$ is a public strategy in $G_\pi^\infty(\delta)$, Φ is a public communication device and $\sigma = (\alpha, \rho)$ is a truthful strategy in $G'_p{}^\infty(\delta, \Phi)$. Choose any $\varepsilon > 0$ and suppose that (G', p, ϕ_{h^t}) is an ε -approximation of (G, π) for each $\phi_{h^t} \in \Phi$. Let $B = \sup_{i, h^t} |w_i^\alpha(h^t) - w_i^\sigma(h^t)| < \infty$. For each public history h^t , we obtain

$$\begin{aligned} &|w_i^\alpha(h^t) - w_i^\sigma(h^t)| \\ &\leq (1 - \delta)\varepsilon + \delta \left| \sum_{y \in Y} w_i^\alpha(h^t, y)\pi(y|\alpha(h^t)) - \sum_{y \in Y} w_i^\sigma(h^t, y)p^{\phi_{h^t}}(y|\alpha(h^t)) \right| \end{aligned}$$

$$\begin{aligned} &\leq (1 - \delta)\varepsilon + \delta \left| \sum_{y \in Y} w_i^\alpha(h^t, y) \{ \pi(y|\alpha(h^t)) - p^{\phi_{h^t}}(y|\alpha(h^t)) \} \right| \\ &\quad + \delta \left| \sum_{y \in Y} \{ w_i^\alpha(h^t, y) - w_i^\sigma(h^t, y) \} p^{\phi_{h^t}}(y|\alpha(h^t)) \right| \\ &\leq (1 - \delta)\varepsilon + \delta \|w_i^\alpha(h^t, \cdot)\|_2 \| \pi(\cdot|\alpha(h^t)) - p^{\phi_{h^t}}(\cdot|\alpha(h^t)) \|_2 + \delta B \\ &\leq (1 - \delta)\varepsilon + \delta \sqrt{|Y|} \kappa \varepsilon + \delta B. \end{aligned}$$

Computing the supremum of the left hand side, we obtain $B \leq (1 - \delta)\varepsilon + \delta \sqrt{|Y|} \kappa \varepsilon + \delta B$, from which it follows that

$$B \leq \left(1 + \frac{\delta \sqrt{|Y|} \kappa}{1 - \delta} \right) \varepsilon$$

and the proof of the claim is complete. \square

Step 2. Constructing the public communication device.

Claim. Choose any $\lambda \in (0, 1)$. If $\gamma \in (0, \frac{\lambda}{6\sqrt{|Y|}n})$ and if there exists a public coordinating device ϕ for which p^ϕ is γ -regular, then for any pure strategy PPE α^* of $G_\pi^\infty(\delta)$, there exists a collection of public coordinating devices $\{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ such that, for each public history $h^t \in Y^t$, truthful reporting is a Bayesian Nash equilibrium in the one-shot information revelation game

$$(G^t, p, (1 - \lambda)\phi + \lambda\phi_{h^t}, w(h^t, \cdot), \alpha^*(h^t)),$$

where $\{w(h^t) : h^t \in Y^t, t \geq 0\}$ is the collection of continuation payoffs in $G_p^{\infty}(\delta, \Phi^\lambda)$ that are generated given the public communication device $\Phi^\lambda = \{(1 - \lambda)\phi + \lambda\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ and any truthful strategy profile $\sigma^* = ((\alpha_1^*, \rho_1^*), \dots, (\alpha_n^*, \rho_n^*))$ for $G_p^{\infty}(\delta, \Phi^\lambda)$.

