



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 05-024

“Informational Smallness and Private Monitoring in Repeated Games”

by

Richard McLean, Ichiro Obara, and Andrew Postelwaite

<http://ssrn.com/abstract=796946>

Informational Smallness and Private Monitoring in Repeated Games*

Richard McLean
Rutgers University

Ichiro Obara
UCLA

Andrew Postlewaite
University of Pennsylvania

First version: May 2001

This version: July 20, 2005

Abstract

For repeated games with noisy private monitoring and communication, we examine *robustness* of perfect public equilibrium/subgame perfect equilibrium when private monitoring is “close” to some public monitoring. Private monitoring is “close” to public monitoring if the private signals can generate approximately the same public signal once they are aggregated. Two key notions on private monitoring are introduced: *Informational Smallness* and *Distributional Variability*. A player is informationally small if she believes that her signal is likely to have a small impact when private signals are aggregated to generate a public signal. Distributional variability measures the variation in a player’s conditional beliefs over the generated public signal as her private signal varies. When informational size is small relative to distributional variability (and private signals are sufficiently close to public monitoring), a uniformly strict equilibrium with public monitoring remains an equilibrium with private monitoring and communication.

To demonstrate that uniform strictness is not overly restrictive, we prove a uniform folk theorem with public monitoring which, combined with our robustness result, yields a new folk theorem for repeated games with private monitoring and communication.

Keywords: Communication, Informational size, Perfect Public Equilibrium, Private monitoring, Public monitoring, Repeated games, Robustness

JEL Classifications: C72, C73, D82

*We thank George Mailath for very helpful discussions. We also thank the participants of the 2002 North American summer meetings of the Econometric Society at UCLA and 2004 Second World Congress of Game Theory at Marseille, at which this paper was presented, for helpful comments. Postlewaite gratefully acknowledges support from National Science Foundation grant SBR-9810693.

1 Introduction

How groups effect cooperation is one of the most important social phenomena, and also one of the least understood. 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 about what is commonly known to the members of the group. It is typically assumed that all agents involved in a long term relationship share the same public information, either perfectly or imperfectly, through which each agent is monitored by the other agents. For the case in which each agent can observe all other agents' actions directly (perfect monitoring), Aumann and Shapley [5] and Rubinstein [16] proved a folk theorem without discounting, and Fudenberg and Maskin [9] proved a folk theorem with discounting. For the case in which each agent observes a noisy public signal (imperfect public monitoring), Abreu, Pearce and Stacchetti [1] characterized the set of pure strategy sequential equilibrium payoffs for a fixed discount factor, and Fudenberg, Levine, and Maskin [10] proved a folk theorem with discounting.

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 no 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 each agent as having only partial information about the environment and, most importantly, agents may not know what information other agents have. Players may communicate their partial information to other players in order to build a "consensus" about the current situation, which can be used to coordinate future behavior. In this view, repeated games with public information can be thought of as a reduced form of a more complex interaction involving the revelation of agents' private information.

This point of view leads us to examine *robustness* of equilibria with public monitoring when monitoring is private, but "close" to public monitoring when communication is allowed. Through communication, private signals are aggregated. Private monitoring is informationally close to public monitoring if the private signals, in aggregate, can generate approximately the same public signal distribution. One can think of this as a situation in which information contained in the public signal is dispersed among agents

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

in the form of private signals. We ask whether an equilibrium with public monitoring remains an equilibrium with respect to the public signal generated from private monitoring and communication, and whether (and how) players can be induced to reveal their private information.

When private monitoring approximates public monitoring and private signals are revealed truthfully, the game resembles a repeated game with public monitoring, which we understand. Hence, our main focus is on players' truth-telling constraints. The revelation of private information can be problematic, as can be seen in an equilibrium in which there is a simple trigger strategy to support collusion. In a private monitoring setting with communication, the trigger strategy is based on the announcements of all players. It is clear that players will not want to reveal any private information that may trigger mutual punishment.

In this paper, we provide a sufficient condition on the private monitoring structure that assures that truthful revelation of private signals can be induced. The condition is related to the concepts of *informational size* and *distributional variability* introduced in McLean and Postlewaite [15] (hereafter, MP).

To illustrate these ideas, consider an imperfect public monitoring problem G and a private monitoring problem \hat{G} . In G , each action profile a generates a public signal y from a set Y with probability $\pi(y|a)$. In \hat{G} , 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 \hat{G} , we augment the model with a “coordinating device” ϕ that associates a (possibly random) public signal in the set Y with each private message profile. In this expanded game, players choose an action profile a , and then observe their private signals. Upon observing their private signals, they make a public announcement regarding these private signals. The publicly announced signal profile is then used in conjunction with ϕ to generate a public signal. Note that the generated public signal is common knowledge.²

Suppose that players truthfully reveal their private signals in the expanded game with public announcements and coordinating device ϕ . Then each action profile a induces a distribution on Y where the public signal y is chosen with probability $\pi^\phi(y|a) := \sum_{s \in S} \phi(y|s)p(s|a)$. If for each action vector a the resulting distribution on public signals $\pi^\phi(\cdot|a)$ is close to the distribution on public signals $\pi(\cdot|a)$ in the public monitoring problem G , we say that the expanded private monitoring problem is “close” to the initial

²The announced signal profile \tilde{s} is also assumed to be common knowledge, but this can be dropped without consequence.

public monitoring problem G given the mapping ϕ . Thus, if the players can be induced to truthfully reveal their private signal vector s , players can coordinate on the generated public signal as in the initial public monitoring problem. Roughly speaking, player i is informationally small if her private information is unlikely to have a large effect on the distribution of the generated public signal. An implication of a player being informationally small is that she will have little incentive to misreport her private signal to manipulate the other players' behavior to her advantage. Distributional variability is an index that measures the variation in a player's conditional beliefs over the generated public signal as her private signal varies. When this index is large, it is easier to statistically detect and punish a lie.

Players are naturally informationally small in many settings. Suppose, for example, that there are many players who observe conditionally independent and identically distributed noisy private signals of a hidden signal. If ϕ maps each vector of signals into the posterior distribution on public signals, each player is then informationally small by the law of large numbers. Alternatively, with the same function ϕ , agents receiving conditionally i.i.d. signals about an underlying state of the world will be informationally small if their signals are very precise even if the number of agents is small. When a player's informational size is small relative to the distributional variability, it is possible to perturb ϕ slightly to provide players the necessary incentives to induce truthful revelation.

We will prove a robustness theorem that says, “roughly”, the following. Consider a public monitoring problem and a related private monitoring problem which is close given the mapping ϕ as above, and a perfect public equilibrium σ of the public monitoring problem. Then if all players are sufficiently informationally small relative to their distributional variability, there exists public history dependent $\phi^{h^t}(\cdot|s)$, which is a small perturbation of ϕ at each history h^t , such that σ remains an equilibrium with respect to the public signals generated by $\phi^{h^t}(\cdot|s)$ with some truth-telling strategy. In short, under the conditions described, the equilibrium strategies for the public monitoring problem can be supported as equilibria in the related private monitoring problem if they can communicate. We say “roughly” in the statement above because we prove this not for all equilibria of the public monitoring problem, but only *uniformly strict perfect public equilibria*. A uniformly strict perfect public equilibrium is an equilibrium in which each player loses more than a fixed amount by deviating at any history. A trigger strategy with strict incentive is one example of uniformly strict PPE. Such equilibrium remains an equilibrium even when public signal distribution is slightly perturbed, thus when private monitoring is enough close to public

monitoring.

The intuition for the result is similar to that in MP [15], but there is one complication which is particular to repeated games. When players are informationally small, the necessary punishment to induce truthful reporting is small. However, any punishment needs to be endogenously generated by continuation payoffs in repeated games rather than through side payments, and it is not clear in general whether there exist sufficiently large punishments to induce truth-telling after every history. We cannot use arbitrary equilibria to create any punishment we like.

Although uniformly strict perfect public equilibrium is a special class of perfect public equilibrium, we can prove a *uniform folk theorem* with it building on the folk theorem by Fudenberg, Levine and Maskin [10]. More precisely, for any smooth set W in the interior of the feasible and individually rational payoff set, we can find $\eta > 0$ such that every payoff profile in W is supported by $(1 - \delta)\eta$ uniformly strict perfect public equilibria for large enough δ .³ We then combine this uniformly strict folk theorem with our robustness result to prove a folk theorem for repeated games with private monitoring and communication. This folk theorem is not covered by the existing folk theorems with private monitoring and communication.

Related Literature

Our approach is related to Ben-Porath and Kahneman [6]. They prove a folk theorem when a player’s action is perfectly observed by at least two other players. Although their focus is on the folk theorem rather than robustness of equilibria, their analysis is similar to ours. They fix a strategy to support a payoff profile with perfect monitoring, and construct a similar strategy augmented with an announcement strategy to support the same payoff profile in the private monitoring setting. Their strategies employ draconian punishments when a player’s announcement is inconsistent with others’ announcements (“shoot the deviator”). Our paper differs from their paper in many respects. First, since Ben-Porath and Kahneman [6] focused on a folk theorem, they restrict attention to a particular class of strategy profiles, which are similar to the strategy profiles employed in Fudenberg and Maskin [9], while we focus on robustness of strategies in general. Secondly, our paper uses not only perfect monitoring but also imperfect public monitoring as a benchmark. Finally, private signals are noisy in our paper.

Compte [7] and Kandori and Matsushima [11] also consider communication in repeated games with private monitoring. These papers provide

³A $(1 - \delta)\eta$ uniformly strict perfect public equilibrium is a perfect public equilibrium in which every player loses at least $(1 - \delta)\eta$ by deviating at any history.

sufficient conditions on the private monitoring structure for a folk theorem. Compte [7] assumes that players' private signals are conditionally independent, while Kandori and Matsushima [11] assume that player i 's deviation and player j 's deviation can be statistically distinguished based on the private signals of the rest of players. These conditions are not implied by the condition we impose on the information structure.

Aoyagi [4] recently proved a Nash-threat folk theorem in a setting similar to Ben-Porath and Kahneman [6], but with noisy private monitoring. In his paper, each player is monitored by a subset of players. Private signals are noisy and reflect the action of the monitored player very accurately when they are jointly evaluated. That is, private monitoring is jointly almost perfect. In his paper, players have access to a more general coordination device than ours, namely, mediated communication. On the other hand, his result applies to the case with only two players, while many folk theorems, including ours, require more than two players.