Proof of Step 2. We use the mechanism of Proposition 1 to induce truthful reporting. Define $\psi_i : A \times S \rightarrow [0, 1]$ as follows.

$$\psi_i(a, s) := \sum_{y \in Y} \frac{p^\phi(y|a, s_i)}{\|p^\phi(\cdot|a, s_i)\|_2} \cdot \phi(y|s)$$

For any pair $\bar{\mu}_i, \underline{\mu}_i \in \Delta(Y)$, $a \in A$ and $s \in S$, define $\phi_{a, \bar{\mu}_i, \underline{\mu}_i}(\cdot|s) \in \Delta(Y)$ by

$$\phi_{a, \bar{\mu}_i, \underline{\mu}_i}(y|s) := \bar{\mu}_i(y)\psi_i(a, s) + \underline{\mu}_i(y)(1 - \psi_i(a, s))$$

and $\phi_{a, \bar{\mu}, \underline{\mu}} := \frac{\sum_{i=1}^n \phi_{a, \bar{\mu}_i, \underline{\mu}_i}}{n}$.

Let $K = (\Delta(Y) \times \Delta(Y))^n$ and let M denote the product of countably many copies of K indexed by the elements of $H := \bigcup_{t \geq 0} Y^t$. Suppose that α^* is a PPE in $G_\pi^\infty(\delta)$. For each

$$\mu = (\mu(h^t))_{h^t \in H} = ((\bar{\mu}_1(h^t), \underline{\mu}_1(h^t)), \dots, (\bar{\mu}_n(h^t), \underline{\mu}_n(h^t)))_{h^t \in H} \in M$$

where

$$\mu(h^t) = (\bar{\mu}_1(h^t), \underline{\mu}_1(h^t)), \dots, (\bar{\mu}_n(h^t), \underline{\mu}_n(h^t)) \in R^{2n|Y|},$$

define a public communication device given $\lambda \in (0, 1)$ as follows:

$$\Phi(\mu, \lambda) = \{(1 - \lambda)\phi + \lambda\phi_{\alpha^*(h^t), \bar{\mu}(h^t), \underline{\mu}(h^t)} : h^t \in Y^t, t \geq 0\}.$$

For each i , choose any reporting strategy ρ_i^* so that $\sigma_i^* = (\alpha_i^*, \rho_i^*)$ is truthful. Then the strategy profile σ^* and the public communication device $\Phi(\mu, \lambda)$ define continuation payoffs in the game $G_p^{\infty}(\delta, \Phi(\mu, \lambda))$ at every public history. For each $\mu \in M$, let $w_i^{\sigma^*}(h^t; \mu, \lambda)$ denote player i 's continuation payoff in $G_p^{\infty}(\delta, \Phi(\mu, \lambda))$ at public history h^t . Next, define for each $\mu \in M$ a subset $\Gamma(\mu) \subset M$ as follows: $\mu' \in \Gamma(\mu)$ if and only if, for all $h^t \in Y^t$ and $i \in N$,

$$\bar{\mu}'_i(h^t) \in \arg \max_{q \in \Delta(Y)} \sum_{y \in Y} q(y) w_i^{\sigma^*}((h^t, y); \mu, \lambda)$$

and

$$\underline{\mu}'_i(h^t) \in \arg \min_{q \in \Delta(Y)} \sum_{y \in Y} q(y) w_i^{\sigma^*}((h^t, y); \mu, \lambda).$$

Note that μ' need not be the same μ that we started with. Now we show that the correspondence $\Gamma : M \rightarrow M$ has a fixed point by applying the Fan–Glicksberg fixed point theorem. First, let X denote the Cartesian product of countably many copies of $\mathbb{R}^{2n|Y|}$ indexed by $H := \bigcup_{t \geq 0} Y^t$. Since $\mathbb{R}^{2n|Y|}$ is a Hausdorff locally convex topological vector space, it follows that X is a Hausdorff locally convex topological vector space with respect to the product topology. (Theorem 5.1 and Lemma 5.54 in Aliprantis and Border [2].) Since $K \subseteq \mathbb{R}^{2n|Y|}$ is nonempty, convex and compact, we conclude that M is a nonempty, convex, compact subset in X .

Since Γ is nonempty, convex valued, and compact valued, we need only verify that Γ is upper hemicontinuous. Upper hemicontinuity will follow from Berge's Theorem if we can establish that $\mu \mapsto w_i^{\sigma^*}((h^t, y); \mu, \lambda)$ is a continuous real valued function on M for each t and $h^t \in Y^t$ since the product of compact valued upper hemicontinuous correspondences from M , to K is upper hemicontinuous with respect to the product topology on X (Theorem 16.28 in [2]). To see that $\mu \mapsto w_i^{\sigma^*}((h^t, y); \mu, \lambda)$ is continuous, first define

$$\phi_{h^t}^{\mu} = (1 - \lambda)\phi + \lambda\phi_{\alpha^*(h^t), \bar{\mu}(h^t), \underline{\mu}(h^t)}$$

and let d_1 denote the ℓ_1 metric on $\mathbb{R}^{2n|Y|}$. Note that

$$\begin{aligned} & |\phi_{h^t}^{\mu'}(y|s) - \phi_{h^t}^{\mu}(y|s)| \\ & \leq \frac{\lambda}{n} \sum_{i=1}^n (|\bar{\mu}_i(h^t)(y) - \bar{\mu}'_i(h^t)(y)| \psi_i(\alpha^*(h^t), s) \\ & \quad + |\underline{\mu}_i(h^t)(y) - \underline{\mu}'_i(h^t)(y)| [1 - \psi_i(\alpha^*(h^t), s)]) \\ & \leq \frac{\lambda}{n} \sum_{i=1}^n (|\bar{\mu}_i(h^t)(y) - \bar{\mu}'_i(h^t)(y)| + |\underline{\mu}_i(h^t)(y) - \underline{\mu}'_i(h^t)(y)|) \end{aligned}$$

so that

$$\begin{aligned} & |p^{\phi_{h^t}^{\mu'}}(y|\alpha^*(h^t)) - p^{\phi_{h^t}^{\mu}}(y|\alpha^*(h^t))| \\ & \leq \sum_{s \in S} |\phi_{h^t}^{\mu'}(y|s) - \phi_{h^t}^{\mu}(y|s)| p(s|\alpha^*(h^t)) \\ & \leq \max_{s \in S} |\phi_{h^t}^{\mu'}(y|s) - \phi_{h^t}^{\mu}(y|s)| \end{aligned}$$

$$\leq \frac{\lambda}{n} \sum_{i=1}^n (|\bar{\mu}_i(h^t)(y) - \underline{\mu}_i(h^t)(y)| + |\underline{\mu}_i(h^t)(y) - \underline{\mu}'_i(h^t)(y)|)$$

Therefore,

$$\|p^{\phi_{h^t} \mu}(\cdot | \alpha^*(h^t)) - p^{\phi_{h^t} \mu'}(\cdot | \alpha^*(h^t))\|_1 \leq \frac{\lambda}{n} d_1(\mu(h^t), \mu'(h^t)).$$

Next, note that the product topology on X is metrizable (recall that X is a countable product) with well known metric \hat{d} as defined, for example, in Theorem 3.24 of [2]. For this metric \hat{d} , the following condition holds: for each $\xi > 0$ and each t , there exists a $\zeta > 0$ (depending on ξ and t) such that for each $t' \leq t$ and each $h^{t'} \in Y^{t'}$,

$$\hat{d}(\mu, \mu') < \zeta \Rightarrow d_1(\mu(h^{t'}), \mu'(h^{t'})) < \xi.$$

That is, we can make $\mu(h^{t'})$ and $\mu'(h^{t'})$ as close with respect to d_1 as we wish for any $h^{t'}$ with $0 \leq t' \leq t$ by making μ and μ' close enough with respect to \hat{d} . Consequently, we can make $p^{\phi_{h^{t'}} \mu}(\cdot | \alpha^*(h^{t'}))$ and $p^{\phi_{h^{t'}} \mu'}(\cdot | \alpha^*(h^{t'}))$ as close in the ℓ_1 norm as we wish for any $h^{t'}$ with $0 \leq t' \leq t$ by making μ and μ' close enough with respect to \hat{d} . Since the stage game payoff is bounded and we use a discounted average payoff criterion, we conclude that $\mu \mapsto w_i^{\sigma^*}((h^t, y); \mu, \lambda)$ is a continuous real valued function on M for each $h^t \in H$. Therefore, we conclude that $\Gamma : M \rightarrow M$ is nonempty valued, convex valued, compact valued and upper hemicontinuous. Applying the Fan–Glicksberg fixed point theorem, we can find $\mu^* \in M$ such that $\mu^* \in \Gamma(\mu^*)$. In particular,