Mailath and Morris [13] is close in spirit to our paper. They also focus on robustness of equilibrium when a public monitoring structure is perturbed, but without communication. One of their assumptions is that private monitoring is *almost public*. When private monitoring is almost public, the space of private signals for each player coincides with the space of public signals and $|\Pr(w = (y, \dots, y) | a) - \pi(y|a)|$ is small, where w is a private signal profile. In a subsequent paper [14], Mailath and Morris introduce a more general notion of almost public monitoring, and refer to their previous notion of almost publicness as *minimally almost public*. In their new definition, *almost publicness*, the space of private signals for each player can be larger than the space of public signals. As in our paper, they consider a mapping from private signal profiles to public signals. Almost publicness in this general sense implies our notion of closeness, but not informational smallness. Minimally almost public monitoring implies informational smallness as well when there are at least three players.⁴ We discuss this in more detail below. They show that perfect public equilibria without bounded recall is generally not robust to a perturbation of public monitoring without communication even when private monitoring is almost public. Since there are many uniformly strict equilibria without bounded recall, our result, together with their result, suggests that communication is essential for robustness of perfect public equilibria.

⁴When there are only two players, both of them may not be informationally small even when the monitoring is almost public. As usual, it is more difficult to induce truthful revelation of private information with two players.

Fudenberg and Levine [8] prove a folk theorem for repeated games with communication when the game has *almost perfect messaging*. This notion also generalizes almost public monitoring from Mailath and Morris [13] by allowing many irrelevant private signals. Their result also covers the case with two players.

Anderlini and Lagunoff [3] consider 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 robustness of equilibrium.

The model is described in Section 2 and the concepts of informational smallness and informational variability are introduced in Section 3. Section 4 is devoted to an example to illustrate the idea of our main theorem. Section 5 states and proves our robustness result. In Section 6, we prove uniformly strict folk theorem with imperfect public monitoring, and combine it with our robustness result to prove a folk theorem with private monitoring and communication. Section 7 concludes.

2 Preliminaries

2.1 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 $a = (a_1, \dots, a_n) \in \Pi_i A_i := A$. Actions are not observable, but players observe a stochastic public signal from a finite set Y ($|Y| = m$). The probability distribution on Y given a is denoted by $\pi(\cdot|a)$. We do not assume full support or common support, that is, $\{y \in Y | \pi(y|a) > 0\}$ can depend on $a \in A$. This allows both perfect monitoring ($Y = A$ and $\pi(y|a) = 1$ if $y = a$) and imperfect public monitoring. The stage game payoff for player i is $g_i(a)$. We assume that y is the only available signal and players do not obtain any additional information from their own payoffs.⁵ We call this stage game G . We normalize the payoffs so that each player's pure strategy minmax payoff is 0. The feasible payoff set is $V = \text{co}\{g(a) | a \in A\}$ and $V^* = \{v \in V | v \gg \mathbf{0}\}$ denotes the individually rational and feasible payoff set. Note that pure minmax is used here instead of mixed minmax, which is usually smaller.

⁵This is satisfied, for example, if we regard $g_i(a)$ as an expected payoff given by $g_i(a) = \sum_y u_i(a_i, y) \pi(y|a)$, where $u_i(a_i, y)$ is player i 's realized payoff.

Player i 's private history is $h_i^t = (a_i^0, \dots, a_i^{t-1}) \in H_i^t$ and a public history is $h^t = (y^0, \dots, y^{t-1}) \in H^t$ with $H_i^0 = H^0 = \{\emptyset\}$. Player i 's pure strategy is $\sigma_i = \{\sigma_i^t\}_{t=0}^\infty$, where σ_i^t is a mapping from $H_i^t \cup H^t$ to A_i . A strategy profile is $\sigma = \{\sigma_i\}_{i \in N} \in \Sigma$. Player i 's discounted average payoff given σ is $V_i(\sigma, \delta) = (1 - \delta) \sum_{t=0}^\infty \delta^t E[g_i(a^t) | \sigma]$. We denote this repeated game associated with the stage game G by $G^\infty(\delta)$.

A strategy is *public* if it only depends on H^t . A profile of public strategies constitute a *perfect public equilibrium* if, after every public history, the continuation (public) strategy profile constitutes a Nash equilibrium (Fudenberg, Levine, Maskin[10]). Note that a perfect public equilibrium is a subgame perfect equilibrium when the stage game is of perfect information. Since we focus on perfect public equilibrium, we ignore private histories and denote a strategy simply by $\sigma^t(h^t)$ instead of $\sigma^t(h_i^t, h^t)$.⁶

We need to introduce the notion of *uniformly strict perfect public equilibrium* for $G^\infty(\delta)$. Let $w_i^\sigma(h^{t+1})$ be player i 's continuation payoff from period $t+1$ at public history $h^{t+1} = (h^t, y^t)$ given σ .

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

$$\begin{aligned} & (1 - \delta) g_i(\sigma^t(h^t)) + \delta \sum_{y \in Y} \pi(y | \sigma^t(h^t)) w_i^\sigma((h^t, y)) - \eta \\ & \geq (1 - \delta) g_i(a'_i, \sigma_{-i}^t(h^t)) + \delta \sum_{y \in Y} \pi(y | a'_i, \sigma_{-i}^t(h^t)) w_i^\sigma((h^t, y)) \end{aligned}$$

for all $a'_i \neq \sigma_i^t(h^t)$, h^t and $i \in N$.

This means that any player loses by more than η if she deviates at any history. It is a kind of robustness requirement, but is not as restrictive as it may seem. We will show later that a folk theorem is obtained under standard assumptions even within this class of PPE.

2.2 Imperfect Private Monitoring

Fix a stage game G with public monitoring. A corresponding stage game \hat{G} with private monitoring has the same set of players and the same action sets. The stage game payoff for player i is denoted $g_i^p(a)$. Each player

⁶Thus we ignore *private strategies*. But this is without loss of generality for pure strategies. A mixed private strategy can be sometimes beneficial. See Kandori and Obara [12].

receives a private signal s_i from a finite set S_i instead of public signal. Feasible payoff V_p and feasible and individually rational payoff set V_p^* is similarly defined. A private signal profile is $s = (s_1, \dots, s_n) \in \prod_i S_i = S$. The conditional distribution on S given a is $p(s|a)$. We assume that 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 $p(s_{-i}|a, s_i)$ be the conditional probability of $s_{-i} \in S_{-i}$ given (a, s_i) .

Players can communicate with each other every period. As suggested in the introduction, we model communication as a way to aggregate players' private signals to generate a public signal. Let $\phi : S \rightarrow \Delta Y$ be a *coordination device* which converts players' messages to a public signal. With a slight abuse of notation, we use $\phi(y|s)$ to denote the probability of y given s . This is a model of unmediated communication based on players' direct message. Assuming truthful revelation of private signals of players $j \neq i$, player i 's conditional distribution on Y generated by ϕ given (a_i, s_i) and her announcement \tilde{s}_i is given by $p^\phi(y|a, s_i, \tilde{s}_i) = \sum_{s_{-i}} p(s_{-i}|a, s_i) \phi(\tilde{s}_i, s_{-i})(y)$. We often use $p^\phi(y|a, s_i)$ for $p^\phi(y|a, s_i, s_i)$ to economize on notation.

Let $(h^t, m^t) = (y^0, \dots, y^{t-1}, \tilde{s}^0, \dots, \tilde{s}^{t-1}) \in H^t = Y^t \times S^t$ be a public history and $h_i^t = (a_i^0, \dots, a_i^{t-1}, s_i^0, \dots, s_i^{t-1}) \in H_i^t = A_i^t \times S_i^t$ be player i 's private history.⁷ Player i 's (pure) strategy consists of two components: $\sigma_i^{a,t}$ (action) and $\sigma_i^{s,t}$ (message), $t = 1, 2, \dots$, where $\sigma_i^{a,t} : H_i^t \times H^t \rightarrow A_i$ and $\sigma_i^{s,t} : H_i^t \times H^t \times A_i \times S_i \rightarrow S_i$. Let $\sigma^a = \{\sigma_i^a\}_{i \in N} \in \Sigma^a$ and $\sigma^s = \{\sigma_i^s\}_{i \in N} \in \Sigma^s$.

A message strategy is *truth-telling* if each player reports her private signal truthfully *on the equilibrium path* ($\sigma_i^{s,t}(h_i^t, h^t, m^t, \sigma_i^{a,t}(h_i^t, h^t, m^t), s_i^t) = s_i^t$). A truth-telling message strategy is denoted by σ_i^{s*} . Note that there are many truth-telling strategies because we allow players to lie off the equilibrium path.

We say that a strategy is *public* if $\sigma_i^{a,t}$ only depends on $h^t = (y^0, \dots, y^{t-1})$ and $\sigma_i^{s,t}$ on (y^0, \dots, y^{t-1}) and (a_i^t, s_i^t) . The set of public strategies is denoted Σ_P^a . Note that a public strategy does not depend on the announced message $m^t = (\tilde{s}^0, \dots, \tilde{s}^{t-1})$ even though it is public information.⁸ It only depends on public signal $h^t = (y^0, \dots, y^{t-1})$ generated via ϕ . This is to facilitate comparison between an equilibrium with public monitoring and the corre-

⁷ Whether a message profile can be observed publicly or not does not matter. Our equilibrium, which depends on only y , remains an equilibrium whether it is publicly observable or not. Thus we ignore such private message profiles in the following.

⁸ We do not use public messages in our construction. Indeed we don't need to assume that messages are publicly observable as long as y is publicly observable.

sponding equilibrium with private monitoring. Note that there is a natural one to one relationship between public strategies with public monitoring and public strategies with private monitoring. Since we focus on this class of public strategies, we use $\sigma_i^{a,t}(h^t)$ and $\sigma_i^{s,t}(h^t, \sigma_i^{a,t}(h^t), s_i^t)$ rather than $\sigma_i^{a,t}(h_i^t, h^t, m^t)$ and $\sigma_i^{s,t}(h_i^t, h^t, m^t, \sigma_i^{a,t}(h_i^t, h^t, m^t), s_i^t)$ as before.

Player i 's discounted average payoff is $V_i^p(\sigma, \delta) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t E[g_i^p(a^t) | \sigma]$. We denote this repeated game with private monitoring by $\hat{G}^\infty(\delta)$. We extend the standard definition of perfect public equilibrium to the current setting in a straightforward way. We call a strategy profile $\sigma = (\sigma^a, \sigma^s)$ a perfect public equilibrium if σ is a public strategy and the continuation strategy profile at the beginning of each period constitutes a Nash equilibrium.⁹

2.3 Distance between G and \hat{G} .

We first introduce a notion of closeness between p and π .

Definition 2 p is an ε -approximation of π if there exists a coordination device ϕ such that

$$\left\| \pi(\cdot | a) - p^\phi(\cdot | a) \right\| < \varepsilon \text{ for all } a \in A.$$

When p is an ε -approximation of π , we associate a default coordination device ϕ with p , which satisfies the above definition. This does not mean that we always use such ϕ . Later we construct another ϕ' based on ϕ which is still a good approximation of π .

The following example illustrates a possible choice of ϕ . Suppose that, given an action profile a , the joint distribution of s and y is defined by $q(s|y)\pi(y|a)$, where $q(s|y)$ is the conditional distribution of private signals given y . We can interpret s as a noisy hidden signal of y . Note that y is a sufficient statistic with respect to action profile a . Here, an example of ϕ would be as follows;¹⁰

$$\phi(s) = \arg \max_{y \in Y} q(s|y)$$

that is, $\phi(s)$ is the conditional maximum likelihood estimate of y given s .

⁹This is called 1-public perfect equilibrium in Kandori and Matsushima.[11]

¹⁰In this example (and several examples that follow), ϕ is deterministic, mapping S into Y rather than ΔY .

We say G and \hat{G} are ε -close if p is an ε -approximation of π and the payoffs differ by at most ε .

Definition 3 G and \hat{G} are ε -close if p is an ε -approximation of π and $\max_{i,a} |g_i(a) - g_i^p(a)| \leq \varepsilon$.

3 Informational Size

In this section, we introduce several definitions related to the private monitoring information structure. Let $\bar{\varepsilon}^\phi(a, s_i, s'_i)$ be the minimum $\varepsilon \geq 0$ such that

$$\Pr \left(\left\| \phi(s_i, \tilde{s}_{-i})(\cdot) - \phi(s'_i, \tilde{s}_{-i})(\cdot) \right\| > \varepsilon | a, s_i \right) \leq \varepsilon$$

Definition 4 Player i 's informational size v_i^ϕ is defined as

$$v_i^\phi = \max_a \max_{s_i} \max_{s'_i \in S_i \setminus s_i} \bar{\varepsilon}^\phi(a, s'_i, s_i).$$

This means that the probability of affecting ϕ by more than v_i^ϕ by misrepresentation is less than v_i^ϕ . We say a player is *informationally small* when her informational size is small.

The following is a simple example illustrating this idea.

Example 5 Law of Large Numbers

Consider a binary public signal y on $Y = \{0, 1\}$. It is equally likely that $y = 0$ or $y = 1$. Player i observes only her private signal $s_i \in \{0, 1\}$, which agrees with y with probability $\frac{3}{4}$ and differs with probability $\frac{1}{4}$. The private signals are conditionally independent. Suppose that ϕ is given by majority rule, that is, ϕ takes the value which corresponds to the majority of messages. When the number of players receiving a private signal goes to infinity, each player's informational size goes to 0. Note that this example also serves as an example of ε -approximation of a public signal by private signals.

While each player's informational size become negligible as the number of the players increases in the above example, this is not always the case for the best estimate of y given $\{s_i\}_i$ such as with ϕ above. For example, suppose that $s_i = y$ with probability $3/4$ for $i = 1, \dots, n-1$, but $s_n = y$ with probability 1. When $\tilde{s} = (0, \dots, 0, 1)$, the best estimate of y is 1 provided that player n tells the truth. Thus ϕ should depend solely on s_n if ϕ is to be the best estimate of y . However, then player n is not informationally

small. Consequently, although we lose some information, we may choose a different ϕ for which every player is informationally small and $p^\phi(\cdot|a)$ is still a good approximation of $\pi(\cdot|a)$. In this example, one such ϕ could be constructed by generating a new signal \hat{s}_n that adds noise to player n 's message ($\hat{s}_n = \tilde{s}_n$ with probability $3/4$) and applying majority rule to $(\tilde{s}_1, \dots, \tilde{s}_{n-1}, \hat{s}_n)$.

Remark. Almost publicness in Mailath and Morris [13] is related to ε -approximation and informational size. Suppose that private monitoring is almost public, that is, $|p((y, \dots, y)|a) - \pi(y|a)| \approx 0$ for all $y \in Y$. Define $\phi(s)$ to be that value of y observed by the largest number of the players. Then clearly p^ϕ approximates π well. Moreover players are informationally small when there are at least three players.

There is a more general notion of almost publicness in [14]. When monitoring is almost public in this sense, a good approximation is implied by almost publicness with the same ϕ , but informational smallness is not. With small probability, each player may observe some private signal for which informational size is large.¹¹

On the other hand, it is possible that private monitoring is close to public monitoring in our sense, but not almost public in their sense (see the next example).

The next example shows that a large number of players is not necessary for informational smallness.

Example 6 *Complementary Information*

Public and private signals are again binary, $Y = \{0, 1\}$, and $S_i = \{0, 1\}$. There are six players. When the true public signal is 1, the private signal profile is such that three players receive the signal 0 and three players receive the signal 1. Each such profile of signals is equally likely. When the

¹¹Almost publicness in Mailath and Morris [13] is also based on a mapping from private signal profile to a public signal (they call their old notion of almost publicness “minimally almost public”). Let $f_i : S_i \rightarrow Y \times \{\emptyset\}$ be such a mapping for player i . Note that some private signals may not be mapped to any public signal. Then they say private monitoring is ε -close to public monitoring if there exists f_i such that $|\Pr(f_i(s_i) = y, i = 1, \dots, n) - \pi(y|a)| \leq \varepsilon$ for any a, y and $\Pr(\{s_{-i} : f_j(s_j) = y, j \neq i\} | a, s_i) \geq 1 - \varepsilon$ for any $a, s_i \in f_j^{-1}(y)$, and i .

The reason why informational smallness fails is that there is no restriction on player i 's conditional belief given s_i such that $f_i(s_i) = \emptyset$.

However, we believe that it is possible to relax our notion of informational smallness so that almost public monitoring implies informational smallness without affecting our main results.

true public signal is 0, the private signal profile is either $(1, 1, 1, 1, 1, 1)$ or $(0, 0, 0, 0, 0, 0)$, each with probability $\frac{1}{2}$. In this example, each player is equally likely to receive the signal 0 or 1 for any realization of public signal, hence her signal alone provides no information about the true public signal. But when players' signals are aggregated, a private signal profile completely reveals the true public signal. Each player's informational size is exactly 0 here.¹² This information structure has the property that any single agent's private information is meaningful only when combined with the other agents' information.

The following lemma follows immediately from the definition of informational size.

Lemma 7 1. $E [\|\phi(s_i, s_{-i})(\cdot) - \phi(s'_i, s_{-i})(\cdot)\| | a, s_i] \leq (1 + \sqrt{2}) v_i^\phi$ for all $a \in A, s_i, s'_i \in S_i$.

2. $|p^\phi(y|a, s_i) - p^\phi(y|a, s_i, s'_i)| \leq 2v_i^\phi$ for all $a \in A, s_i, s'_i \in S_i, y \in Y$

Proof. (1)

$$\begin{aligned} E [\|\phi(s_i, s_{-i})(\cdot) - \phi(s'_i, s_{-i})(\cdot)\| | a, s_i] &\leq (1 - v_i^\phi) v_i^\phi + v_i^\phi \sqrt{2} \\ &\leq (1 + \sqrt{2}) v_i^\phi \end{aligned}$$

(2)

$$\begin{aligned} |p^\phi(y|a, s_i) - p^\phi(y|a, s_i, s'_i)| &\leq (1 - v_i^\phi) v_i^\phi + v_i^\phi \\ &\leq 2v_i^\phi \end{aligned}$$

■

Next we introduce distributional variability.

Definition 8 (*Distributional Variability of player i*)

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

¹²The function ϕ maps vectors of private signals with all or all but one signal being the same into $y = 0$, and all the other vectors into $y = 1$.

This index measures the extent to which player i 's conditional belief is affected by her private signal. Note that $\Lambda_i^\phi = 0$ if private signals are independent given some action profile. One implication of this is

$$\begin{aligned}\Lambda_i^\phi &\leq \left\| \frac{p^\phi(\cdot|a, s_i)}{\|p^\phi(\cdot|a, s_i)\|} - \frac{p^\phi(\cdot|a, s'_i)}{\|p^\phi(\cdot|a, s'_i)\|} \right\|^2 \\ &= 2 \left(1 - \frac{p^\phi(\cdot|a, s_i) \cdot p^\phi(\cdot|a, s'_i)}{\|p^\phi(\cdot|a, s_i)\| \|p^\phi(\cdot|a, s'_i)\|} \right)\end{aligned}$$

for any a, s_i, s'_i .

Finally, let $v^\phi = \max_i v_i^\phi$ and $\Lambda^\phi = \min_i \Lambda_i^\phi$.

4 An Example

The following example is meant to illustrate our approach behind the general theorem in the next section. Consider the following game with the “law of large numbers” information structure discussed in the last section. There are $2n + 1$ players and the stage game payoff (for both public monitoring and private monitoring) is as follows:

	$a_{-i} = \mathbf{C}$	$a_{-i} \neq \mathbf{C}$
$a_i = C$	1	-1
$a_i = D$	2	0

where $a_{-i} = \mathbf{C}$ means that all players $j \neq i$ choose C . The distribution of the public signal \mathbf{y} is;

$$Prob(\mathbf{y} = 1|a) = \begin{cases} \frac{2}{3} & \text{if } a = \mathbf{C} \\ \frac{1}{3} & \text{if } a \neq \mathbf{C} \end{cases}$$

Thus 1 is more likely to occur when every player is cooperating, while 0 is more likely when any player defects. We assume that $s_i, i = 1, \dots, 2n + 1$ are conditionally independent given y and $s_i = y$ with probability $\frac{3}{4}$ as before. Note that the stage game is defined so as to be essentially independent of the number of players.