$$\bar{\mu}_i^*(h^t) \in \arg \max_{q \in \Delta(Y)} \sum_{y \in Y} q(y) w_i^{\sigma^*}((h^t, y); \mu^*, \lambda)$$

and

$$\underline{\mu}_i^*(h^t) \in \arg \min_{q \in \Delta(Y)} \sum_{y \in Y} q(y) w_i^{\sigma^*}((h^t, y); \mu^*, \lambda)$$

by construction for every $h_t \in H = Y^t$ and $i \in N$.

By Proposition 1, if $\gamma \in (0, \frac{\lambda}{6\sqrt{|Y|n}})$ and p^ϕ is γ -regular, then truthful reporting is a Bayesian Nash equilibrium in the one-shot information revelation game

$$(G^t, p, (1 - \lambda)\phi + \lambda\phi_{\alpha^*(h^t), \bar{\mu}^*(h^t), \underline{\mu}^*(h^t)}, w^{\sigma^*}((h^t, \cdot); \mu^*, \lambda), \alpha^*(h^t))$$

for every $h_t \in H = Y^t$. Defining $\phi_{h^t} := \phi_{\alpha^*(h^t), \bar{\mu}^*(h^t), \underline{\mu}^*(h^t)}$ and

$$\Phi^\lambda := \{(1 - \lambda)\phi + \lambda\phi_{h^t} : h^t \in Y^t, t \geq 0\} (= \Phi(\mu^*, \lambda))$$

completes the proof of the claim. \square

Step 3. Checking all one-period deviations.

Claim. For each $\eta > 0$, there exist an $\varepsilon > 0$ and $\gamma > 0$ with the property stated in the theorem: if (G^t, p, ϕ) is an ε -approximation of (G, π) and p^ϕ is γ -regular, then for any η -uniformly strict PPE α^* of $G_\pi^\infty(\delta)$, there exists a public communication device Φ and a reporting strategy ρ^* such that $\sigma^* = (\alpha^*, \rho^*)$ is a truthful PPE of $G_p'^\infty(\delta, \Phi)$.

Proof of Step 3. Fix $\eta > 0$. Choose $\lambda \in (0, 1)$, $\varepsilon > 0$, and $\gamma \in (0, \frac{\lambda}{6\sqrt{|Y|n}})$ such that the following strict inequality is satisfied

$$2\varepsilon\sqrt{Y}\kappa + 2\left(1 + \frac{\delta\sqrt{|Y|}\kappa}{1-\delta}\right)\varepsilon + 4\lambda(\kappa + \varepsilon) + 6\kappa\gamma < \frac{\eta}{3}. \tag{2}$$

Suppose that (G', p, ϕ) is an ε -approximation of (G, π) and p^ϕ is γ -regular. Let α^* be any η -uniformly strict PPE of $G_\pi^\infty(\delta)$. By Step 2, there exists a collection $\{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ and a public communication device $\Phi^\lambda = \{\phi_{h^t}^\lambda : h^t \in Y^t, t \geq 0\}$, where $\phi_{h^t}^\lambda = (1 - \lambda)\phi + \lambda\phi_{h^t}$, such that truthful reporting is optimal in $G_p^\infty(\delta, \Phi^\lambda)$ at each public history h^t given α^* . Next, define a reporting strategy ρ^* as follows¹⁵: for $h^t \in H, a_i \in A_i, s_i \in S_i$,

$$\rho^*(s_i|h^t, a_i) = \begin{cases} s_i & \text{if } a_i = \alpha_i^*(h^t) \\ \text{any optimal report} & \text{if } a_i \neq \alpha_i^*(h^t). \end{cases}$$