Suppose that a simple grim trigger strategy profile σ^{trig} is a strict equilibrium for a fixed δ .¹³ We show that σ^{trig} and a truth-telling strategy is an equilibrium with respect to public signals generated by some ϕ when the number of the players is large.

¹³Note that this strict trigger strategy equilibrium is a uniformly strict PPE.

Let $\alpha(s)$ be the number of 0's in message profile s and let $\Pr(\mathbf{y} = 1|a = \mathbf{C}, \alpha(s) = \alpha)$ be the probability that the true public signal y is 1 given that \mathbf{C} is played and players observe α 0's. Note that $\Pr(\mathbf{y} = 0|a = \mathbf{C}, \alpha(s) = n') > \Pr(\mathbf{y} = 1|a = \mathbf{C}, \alpha(s) = n')$ for $n' > n$ and $\Pr(\mathbf{y} = 0|a = \mathbf{C}, \alpha(s) = n') < \Pr(\mathbf{y} = 1|a = \mathbf{C}, \alpha(s) = n')$ for $n' \leq n$. Define $\phi : S \rightarrow \{0, 1\}$ as follows:

$$\phi(s) = \begin{cases} 1 & \text{if } \alpha(s) \leq n \\ 0 & \text{if } \alpha(s) > n \end{cases}$$

Of course, $p^\phi(\cdot|a)$ is a good approximation of the original public signal distribution if the number of the players is large and true private signals are revealed.¹⁴ Thus σ^{trig} is incentive compatible as long as players announce their private signals truthfully.¹⁵

Thus the question is whether we can induce players to tell the truth. Note that ϕ does not work as it is. No player has incentive to send a negative signal ($\tilde{s}_i = 0$) because the probability of the mutual punishment goes up. Therefore we need to modify ϕ to induce truthful revelation of players' private signals, while still generating a good approximation of the original public signal distribution. To be more specific, we first construct another function $\phi' : S \rightarrow \Delta Y$ which is useful for punishing misrepresentation. This ϕ' picks each player randomly and tests whether she is announcing truthfully. Then we construct a new ϕ^* by taking a linear combination of ϕ and ϕ' , with most of the weight put on ϕ . One interpretation is that the distribution of y is determined by ϕ most of the time, but occasionally by ϕ' . In this example, each player's incentive to send a false message to manipulate public signal is bounded by her informational size. On the other hand, the size of expected punishment to deter such misrepresentation is measured by $\frac{1}{n} \times \text{distributional variability}$. This $\frac{1}{n}$ comes from the fact that each player is picked by ϕ' with equal probability. When n is large, the informational size decreases exponentially with respect to n , thus the former effect is dominated by the latter effect above some critical n . In our general theorem as well as this example, the ratio of the first effect and the second effect plays a critical role.

¹⁴Note also that the players are informationally small with this ϕ .

¹⁵We do not require players to tell the truth after deviating to a non-equilibrium action within the period. In fact, it may not be optimal to announce a true private signal in such a history. However, we know that such joint deviation in action and message is not profitable anyway when players are informationally small (it is not so different from a simple deviation in action with truth-telling). Thus we ignore such deviations and do not specify sequentially optimal announcement after a deviation within the period.

In the following, we explicitly construct ϕ' and ϕ^* to induce truth-telling. When $\phi(\tilde{s}) = y \in \{0, 1\}$ in the cooperative phase, $\phi'(\tilde{s})$ is 1 with probability $\frac{1}{2n+1} \left(\sum_{j=1}^{2n+1} \frac{p^\phi(y|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|} \right)$ and $\phi'(\tilde{s})$ is 0 with probability $1 - \frac{1}{2n+1} \left(\sum_{j=1}^{2n+1} \frac{p^\phi(y|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|} \right)$.¹⁶ The j th term $\frac{p^\phi(y|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|}$ is tailored to punish player j when player j lies; the expected value of this term decreases when player j lies.

Now let $\phi^* = (1 - \lambda)\phi + \lambda\phi'$. We choose λ small enough so that $p^{\phi^*}(y|a)$ is still close to $\pi(y|a)$ and σ^{trig} is an equilibrium when private signals are revealed truthfully. Given message profile \tilde{s} , players stay in the cooperative phase with probability $(1 - \lambda)\phi(\tilde{s}) + \frac{\lambda}{2n+1} \left(\sum_{j=1}^{2n+1} \frac{p^\phi(\phi(\tilde{s})|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|} \right)$. Since the continuation payoff is 0 once mutual punishment starts, the continuation payoff after the announcement \tilde{s} in the cooperative phase is simply

$$\left\{ (1 - \lambda)\phi(\tilde{s}) + \frac{\lambda}{2n+1} \left(\sum_{j=1}^{2n+1} \frac{p^\phi(\phi(\tilde{s})|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|} \right) \right\} \Delta$$

where $\Delta > 0$ is the continuation payoff from playing σ^{trig} .¹⁷

We first focus on the term $\frac{p^\phi(\phi(\tilde{s})|\mathbf{C}, \tilde{s}_i)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_i)\|}$ for player i . Given that player i 's true signal is s_i , the expected loss (the reduction of probability of $y = 1$) from this term when lying is as follows.

$$\begin{aligned} & \left(\sum_{y=0,1} \frac{p^\phi(y|\mathbf{C}, s_i)}{\|p^\phi(\cdot|\mathbf{C}, s_i)\|} p^\phi(y|\mathbf{C}, s_i) - \sum_{y=0,1} \frac{p^\phi(y|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} p^\phi(y|\mathbf{C}, s_i, s'_i) \right) \\ &= \left(\sum_{y=0,1} \left(\frac{p^\phi(y|\mathbf{C}, s_i)}{\|p^\phi(\cdot|\mathbf{C}, s_i)\|} - \frac{p^\phi(y|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \right) p^\phi(y|\mathbf{C}, s_i) \right. \\ & \quad \left. + \sum_{y=0,1} \frac{p^\phi(y|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \{p^\phi(y|\mathbf{C}, s_i) - p^\phi(y|\mathbf{C}, s_i, s'_i)\} \right) \Delta \end{aligned}$$

Note that $\sum_{y=0,1} \frac{p^\phi(y|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \{p^\phi(y|\mathbf{C}, s_i) - p^\phi(y|\mathbf{C}, s_i, s'_i)\}$ is bounded

¹⁶This probability is well defined because $\frac{p^\phi(y'|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|} \leq 1$.

¹⁷More precisely, Δ depends on n with private monitoring, but we omit that dependence because it converges to the original trigger strategy equilibrium payoff (only approximately because of λ) as $n \rightarrow \infty$.

below by $-\frac{2v^\phi}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|}$ (Lemma 7). Thus the expected loss is at least

$$\begin{aligned}
& \sum_{y=0,1} \left(\frac{p^\phi(y|\mathbf{C}, s_i)}{\|p^\phi(\cdot|\mathbf{C}, s_i)\|} - \frac{p^\phi(y|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \right) p^\phi(y|\mathbf{C}, s_i, s_i) - \frac{2v^\phi}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \\
&= \|p^\phi(\cdot|\mathbf{C}, s_i)\| \left(1 - \frac{p^\phi(\cdot|\mathbf{C}, s_i) \cdot p^\phi(\cdot|\mathbf{C}, s'_i)}{\|p^\phi(\cdot|\mathbf{C}, s_i)\| \|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \right) - \frac{2v^\phi}{\|p^\phi(\cdot|\mathbf{C}, s'_i)\|} \\
&\geq \frac{1}{2\sqrt{2}} \Lambda^\phi - 2\sqrt{2}v^\phi
\end{aligned}$$

On the other hand, the expected future gain by lying from the other terms such as $(1-\lambda)\phi(\tilde{s})$ and $\frac{p^\phi(\phi(\tilde{s})|\mathbf{C}, \tilde{s}_j)}{\|p^\phi(\cdot|\mathbf{C}, \tilde{s}_j)\|}$, $j \neq i$ can be shown to be at most $(1-\lambda)2v^\phi$ and $(1+\sqrt{2})v^\phi$ respectively (which also follows from Lemma 7). Hence all the truth-telling constraints are satisfied if the following condition is satisfied.

$$(1-\lambda)2v^\phi + \lambda \frac{2n}{2n+1} (1+\sqrt{2})v^\phi \leq \frac{\lambda}{2n+1} \left(\frac{1}{2\sqrt{2}} \Lambda^\phi - 2\sqrt{2}v^\phi \right).$$

We show that this condition is satisfied as $n \rightarrow \infty$ by deriving approximate values of Λ^ϕ and v^ϕ as $n \rightarrow \infty$.

First, the conditional distribution of the generated public signal is approximately the same as the conditional distribution of the true public signal for large n , thus Λ^ϕ is approximately a constant and given by

$$\sqrt{\left(\frac{1}{7}\right)^2 + \left(\frac{6}{7}\right)^2} \sqrt{\left(\frac{3}{5}\right)^2 + \left(\frac{2}{5}\right)^2} - \left(\frac{1}{7}, \frac{6}{7}\right) \left(\frac{3}{5}, \frac{2}{5}\right) > 0.$$

Next we compute v^ϕ . Since player i 's announcement can change the value of ϕ only when she is pivotal ($\alpha(\tilde{s}_{-i}) = n$), her informational size is $v^\phi = \text{Prob}(\alpha(\tilde{s}_{-i}) = n|\mathbf{C}, s_i)$, which can be computed as follows,¹⁸

$$\begin{aligned}
& \text{Prob}(\alpha(\tilde{s}_{-i}) = n|\mathbf{C}, s_i) \\
&= \text{Prob}(\mathbf{y} = s_i|s_i) \binom{2n}{n} \left(\frac{3}{4}\right)^n \left(\frac{1}{4}\right)^n \\
&\quad + \text{Prob}(\mathbf{y} \neq s_i|s_i) \binom{2n}{n} \left(\frac{3}{4}\right)^n \left(\frac{1}{4}\right)^n \\
&= \frac{(2n)!}{n!n!} \left(\frac{3}{4}\right)^n \left(\frac{1}{4}\right)^n.
\end{aligned}$$

¹⁸Note that the formula is independent of the choice of s_i (and a as well) for this example.

Since $n! \sim n^{n+\frac{1}{2}} e^{-n} \sqrt{2\pi}$ for large n by Sterling's formula, we obtain

$$\begin{aligned} \Pr(\alpha(s_{-i}) = n | s_i = 1) &\sim \frac{(2n)^{2n+\frac{1}{2}} e^{-2n} \sqrt{2\pi}}{\left(n^{n+\frac{1}{2}} e^{-n} \sqrt{2\pi}\right)^2} \left(\frac{3}{4}\right)^n \left(\frac{1}{4}\right)^n \\ &= \frac{1}{\sqrt{n\pi}} \left(\frac{3}{4}\right)^n. \end{aligned}$$

Therefore informational size converges to 0 at least at an exponential rate. This implies that, while the lower bound of the expected loss by lying $\left(\frac{\lambda}{2n+1} \left(\frac{1}{2\sqrt{2}} \Lambda^\phi - 2\sqrt{2} v^\phi\right)\right)$ is converging to 0 at the rate of $\frac{1}{n}$, the bound of the expected gain from the other terms $\left((1-\lambda) 2v^\phi + \lambda \frac{2n}{2n+1} (1+\sqrt{2}) v^\phi\right)$ is converging to 0 at a much faster rate.

Therefore, if n is large enough, truth-telling constraints are satisfied. Then ϕ^* approximates the original public signal distribution well for a small λ , thus σ^{trig} remains an equilibrium when the public signal is generated by ϕ^* .

5 Robustness Result

We first introduce the formal notion of robustness of perfect public equilibria. Our requirement is that an equilibrium public strategy for $G^\infty(\delta)$ remains an equilibrium of $\hat{G}^\infty(\delta)$ when combined with some truth-telling strategy if \hat{G} is ε -close to G .

Definition 9 *A perfect public equilibrium σ of $G^\infty(\delta)$ is ε -robust if, for every ε -close \hat{G} , there exists a coordination device $\phi^{h^t} : S \rightarrow \Delta Y$ for each h^t such that $\sigma' = (\sigma, \sigma^{s*})$ is a perfect public equilibrium of $\hat{G}^\infty(\delta)$ with some truth-telling strategy σ^{s*} .*

We prove that every η -uniformly strict equilibrium is ε -robust if informational size is relatively small compared to distributional variability. The proof is a generalization the proof in the previous example. The main difficulty arises when the continuation payoffs of a target PPE does not have enough variability to be used as a device for punishment.

When informational size is small, no player has an incentive to deviate from the equilibrium action when other players tell the truth. This is because the target equilibrium is a uniformly strict equilibrium, and stage game payoffs and continuation payoffs with private monitoring become uniformly

close to the ones with public monitoring when \hat{G} is ε -close to G . Joint deviation in action and message within a period is not profitable either (footnote 13). Thus we only need to check whether players have incentive to reveal their private signals truthfully on the equilibrium path.

At each history, we use ϕ with high probability to ensure that the generated public signal distribution is close to the original public signal distribution, but we use punishment schemes tailored for each player with small probability to provide players with an incentive to reveal their private signals truthfully on the equilibrium path. When punishment schemes are used, each player is picked with equal probability as in the previous example. We would like to define a punishment scheme ϕ'_i for player i which has the following properties: (1) it generates public signals which correspond to either the largest expected continuation payoff for player i or the smallest expected continuation payoff for player i and (2) the probability of the former is maximized by revealing her private signal truthfully. This generalizes our previous construction. In the previous example, the largest continuation payoff corresponds to the equilibrium continuation payoff in the cooperative phase and the smallest continuation payoff corresponds to the stage-game Nash equilibrium payoff.

However, there is one difficulty with such a generalization. In the example, the stage game Nash equilibrium payoff is always 0 and the equilibrium continuation payoff in the cooperative phase is a distinct positive number if the probability λ to use punishment schemes is set small. Thus which public signal leads to a better continuation payoff is clear and we were able to create an endogenous punishment by adjusting the probability of public signal leading to these continuation payoffs. This is not so straightforward in general case. Note that each player's expected continuation payoff at period t depends on how punishment schemes are constructed from period $t + 1$ on. Thus whether a particular signal in one period is a good signal or bad signal may depend on the choice of punishment schemes in the future. If there is a uniform lower bound on the difference between the largest continuation payoff and the smallest continuation payoff over all public histories in the original PPE, then we can choose λ small enough so that this wedge can be always used to induce truth-telling as before. But it could be the case that this wedge may not be uniformly bounded below. This is not necessarily a pathological case. For example, such a wedge is not needed if a stage game Nash equilibrium is played in the current period. Then the choice of punishment scheme at period t is intrinsically related to the choice of punishment schemes from period $t + 1$ on. This defines mapping from an infinite sequence of punishment scheme to itself. Thus we need to find a

fixed point in infinite dimensional space. We apply Glicksberg fixed point theorem to address this problem.

Once such punishment schemes at each history are constructed, the rest of the proof is exactly the same as before. If player i lies on the equilibrium path, then she is punished by a higher probability of bad public signal. This expected loss from lying is a function of informational variability of ϕ . On the other hand, player i might gain from either when the punishment schemes are not used (ϕ is used) or the punishment schemes for the other players are used. But the expected future gain from such terms is bounded by her informational size with respect to ϕ . Thus if player i 's informational variability is large compared to her informational size, the first effect dominates the second effect, hence each player has incentive to announce her message truthfully on the equilibrium path.

A detailed proof of the following theorem is in the appendix.

Theorem 10 *Fix δ and G . For any $\eta > 0$, there exists $\gamma, \varepsilon > 0$ such that if $v^\phi \leq \gamma \Lambda^\phi$, then every η -uniformly strict pure strategy perfect public equilibrium of $G^\infty(\delta)$ is ε -robust*

Note that γ and ε are taken with respect to all η -uniformly strict pure strategy perfect public equilibria. On the other hand, the coordination device for supporting each equilibrium can vary by definition of ε -robustness.

6 Folk Theorem

We can use our robustness result to obtain a new folk theorem with private monitoring and communication. As we noted, there exist several folk theorems with private monitoring. However, the available sufficient conditions on information structure are not necessarily satisfied in some interesting cases.

Kandori and Matsushima [11] assume that player i 's deviation and player j 's deviation are distinguishable by the rest of players. More precisely, they assumed that for all $a \in A$, $i \neq j$,

$$co(\{p(s_{-i,j}|a'_i, a_{-i}) | a'_i \neq a_i\}) \cap co(\{p(s_{-i,j}|a''_j, a_{-j}) | a''_j \neq a_j\}) = \{p(s_{-i,j}|a)\}.$$

In our example in Section 4, each player's deviation has the same effect on the distribution of the public signal, hence the distribution of the private signals. Thus deviations by two different players are not distinguishable in this sense.

Compte [7] assumes that players' private signals are conditionally independent, that is, for all $a \in A$, i , s_i ,

$$p(s_{-i}|a, s_i) = p(s_{-i}|a).$$

This condition is not satisfied in our example, either. Each player's private signal is correlated with the public signal, hence players' private signals are correlated. Finally, it is clear that our example is not almost public monitoring in the sense of Mailath and Morris [13].

In order to apply our robustness result, we first prove a folk theorem based on uniformly strict perfect public equilibria for repeated games with imperfect public monitoring.¹⁹ The proof of this extension of the Fudenberg, Levine and Maskin ([10]) folk theorem is left to the appendix. We should note here that we use pure minmax payoff instead of the mixed one.

Let $E(\delta, \eta)$ be the set of η -uniformly strict pure perfect public equilibrium payoffs given δ . We borrow the following standard assumptions from Fudenberg, Levine, and Maskin ([10]),²⁰

Assumption A

- All pure action profiles have pairwise full rank.²¹
- The dimension of V^* is n .

With this assumption, the following uniform folk theorem is obtained.

Theorem 11 (*Uniform Folk Theorem*) *Suppose that Assumption A holds. Then, for any smooth set W in the interior of V^* , there exists $\underline{\delta} \in (0, 1)$ and $\eta > 0$ such that $W \in E(\delta, (1 - \delta)\eta)$ for any $\delta \in (\underline{\delta}, 1)$.*

Proof. See the appendix. ■

Combining this folk theorem with uniformly strict equilibrium, our robustness result implies the following folk theorem.

¹⁹It is relatively straightforward to prove a uniform folk theorem (with respect to pure strategy minmax) for repeated games with perfect information following the construction of Fudenberg and Maskin [9].

²⁰This assumption is for convenience and can be replaced with weaker assumptions.

²¹A pure action profile a has pairwise full rank if the distributions of y given a , (a'_i, a_{-i}) , $a'_i \neq a_i$ and (a'_j, a_{-j}) , $a'_j \neq a_j$ have the maximal dimension for all $i, j \neq i$. That is, for all $i, j \neq i$

$$\text{rank}(\{\pi(\cdot | (a'_i, a_{-i})) | a'_i \in A_i\}, \{\pi(\cdot | (a'_j, a_{-j})) | a'_j \in A_j\}) = |A_i| + |A_j| - 1$$

Theorem 12 Fix \hat{G} , $w \in \text{int}V^*$ and $\kappa > 0$. Suppose that there exists G that is 0-close to \hat{G} and satisfies Assumption A. Then there exists $\underline{\delta} \in (0, 1)$ such that, for any $\delta \in (\underline{\delta}, 1)$, there exists γ for which a perfect public equilibrium σ^* of $\hat{G}^\infty(\delta)$ exists and satisfies

$$\|w - V^P(\sigma^*, \delta)\| < \kappa$$

if $v^\phi \leq \gamma \Lambda^\phi$.

Proof. Pick a $\eta(1 - \delta)$ -uniformly strict PPE of $G^\infty(\delta)$ such that its payoff is w . Such a PPE exists for every δ above some critical value $\underline{\delta} \in (0, 1)$ by Theorem 11. For each $\delta \in (\underline{\delta}, 1)$, we can find γ such that if $v^\phi \leq \gamma \Lambda^\phi$, then this PPE is robust. Moreover, the probability of using punishment schemes becomes negligible for small γ , thus the equilibrium payoff will be within $\frac{\kappa}{2}$ of w if γ is chosen small enough. Therefore this PPE (combined with some truth-telling strategy) generates a payoff within κ of w for $\hat{G}^\infty(\delta)$. ■

Remark. 0-closeness cannot be replaced by ε -closeness because the choice of ε depends on the choice of δ in our robustness result. As $\delta \rightarrow 1$, ε needs to converge to 0. If ε is bounded away from 0 as $\delta \rightarrow 1$, then continuation payoffs with private monitoring can be very different from the payoffs with public monitoring because error can accumulate in the long run. Thus, ε needs to be set to 0 to obtain a uniform result with respect to δ .