We will show that $\sigma^* = (\alpha^*, \rho^*)$ is a truthful PPE of the public communication extension $G_p^{\infty}(\delta, \Phi^\lambda)$. It is clearly truthful by definition. To verify sequential rationality constraints, we apply the principle of optimality and check one-period deviations at every public history $h^t \in H$ at the beginning of each period. We must check two types of deviations: those which involve a deviation at the action stage and those that do not. By construction of Φ^λ , honest reporting is optimal when the equilibrium action is played within the same period, i.e., when $\alpha^*(h^t)$ is played given public history h^t . Consequently, the second type of deviation is not profitable. To complete the argument, we must show that, for any $h^t \in H = Y^t$ and any $i \in N$, player i cannot profitably deviate by first choosing an action a_i different from $\alpha_i^*(h^t)$ and then choosing a report contingent on a_i .

Let

$$E^\pi[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t)] = \sum_y w_i^{\alpha^*}(h^t, y)\pi(y|a_i, \alpha_{-i}^*(h^t)).$$

Since α^* is η -uniformly strict, at every h^t and for every $i \in N$, the following inequality must be satisfied for every $a_i \neq \alpha_i^*(h^t)$

$$(1 - \delta)(g_i(a_i, \alpha_{-i}^*(h^t)) - g_i(\alpha^*(h^t))) + \eta \leq \delta\{E^\pi[w_i^{\alpha^*}(h^t, \cdot)|\alpha^*(h^t)] - E^\pi[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t)]\}. \tag{3}$$

We compare the left hand side and the right hand of this inequality with the corresponding terms in $G_p^{\infty}(\delta, \Phi^\lambda)$.

To begin, note that (2) implies that $2\varepsilon < \frac{\eta}{3}$. Since (G', p, ϕ) is an ε -approximation of (G, π) , it follows that

$$g'_i(a_i, \alpha_{-i}^*(h^t)) - g'_i(\alpha^*(h^t)) \leq g_i(a_i, \alpha_{-i}^*(h^t)) - g_i(\alpha^*(h^t)) + 2\varepsilon < g_i(a_i, \alpha_{-i}^*(h^t)) - g_i(\alpha^*(h^t)) + \frac{\eta}{3}. \tag{4}$$

Next, we show that

¹⁵ We do not derive the optimal reporting strategy off the equilibrium path explicitly, as it is not needed for our proof.

$$\begin{aligned}
 & E^{\phi_{h^t}^\lambda} [w_i^{\sigma^*}(h^t, \cdot) | \alpha^*(h^t)] - E^{\phi_{h^t}^\lambda} [w_i^{\sigma^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] \\
 & > E^\pi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\pi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] - \frac{\eta}{3}.
 \end{aligned} \tag{5}$$

To see this, note that player i 's expected loss at h^t in $G_p'^\infty(\delta, \Phi^\lambda)$ satisfies, for any deviation $a_i \neq \alpha_i^*(h^t)$ and $r_i \in R_i$,

$$\begin{aligned}
 & E^{\phi_{h^t}^\lambda} [w_i^{\sigma^*}(h^t, \cdot) | \alpha^*(h^t)] - E^{\phi_{h^t}^\lambda} [w_i^{\sigma^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] \\
 & \geq (1 - \lambda) E^\phi [w_i^{\sigma^*}(h^t, \cdot) | \alpha^*(h^t)] - (1 - \lambda) E^\phi [w_i^{\sigma^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] - 2\lambda(\kappa + \varepsilon) \\
 & \geq E^\phi [w_i^{\sigma^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\phi [w_i^{\sigma^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] - 4\lambda(\kappa + \varepsilon) \\
 & \geq E^\phi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] \\
 & \quad - 2 \left(1 + \frac{\delta \sqrt{|Y| \kappa}}{1 - \delta} \right) \varepsilon - 4\lambda(\kappa + \varepsilon) \quad (\text{by Step 1}) \\
 & = \{ E^\phi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] \} \\
 & \quad + \{ E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] - E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] \} \\
 & \quad - 2 \left(1 + \frac{\delta \sqrt{|Y| \kappa}}{1 - \delta} \right) \varepsilon - 4\lambda(\kappa + \varepsilon).
 \end{aligned}$$

The first term on the right hand side of this expression can be bounded from below as follows:

$$\begin{aligned}
 & E^\phi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] \\
 & = \sum_{y \in Y} w_i^{\alpha^*}(h^t, y) (p^\phi(y | \alpha^*(h^t)) - \pi(y | \alpha^*(h^t))) \\
 & \quad + \sum_{y \in Y} w_i^{\alpha^*}(h^t, y) (\pi(y | \alpha^*(h^t)) - \pi(y | a_i, \alpha_{-i}^*(h^t))) \\
 & \quad + \sum_{y \in Y} w_i^{\alpha^*}(h^t, y) (\pi(y | a_i, \alpha_{-i}^*(h^t)) - p^\phi(y | a_i, \alpha_{-i}^*(h^t))) \\
 & \geq E^\pi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\pi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] \\
 & \quad - \|w_i^{\alpha^*}(h^t, \cdot)\|_2 \|p^\phi(\cdot | \alpha^*(h^t)) - \pi(\cdot | \alpha^*(h^t))\|_2 \\
 & \quad - \|w_i^{\alpha^*}(h^t, \cdot)\|_2 \|\pi(\cdot | a_i, \alpha_{-i}^*(h^t)) - p^\phi(\cdot | a_i, \alpha_{-i}^*(h^t))\|_2 \\
 & \geq E^\pi [w_i^{\alpha^*}(h^t, \cdot) | \alpha^*(h^t)] - E^\pi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] - 2\epsilon \sqrt{Y} \kappa.
 \end{aligned}$$

For the second term on the right hand side, we have

$$\begin{aligned}
 & E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t)] - E^\phi [w_i^{\alpha^*}(h^t, \cdot) | a_i, \alpha_{-i}^*(h^t), r_i] \\
 & = \sum_{y \in Y} w_i^{\alpha^*}(h^t, y) (p^\phi(y | a_i, \alpha_{-i}^*(h^t)) - p^\phi(y | a_i, \alpha_{-i}^*(h^t), r_i)) \\
 & \geq -\kappa \sum_{y \in Y} |p^\phi(y | a_i, \alpha_{-i}^*(h^t)) - p^\phi(y | a_i, \alpha_{-i}^*(h^t), r_i)| \\
 & = -\kappa \sum_{y \in Y} \left| \sum_{s_i} [p^\phi(y | a_i, \alpha_{-i}^*(h^t), s_i) - p^\phi(y | a_i, \alpha_{-i}^*(h^t), s_i, r_i(s_i))] p(s_i | a_i, \alpha_{-i}^*(h^t)) \right|
 \end{aligned}$$

$$\begin{aligned}
 &\geq -\kappa \sum_{y \in Y} \sum_{s_i} |p^\phi(y|a_i, \alpha_{-i}^*(h^t), s_i) - p^\phi(y|a_i, \alpha_{-i}^*(h^t), s_i, r_i(s_i))| p(s_i|a_i, \alpha_{-i}^*(h^t)) \\
 &= -\kappa \sum_{s_i} \|p^\phi(\cdot|a_i, \alpha_{-i}^*(h^t), s_i) - p^\phi(\cdot|a_i, \alpha_{-i}^*(h^t), s_i, r_i(s_i))\|_1 p(s_i|a_i, \alpha_{-i}^*(h^t)) \\
 &\geq -3\kappa \sum_{s_i} v_i^\phi(s_i, r_i(s_i), a) p(s_i|a_i, \alpha_{-i}^*(h^t)) \quad (\text{by Lemma 1}) \\
 &\geq -3\kappa\gamma \sum_{s_i} \Lambda_i^\phi(s_i, r_i(s_i), a) p(s_i|a_i, \alpha_{-i}^*(h^t)) \quad (\text{by } \gamma\text{-regularity})
 \end{aligned}$$