Note that γ depends on δ . Since we need to approximate the original public distribution more and more closely as $\delta \rightarrow 1$, the probability of punishment schemes (λ) as well as ε needs to converge to 0 as $\delta \rightarrow 1$. This requires γ to be also very small, indeed converging to 0 as $\delta \rightarrow 1$ and $\lambda \rightarrow 0$.

7 Conclusion

As soon as we depart from the assumption of public monitoring, coordination among the players becomes a nontrivial problem even when public monitoring is only slightly perturbed since we lose a common knowledge history on which players can coordinate. The main idea in this paper is to allow players to communicate with each other each period to build a public history on which they can coordinate. After the actions have been taken in a given round and players have received their private signals, they announce those signals. The coordinating device ϕ maps the announcements into a distribution on the set of public signals Y . If the public signal $\pi(y|a)$ is deterministic, there will exist a deterministic ϕ that aggregates private signals

in a way that approximates the public signal. In this case, ϕ is unnecessary as the players can as easily coordinate on the vector of announced signals s as on $\phi(s)$. When $\pi(y|a)$ is nondegenerate, ϕ serves as a public randomizing device, which could be replaced by jointly controlled lotteries.²²

Once there exists at least the possibility of aggregating private signals to approximate the public signal, we investigate the question of robustness of perfect public equilibria in such close private monitoring settings. We identify a sufficient condition on informational size relative to distributional variation under which any η -uniformly strict perfect public equilibrium is robust with respect to such perturbation. How one induces truth-telling is the most critical issue for robustness. When each player is informationally small relative to distributional variability, we can modify the coordination device slightly to induce all players to reveal their true private information, while still generating a public signal that is close to the original public signal. In other words, the vector of private signals can be truthfully elicited, and the players can use the vector of private signals to coordinate similar to the way they coordinate in the perfect public equilibrium.

The robustness of perfect public equilibria holds for any discount factor, but only for uniformly strict perfect public equilibria. We show, however, that the robustness for even for this restricted class of perfect public equilibria is enough to establish a folk theorem.

8 Appendix

8.1 Proof of Proposition 10

- **Step 1: Upper Bound of Continuation Payoff Variations with Private Monitoring**

Assume that players always send their true private signals. We first prove that if p is an ε -approximation of π with ϕ^{h^t} at each history, then any action strategy $\sigma \in \Sigma$ generates almost the same continuation payoff for $G^\infty(\delta)$ and $\hat{G}^\infty(\delta)$ after the same public history for small ε .

Take any $\sigma \in \Sigma$ and let $w_i^\sigma(h^t)$ and $w_{i,p}^\sigma(h^t)$ be player i 's continuation payoff after h^t for $G^\infty(\delta)$ and $\hat{G}^\infty(\delta)$ respectively given σ , $\{\phi^{h^t}\}_{h^t \in \cup_{t=0}^\infty H_t}$, and any truth-telling strategy profile $\sigma^{s*} \in \Sigma^s$. Let $M = \sup_{i,h^t,\sigma} |w_i^\sigma(h^t) - w_{i,p}^\sigma(h^t)|$.

²²See, e.g., Aumann, Maschler and Stearns (1968).

Then,

$$\begin{aligned}
& |w_i^\sigma(h^t) - w_{i,p}^\sigma(h^t)| \\
&= \delta \left| \sum_{y \in Y} w_i^\sigma((h^t, y)) \pi(y|\sigma(h^t)) - \sum_{y \in Y} w_{i,p}^\sigma((h^t, y)) p^{\phi^{h^t}}(y|\sigma(h^t)) \right| \\
&\leq \delta \left| \sum_{y \in Y} w_i^\sigma((h^t, y)) \left\{ \pi(y|\sigma(h^t)) - p^{\phi^{h^t}}(y|\sigma(h^t)) \right\} \right| \\
&\quad + \delta \left| \sum_{y \in Y} \{w_i^\sigma((h^t, y)) - w_{i,p}^\sigma((h^t, y))\} p^{\phi^{h^t}}(y|\sigma(h^t)) \right| \\
&\leq \delta \varepsilon \bar{g} + \delta M
\end{aligned}$$

Taking the supremum of the left hand side, we obtain

$$\begin{aligned}
M &\leq \delta \varepsilon \bar{g} + \delta M \\
&\Updownarrow \\
M &\leq \frac{\delta \varepsilon \bar{g}}{1 - \delta}.
\end{aligned}$$

This implies that continuation payoffs on and off the equilibrium path are almost the same at each history for G^∞ and \hat{G}^∞ if \hat{G}^∞ is ε -close to G^∞ . This implies the following: for a given $\eta > 0$ and $\delta \in (0, 1)$, take any η -uniformly strict PPE σ_i in G^∞ , then there is no incentive to deviate from σ_i in ε -close \hat{G}^∞ if $\varepsilon > 0$ is small enough and players announce their private signals truthfully.

• **Step 2: Truth-Telling (on the equilibrium path)**

We construct $\phi^{h^t} : S \rightarrow \Delta Y$ for each $h^t \in H^t$, each of which is close to ϕ and induces the players to reveal their private signals truthfully at each history on the equilibrium path (that is, after the equilibrium action is taken). In the following, we construct $\phi^{h^{t'}} : S \rightarrow \Delta Y$ for each h^t to induce truthful revelation of private information of all the players and define ϕ^{h^t} as $(1 - \lambda) \phi + \lambda \phi^{h^{t'}}$. Note that, since G is ε close to \hat{G} , p is 2ε -approximation of π with $(1 - \lambda) \phi + \lambda \phi'$ for any $\phi' : S \rightarrow \Delta Y$ if λ is set small enough.

First, define the following functions $\psi_i : A \times S \rightarrow (0, 1)$ using ϕ

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

Next pick a η -strictly uniform perfect public equilibrium σ^* for $G^\infty(\delta)$ and define $\phi_i^{h^{t'}} : S \rightarrow \Delta Y$ for each h^t and i as follows,

$$\phi_i^{h^{t'}}(s) = \psi_i(a^*, s) \cdot \bar{\mu}_i(h^t) + (1 - \psi_i(a^*, s)) \underline{\mu}_i(h^t) \quad (1)$$

where both $\bar{\mu}_i(h^t)$ and $\underline{\mu}_i(h^t)$ are distributions on Y and $a^* = \sigma_i^*(h^t)$. Let Z be the set of all such collection of distributions $(\bar{\mu}_i(h^t), \underline{\mu}_i(h^t))$ on $\Delta Y \times \Delta Y$ for $h^t \in \cup_{t=0}^\infty H^t$ and $i = 1, \dots, n$.

Consider the following correspondence from Z to itself. For each $z = \{(\bar{\mu}_i(h^t), \underline{\mu}_i(h^t))\}_{h^t \in \cup_{t=0}^\infty H^t, i \in N} \in Z$, compute the expected average continuation payoff $V_i^z(h^t)$ for each player i at every public history based on $\{\phi_i^{h^{t'}}(s)\}_{h^t \in \cup_{t=0}^\infty H^t, i \in N}$ defined in (1), assuming that every player plays according to σ^* and announces her private signal truthfully. Then, for each z , assign the set of $z' = \{(\bar{\mu}'_i(h^t), \underline{\mu}'_i(h^t))\}_{h^t \in \cup_{t=0}^\infty H^t, i \in N} \in Z$ which satisfies the following property:

$$\begin{aligned} \bar{\mu}'_i(h^t)(y) &\begin{cases} \geq 0 \text{ if } V_i^z(h^t, y) \geq V_i^z(h^t, y') \text{ for all } y' \neq y \\ = 0 \text{ otherwise} \end{cases} \\ \underline{\mu}'_i(h^t)(y) &\begin{cases} \geq 0 \text{ if } V_i^z(h^t, y) \leq V_i^z(h^t, y') \text{ for all } y' \neq y \\ = 0 \text{ otherwise} \end{cases} \end{aligned}$$

That is, $\bar{\mu}'_i(h^t)$ puts all mass on the y that maximizes $V_i^z(h^t, y)$ and $\underline{\mu}'_i(h^t)$ has all mass on y that minimizes $V_i^z(h^t, y)$.

Since Z is a compact convex subset of a locally convex space ($= \mathfrak{R}^\infty$) and the above correspondence is nonempty, convex-valued, and closed-graph (in product topology), we can apply Glicksberg Fixed Point Theorem (Theorem 16.51 [?]) to obtain $z^* = \left\{ \left(\bar{\mu}_i^*(h^t), \underline{\mu}_i^*(h^t) \right) \right\}_{h^t \in \bigcup_{t=0}^\infty H^t, i \in I}$ to satisfy

$$\begin{aligned} \bar{\mu}_i^*(h^t)(y) & \begin{cases} \geq 0 & \text{if } V_i^{z^*}(h^t, y) \geq V_i^{z^*}(h^t, y') \text{ for all } y' \neq y \\ = 0 & \text{otherwise} \end{cases} \\ \underline{\mu}_i^*(h^t)(y) & \begin{cases} \geq 0 & \text{if } V_i^{z^*}(h^t, y) \leq V_i^{z^*}(h^t, y') \text{ for all } y' \neq y \\ = 0 & \text{otherwise} \end{cases} \end{aligned}$$

Define $\phi_i^{h^t*}$ by (1) using $\left\{ \left(\bar{\mu}_i^*(h^t), \underline{\mu}_i^*(h^t) \right) \right\}_{h^t \in \bigcup_{t=0}^\infty H^t}$. Let $\phi^{h^t} = (1 - \lambda)\phi + \lambda\phi^{h^t'}$ where $\phi^{h^t'} = \frac{1}{n} \sum_{i=1}^n \phi_i^{h^t*}$. This means that ϕ is used to interpret a profile of message with probability $1 - \lambda$ and, when ϕ is not used, each $\phi_i^{h^t*}$ is used with equal probability. When $\phi_i^{h^t*}$ is used, $\bar{\mu}_i^*(h^t)$ is used with probability $\psi_i(a^*, s)$ and $\underline{\mu}_i^*(h^t)$ with probability $1 - \psi_i(a^*, s)$.

We show that the players do not have incentive to lie on the equilibrium path with such $\left\{ \phi^{h^t} \right\}_{h^t \in \bigcup_{t=0}^\infty H^t}$, if the informational size is sufficiently small relative to the distributional variability with respect to ϕ . Since player i 's expected average continuation payoff associated with the term $\phi_i^{h^t*}$ given s is

$$\begin{aligned} & \psi_i(a^*, s) \sum_{y \in Y} \bar{\mu}_i^*(h^t)(y) V_i^{z^*}(h^t, y) + (1 - \psi_i(a^*, s)) \sum_{y \in Y} \underline{\mu}_i^*(h^t)(y) V_i^{z^*}(h^t, y) \\ = & \psi_i(a^*, s) \left(\max_{y \in Y} V_i^{z^*}(h^t, y) - \min_{y \in Y} V_i^{z^*}(h^t, y) \right) + \min_{y \in Y} V_i^{z^*}(h^t, y) \end{aligned}$$

, player i 's expected future loss from this term by lying is obtained as follows

$$\begin{aligned}
& \Delta \left\{ \sum_{s_{-i} \in S_{-i}} (\psi_i(a^*, s) - \psi_i(a^*, s'_i, s_{-i})) \cdot p(s_{-i}|a^*, s_i) \right\} \\
&= \Delta \sum_{s_{-i} \in S_{-i}} \left\{ \frac{\sum_{y \in Y} \frac{p^\phi(y|a^*, s_i)}{\|p^\phi(\cdot|a^*, s_i)\|} \cdot \phi(y|s)}{-\sum_{y \in Y} \frac{p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s'_i)\|} \cdot \phi(y|s'_i, s_{-i})} \right\} p(s_{-i}|a^*, s_i) \\
&= \Delta \left\{ \sum_{y \in Y} \frac{p^\phi(y|a^*, s_i)}{\|p^\phi(\cdot|a^*, s_i)\|} \cdot p^\phi(y|a^*, s_i) - \sum_{y \in Y} \frac{p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s'_i)\|} \cdot p^\phi(y|a^*, s_i, s'_i) \right\} \\
&= \Delta \left\{ \sum_{y \in Y} \left(\frac{p^\phi(y|a^*, s_i)}{\|p^\phi(\cdot|a^*, s_i)\|} - \frac{p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s'_i)\|} \right) p^\phi(y|a^*, s_i) \right. \\
&\quad \left. + \sum_{y \in Y} \frac{p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s'_i)\|} (p^\phi(y|a^*, s_i) - p^\phi(y|a^*, s_i, s'_i)) \right\}
\end{aligned}$$

where $\Delta = \max_{y \in Y} V_i^{z*}(h^t, y) - \min_{y \in Y} V_i^{z*}(h^t, y) \geq 0$.

Since $p^\phi(y|a^*, s_i) - p^\phi(y|a^*, s_i, s'_i)$ is bounded below by $-2v^\phi$ (by Lemma 7), this is at least

$$\begin{aligned}
& \Delta \left\{ \sum_{y \in Y} \left(\frac{p^\phi(y|a^*, s_i)}{\|p^\phi(\cdot|a^*, s_i)\|} - \frac{p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s'_i)\|} \right) p(y|a^*, s_i) - \frac{2}{\|p^\phi(\cdot|a^*, s'_i)\|} v^\phi \right\} \\
&= \Delta \left\{ \left\| p^\phi(\cdot|a^*, s_i) \right\| \left(1 - \frac{p^\phi(y|a^*, s_i) \cdot p^\phi(y|a^*, s'_i)}{\|p^\phi(\cdot|a^*, s_i)\| \|p^\phi(\cdot|a^*, s'_i)\|} \right) - \frac{2}{\|p^\phi(\cdot|a^*, s'_i)\|} v^\phi \right\} \\
&\geq \Delta \left(\frac{1}{2\sqrt{m}} \Lambda^\phi - 2\sqrt{m} v^\phi \right)
\end{aligned}$$

assuming $\Lambda^\phi > 4m v^\phi$.

Player i might gain from the other terms $\left(= (1 - \lambda) \phi + \frac{\lambda}{n} \sum_{j \neq i} \phi_j^{h^t*} \right)$ by lying. We can compute a bound on the gain from the first term $(= \phi)$.

Let Y' be the subset of Y such that $p^\phi(y|a^*, s_i, s'_i) - p^\phi(y|a^*, s_i) > 0$. Then

$$\begin{aligned}
& \sum_{s_{-i} \in S_{-i}} \left[\sum_{y \in Y} \{ \phi(y|s_i, s'_{-i}) - \phi(y|s) \} V_i^{z*}(h^t, y) \right] p(s_{-i}|a^*, s_i) \\
&= \sum_{y \in Y} \left(p^\phi(y|a^*, s_i, s'_i) - p^\phi(y|a^*, s_i) \right) V_i^{z*}(h^t, y) \\
&\leq \sum_{y \in Y'} \left(p^\phi(y|a^*, s_i, s'_i) - p^\phi(y|a^*, s_i) \right) \max_{y \in Y} V_i^{z*}(h^t, y) \\
&\quad + \sum_{y \in Y/Y'} \left(p^\phi(y|a^*, s_i, s'_i) - p^\phi(y|a^*, s_i) \right) \min_{y \in Y} V_i^{z*}(h^t, y) \\
&\leq \sum_{y \in Y'} \left(p^\phi(y|a^*, s_i, s'_i) - p^\phi(y|a^*, s_i) \right) \Delta \\
&\leq mv^\phi \Delta
\end{aligned}$$

As for the gain from $\phi_j^{h^t*}$, it should be at most

$$\begin{aligned}
& \Delta \sum_{s_{-i} \in S_{-i}} \left[\sum_{y \in Y} \frac{p^\phi(y|a^*, s_j)}{\|p^\phi(\cdot|a^*, s_j)\|} (\phi(y|s'_i, s_{-i}) - \phi(y|s)) \right] p(s_{-i}|a^*, s_i) \\
&\leq \Delta \sum_{s_{-i} \in S_{-i}} \left[\left\| \frac{p^\phi(\cdot|a^*, s_j)}{\|p^\phi(\cdot|a^*, s_j)\|} \right\| \|\phi(s'_i, s_{-i}) - \phi(s)\| \right] p(s_{-i}|a^*, s_i) \\
&= \Delta E [\|\phi(a, s'_i, s_{-i}) - \phi(a, s)\| | a^*, s_i] \\
&\leq \Delta (1 + \sqrt{2}) v^\phi \text{ (By Lemma 7).}
\end{aligned}$$

Therefore, player i' 's expected gain by lying from such terms is bounded by $\Delta \{ (1 - \lambda) mv^\phi + (1 + \sqrt{2}) \lambda \frac{n-1}{n} v^\phi \}$. Thus all the truth-telling constraints are satisfied if the following condition holds:

$$\left\{ (1 - \lambda) m + (1 + \sqrt{2}) \lambda \frac{n-1}{n} \right\} v^\phi \leq \frac{\lambda}{n} \left(\frac{1}{2\sqrt{m}} \Lambda^\phi - 2\sqrt{m} v^\phi \right).$$

Note that this condition is independent of the equilibrium we pick although the definition of ϕ' depends on the particular equilibrium. This condition is satisfied if γ is chosen so that $0 < \gamma < \frac{\lambda}{2(1-\lambda)m^{\frac{3}{2}}n+2(1+\sqrt{2})\lambda\sqrt{m}(n-1)+4\lambda m}$.

• **Step 3:**

For any $\eta > 0$ and $\delta \in (0, 1)$, we can choose ε and λ small enough so that any η -uniformly strict perfect public equilibrium of $G^\infty(\delta)$ is a perfect public equilibrium of $\hat{G}^\infty(\delta)$ given truth-telling by Step 1. Fix such ε and λ . Then we can choose γ as in Step 2 to guarantee truth-telling (on the equilibrium path). Then there is no incentive for one shot deviation with respect to both action and announcement. Finally, as we discussed in the example, any joint deviation with respect to action and announcement within a period is not profitable if the informational size of each player is small enough because of η -uniform strictness. γ can be chosen to satisfy this requirement as well. Then, by the principle of optimality, any η -uniformly strict perfect public equilibrium of $G^\infty(\delta)$ is a perfect public equilibrium of $\hat{G}^\infty(\delta)$ combined with some truth-telling strategy σ^{s*} . This completes the proof. ■

8.2 Proof of Theorem 11 (Uniform Folk Theorem)

We need the following definitions.

Definition 13 Given δ , $a \in A$ is η -enforceable with respect to $W \subset \mathbb{R}^n$ if there exists a function $w : Y \rightarrow W$ such that, for all i , player i loses more than $\eta > 0$ by playing $a'_i \neq a_i$.

Definition 14 v is η -generated by W if there exists $a \in A$ which is η -enforceable with $w : Y \rightarrow W$ and $v = (1 - \delta)g(a) + \delta E[w|a]$. The set of values η -generated by W for a given δ is denoted by $B(\delta, W, \eta)$.

Definition 15 $H(\lambda, k) = \{x \in \mathbb{R}^n | \lambda \cdot x \leq k\}$ for $\lambda \in \mathbb{R}^n, k \in \mathbb{R}$.

Let us first prove the following lemma:

Lemma 16 If $C \subset B(\delta, W, \eta) \cap W$ for a convex set W , then $C \subset B(\delta', W, (1 - \delta')\eta)$ for $\delta' \in (\delta, 1)$.