Since $\Lambda_i^\phi(s_i, r_i(s_i), a)$ is at most 2 by definition (see footnote 9), we have

$$E^\phi[w_i^{\alpha^*}(h^t, \cdot)|\alpha^*(h^t)] - E^\phi[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t)] \geq -6\kappa\gamma.$$

Therefore player i 's expected loss at h^t in $G_p'^\infty(\delta, \Phi^\lambda)$ for any deviation $a_i \neq \alpha_i^*(h^t)$ and r_i is bounded from below by

$$\begin{aligned}
 &E^\pi[w_i^{\alpha^*}(h^t, \cdot)|\alpha^*(h^t)] - E^\pi[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t)] \\
 &\quad - 2\epsilon\sqrt{Y}\kappa - 2\left(1 + \frac{\delta\sqrt{|Y|\kappa}}{1-\delta}\right)\epsilon - 4\lambda(\kappa + \epsilon) - 6\kappa\gamma \\
 &> E^\pi[w_i^{\alpha^*}(h^t, \cdot)|\alpha^*(h^t)] - E^\pi[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t)] - \frac{\eta}{3} \quad (\text{by (2)}).
 \end{aligned}$$

Combining (3), (4) and (5), we conclude that

$$\begin{aligned}
 &(1 - \delta)(g'_i(a_i, \alpha_{-i}^*(h^t)) - g'_i(\alpha^*(h^t))) + \frac{2\eta}{3} \\
 &< \delta(E^{\phi_{h^t}^\lambda}[w_i^{\alpha^*}(h^t, \cdot)|\alpha^*(h^t)] - E^{\phi_{h^t}^\lambda}[w_i^{\alpha^*}(h^t, \cdot)|a_i, \alpha_{-i}^*(h^t), r_i])
 \end{aligned}$$

for every $r_i \in R_i$, $a_i \neq \alpha_i^*(h^t)$, h^t and $i \in N$. Let $\eta' = \frac{2\eta}{3}$. Then we can conclude that the suggested strategy $\sigma^* = (\alpha^*, \rho^*)$ is a η' -uniformly strict truthful PPE of $G_p'^\infty(\delta, \Phi^\lambda)$. This completes the proof of the theorem. \square

References

[1] D. Abreu, D. Pearce, E. Stacchetti, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* 59 (1990) 1041–1063.
 [2] C.D. Aliprantis, K. Border, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer-Verlag, 1999.
 [3] L. Anderlini, R. Lagunoff, Communication in dynastic repeated games: 'whitewashes' and 'coverups', *Econ. Theory* 26 (2005) 265–299.
 [4] M. Aoyagi, Communication equilibria in repeated games with imperfect private monitoring, *Econ. Theory* 25 (2005) 455–475.
 [5] R. Aumann, L. Shapley, Long-term competition: a game theoretic analysis, 1976, mimeo.
 [6] E. Ben-Porath, M. Kahneman, Communication in repeated games with private monitoring, *J. Econ. Theory* 70 (1996) 281–297.
 [7] V. Bhaskar, I. Obara, Belief-based equilibria in the repeated prisoners' dilemma with private monitoring, *J. Econ. Theory* 102 (2002) 40–69.
 [8] O. Compte, Communication in repeated games with imperfect private monitoring, *Econometrica* 66 (1998) 597–626.
 [9] J. Ely, J. Välimäki, A robust folk theorem for the prisoner's dilemma, *J. Econ. Theory* 102 (2002) 84–105.
 [10] J. Ely, J. Hörner, W. Olszewski, Belief-free equilibria in repeated games, *Econometrica* 73 (2005) 377–415.

- [11] F. Forges, An approach to communication equilibria, *Econometrica* 54 (1986) 1375–1385.
- [12] D. Fudenberg, D. Levine, The Nash-threats folk theorem with communication and approximate common knowledge in two player games, *J. Econ. Theory* 132 (2007) 461–473.
- [13] D. Fudenberg, E. Maskin, The folk theorem in repeated games with discounting or with incomplete information, *Econometrica* 54 (1986) 533–554.
- [14] D. Fudenberg, D. Levine, E. Maskin, The folk theorem with imperfect public information, *Econometrica* 62 (1994) 997–1039.
- [15] J. Hörner, W. Olszewski, The folk theorem for games with private, almost-perfect monitoring, *Econometrica* 74 (2006) 1499–1544.
- [16] M. Kandori, Weakly belief-free equilibria in repeated games with private monitoring, *Econometrica* 79 (2010) 877–892.
- [17] M. Kandori, H. Matsushima, Private observation and collusion, *Econometrica* 66 (1998) 627–652.
- [18] M. Kandori, I. Obara, Efficiency in repeated games revisited: the role of private strategies, *Econometrica* 74 (2006) 499–519.
- [19] M. Kandori, I. Obara, Towards a belief-based theory of repeated games with private monitoring: an application of POMDP, 2010, mimeo.
- [20] G. Mailath, S. Matthews, T. Sekiguchi, Private strategies in finitely repeated games with imperfect public monitoring, *Contrib. Theoretical Econ.* 2 (1) (2002).
- [21] G. Mailath, S. Morris, Repeated games with almost-public monitoring, *J. Econ. Theory* 102 (2002) 189–228.
- [22] G. Mailath, S. Morris, Coordination failure in repeated games with almost-public monitoring, *Theoretical Econ.* 1 (2006) 311–340.
- [23] G. Mailath, W. Olszewski, Folk theorems with bounded recall under (almost) perfect monitoring, *Games Econ. Behav.* 71 (2011) 174–192.
- [24] H. Matsushima, Repeated games with private monitoring: two players, *Econometrica* 72 (2004) 823–852.
- [25] R. McLean, I. Obara, A. Postlewaite, Informational smallness and private monitoring in repeated games, 2003, mimeo.
- [26] R. McLean, I. Obara, A. Postlewaite, Informational size and folk theorem with communication for repeated games with private monitoring, 2013, mimeo.
- [27] R. McLean, A. Postlewaite, Informational size and incentive compatibility, *Econometrica* 70 (2002) 2421–2454.
- [28] E. Miyagawa, Y. Miyahara, T. Sekiguchi, The folk theorem for repeated games with observation costs, *J. Econ. Theory* 139 (2008) 192–211.
- [29] R. Myerson, Multistage games with communication, *Econometrica* 54 (1986) 323–358.
- [30] I. Obara, Folk theorem with communication, *J. Econ. Theory* 144 (2009) 120–134.
- [31] M. Piccione, The repeated prisoner’s dilemma with imperfect private monitoring, *J. Econ. Theory* 102 (2002) 70–83.
- [32] C. Phelan, A. Skrzypacz, Beliefs and private monitoring, *Rev. Econ. Stud.* 79 (2012) 1637–1660.
- [33] A. Postlewaite, D. Schmeidler, Implementation in differential information economies, *J. Econ. Theory* 39 (1986) 14–33.
- [34] A. Rubinstein, Equilibrium in supergames with the overtaking criterion, *J. Econ. Theory* 21 (1979) 1–9.
- [35] T. Sekiguchi, Efficiency in repeated prisoner’s dilemma with private monitoring, *J. Econ. Theory* 76 (1997) 345–361.
- [36] T. Sugaya, Folk theorem in repeated games with private monitoring, 2011, mimeo.
- [37] T. Tomala, Perfect communication equilibria in repeated games with imperfect monitoring, *Games Econ. Behav.* 67 (2009) 682–694.
- [38] Y. Yamamoto, Efficiency results in N player games with imperfect private monitoring, *J. Econ. Theory* 135 (2007) 382–413.
- [39] Y. Yamamoto, A limit characterization of belief-free equilibrium payoffs in repeated games, *J. Econ. Theory* 144 (2009) 802–824.