Proof. Suppose that $v \in C$. Then, there exists a which is η -enforceable with respect to W and generates v with $w_\delta(y)$. Fix $\delta' > \delta$. Now define $w_{\delta'} : Y \rightarrow W$ as the following linear combination of v and w_δ ,

$$w_{\delta'}(y) = \frac{\delta' - \delta}{\delta'(1 - \delta)}v + \frac{\delta(1 - \delta')}{\delta'(1 - \delta)}w_\delta(y) \in W.$$

Then, it is easily confirmed that a is $\frac{1 - \delta'}{1 - \delta}\eta$ -enforceable (thus $(1 - \delta')\eta$ -enforceable) with respect to W and generates v with $w_{\delta'}(y)$ for δ' . ■

Remark 17 1. Note that $\|v - w_{\delta'}(y)\| = \frac{\delta(1 - \delta')}{\delta'(1 - \delta)}\|v - w_\delta(y)\|$

2. Suppose that v is η -generated with (a, w) . It is useful to express w in the following way:

$$w_\delta(y) = \bar{w}_\delta + x_\delta(y)$$

where $\bar{w}_\delta = E[w_\delta(y)|a]$ satisfies $v = (1 - \delta)g(a) + \delta\bar{w}_\delta$. By definition, $E[x_\delta(y)|a] = \mathbf{0}$. Then, $w_{\delta'}(y)$ can be expressed as

$$w_{\delta'}(y) = \bar{w}_{\delta'} + \frac{\delta(1 - \delta')}{\delta'(1 - \delta)}x_\delta(y)$$

where $\bar{w}_{\delta'} \in \mathbb{R}^n$ satisfies $v = (1 - \delta')g(a) + \delta'\bar{w}_{\delta'}$.

Now we return to the proof of the theorem.

Proof. Step 1: Self Generation Implies Equilibrium: For a bounded $W \subset \mathbb{R}^n$, $W \subset B(\delta, W, \eta) \Rightarrow W \subset E(\delta, \eta)$.

The proof is similar to the proof in Abreu, Pearce and Stacchetti [1] and is omitted.

Step2: Local uniform strictness is enough.

By step 1 and by Lemma 16, we only need to show that any smooth W in the interior of V^* is included in $B(\delta, W, \eta)$ for some $(\delta, \eta) \in (0, 1) \times (0, \infty)$ to prove the theorem. As in Fudenberg, Levine, and Maskin [10], we show that it is enough to verify a local property because of the compactness of W .

Suppose that for each $v \in W$, there exist $\delta_v \in (0, 1)$ and an open neighborhood U_v of v such that

$$U_v \cap W \subset B(\delta_v, W, \eta_v) \quad (2)$$

for some $\eta_v > 0$. Since W is a compact set, we can take a finite subcover $\{U_{v_i}\}_{i=1}^K$ of W . Define $\underline{\delta} = \max_{i=1, \dots, K} \{\delta_{v_i}\}$ and $\underline{\eta} = \min_{i=1, \dots, K} \{\eta_{v_i}\}$. Then, by the above lemma, $U_{v_i} \cap W \subset B(\underline{\delta}, W, (1 - \underline{\delta})\underline{\eta})$ for $i = 1, \dots, K$. Hence, $\cup_{i=1}^K (U_{v_i} \cap W) = W \subset B(\underline{\delta}, W, (1 - \underline{\delta})\underline{\eta})$.

Step3: Local η -enforceability

We focus on showing local η -enforceability ((2)) at each boundary point of W as the same proof applies to any interior point as well.. Take any $v \in \partial W$. Then, there exists unique outward normal vector λ ($\|\lambda\| \neq 0$). First we show that v is η -enforceable with respect to a half space tangent to the hyperplane at v for some $\eta > 0$. There are two cases.

Case 1: There are at least two components of λ which are nonzero.

Fix any $\delta \in (0, 1)$. It follows from Fudenberg, Maskin and Levine [10] that (1) v can be η -generated with respect to the tangent hyperplane at v for some $\eta > 0$ with some $(a, w_\delta(y))$, and moreover, (2) v is $(1 - \delta')\eta$ -generated with respect to the same hyperplane with $(a, w_{\delta'}(y))$ for any $\delta' \in (\delta, 1)$ and $\|v - w_{\delta'}(y)\|$ converges to 0 at the rate of $(1 - \delta')$ as $\delta' \rightarrow 1$ by Lemma 16 and the above remark.

Case 2: $\exists i \in I$ such that $\lambda_j = 0$ for any $j \neq i$.

Suppose that $\lambda_i > 0$ ($\lambda_i < 0$ can be treated in a similar way). Fix any $\delta \in (0, 1)$. Pick $a \in A$ such that $\lambda \cdot g(a) > \lambda \cdot v$ and player i is playing the best response at a . Since the individual full rank condition is satisfied (which is weaker than pairwise full rank), we can find $w_{\delta,j}$ satisfying (1) and (2) above for any $j \neq i$. However, it is not straightforward to find such $w_{\delta,i}$ for player i . This is because $w_{\delta,i}$ is picked from the tangent hyperplane in the above construction, which requires x_i to be constant. Then if a_i is not a strict best response in the stage game, player i does not have a strict incentive to play a .

Nonetheless we can show that a is η -enforceable for some $\eta > 0$ with respect to the half space $H(\lambda, \lambda \cdot (\frac{1}{2}\bar{w}_\delta + \frac{1}{2}v))$ ($\supset H(\lambda, \lambda \cdot \bar{w}_\delta)$), while generating v with some $w_{\delta,i} = \bar{w}_\delta + x_i$. As $\lambda_j = 0$ for $j \neq i$, we only need to show that x_i can be chosen so that $\lambda_i \cdot (\bar{w}_{\delta,i} + x_i(y)) \leq \lambda_i \cdot (\frac{1}{2}\bar{w}_{\delta,i} + \frac{1}{2}v_i)$, that is, $x_i(y) \leq \frac{1}{2}(v_i - \bar{w}_{\delta,i})$ (> 0) for every $y \in Y$. This can be easily done by finding x_i to satisfy $E[x_i|a] = 0 > E[x_i|a'_i, a_{-i}]$ and multiplying it by some sufficiently small positive number. Note that player i 's incentive constraint holds strictly because player i is playing a best response action at a . Then, following Lemma 16 and the above remark, we can define $w_{\delta'}(y)$ as before so that v is be $(1 - \delta')$ - η -generated with respect to $H(\lambda, \lambda_i \cdot (\bar{w}_{\delta'} + \frac{\delta(1-\delta')}{\delta'(1-\delta)} \times \frac{v_i - \bar{w}_{\delta,i}}{2}))$ for some $\eta > 0$ for $\delta' \in (\delta, 1)$. Note that $w_{\delta'}(y)$ and $H(\lambda, \lambda_i \cdot (\bar{w}_{\delta'} + \frac{\delta(1-\delta')}{\delta'(1-\delta)} \times \frac{v_i - \bar{w}_{\delta,i}}{2}))$ converge to v and $H(\lambda, \lambda \cdot v)$ respectively at the rate of $1 - \delta'$ as $\delta' \rightarrow 1$.

In either case, the distance between $w_{\delta'}(y)$ and v is converging to 0 at the same speed as $(1 - \delta')$. Thus we can take a nested sequence of open balls $B_{\delta'}(v)$ around v for each δ' so that its radius is proportional to $1 - \delta'$ and $w_{\delta'}(y)$ is contained in $B_{\delta'}(v)$. Pick any point in $H(\lambda, \lambda \cdot v) \cap B_{\delta'}(v)$ but outside of W . Since W is smooth, such points converge to the tangent hyperplane at v at the rate of as $(1 - \delta')^2$. On the other hand, all $w_{\delta'}(y)$ are separated away from the tangent hyperplane at v by the order of $(1 - \delta')$. Since Y is finite, this implies that there exists δ'' such that $w_{\delta_v}(y) \in \text{int}(W)$ for any $\delta_v \in (\delta'', 1)$. Thus v can be $(1 - \delta')$ - η -generated with respect to W for large enough δ .

Now pick any such δ_v and take $r \in (0, 1)$ so that $w_{\delta_v}(y) + h \in \text{int}(W)$ for all $\|h\| < \frac{r}{\delta_v}$. Since the incentive constraints are not affected even when a small constant vector is added to the continuation payoff profile, a is still $(1 - \delta_v)$ - η -enforceable with respect to W with $w_{\delta_v}(y) + h$ and generates $v + \delta_v h$, which implies that $B_r(v) \cap W \subset B(\delta_v, W, \eta_v)$ where $\eta_v = (1 - \delta_v)\eta$. \blacksquare

References

- [1] Abreu, D., D. Pearce, and E. Stacchetti, “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring”, *Econometrica* **59**:1041-1063, 1990.
- [2] Aliprantis, C.D. and K. Border, “Infinite Dimensional Analysis: A Hitchhiker’s Guide”, Springer-Verlag, 1999
- [3] Anderlini L. and R. Lagunoff, “Communication in Dynastic Repeated Games: ‘Whitewashes’ and ‘Coverups’”, forthcoming in *Economic Theory*.
- [4] Aoyagi, M., “Communication Equilibria in Repeated Games with Imperfect Private Monitoring”, mimeo, 2000.
- [5] Aumann, R. and L. Shapley, “Long-Term Competition: A Game Theoretic Analysis”, mimeo, 1976.
- [6] Ben-Porath, E. and M. Kahneman, “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory* **70**:281-297, 1996.
- [7] Compte, O., “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, **66**:597-626, 1998.
- [8] Fudenberg, D. and D. Levine, “The Nash Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games”, mimeo, 2002.
- [9] Fudenberg, D. and E. Maskin, “The Folk Theorem in Repeated Games with Discounting of Incomplete Information”, *Econometrica* **54**:533-554, 1986
- [10] Fudenberg, D., D. Levine, and E. Maskin, “The Folk Theorem with Imperfect Public Information”, *Econometrica* **62**:997-1039, 1994.
- [11] Kandori, M. and H. Matsushima, “Private Observation and Collusion”, *Econometrica*, **66**:627-652, 1998.
- [12] Kandori, M. and I. Obara, “Efficiency in Repeated Games revisited: the Role of Private Strategies, UCLA discussion paper, 2003.”
- [13] Mailath, G. and S. Morris, “Repeated Games with Almost-Public Monitoring”, *Journal of Economic Theory*, **102**:189-228, 2002.

- [14] Mailath, G. and S. Morris, “Coordination Failure in Repeated Games with Almost-public Monitoring”, PIER working paper 04-033, 2004
- [15] McLean, R. and A. Postlewaite, “Informational Size and Incentive Compatibility,” *Econometrica*, 70: 2421-2454 2002.
- [16] Rubinstein, A. “Equilibrium in Supergames with the Overtaking Criterion”, *Journal of Economic Theory*, **27**:1-9, 1979.