

MIT Open Access Articles

Efficiency in Games With Markovian Private Information

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Escobar, Juan F., and Juuso Toikka. "Efficiency in Games With Markovian Private Information." *Econometrica* 81.5 (2013): 1887–1934.

As Published: <http://dx.doi.org/10.3982/ecta9557>

Publisher: The Econometric Society

Persistent URL: <http://hdl.handle.net/1721.1/85872>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



EFFICIENCY IN GAMES WITH MARKOVIAN PRIVATE INFORMATION

JUAN F. ESCOBAR AND JUUSO TOIKKA

ABSTRACT. We study repeated Bayesian games with communication and observable actions in which the players' privately-known payoffs evolve according to an irreducible Markov chain whose transitions are independent across players. Our main result implies that, generically, any Pareto-efficient payoff vector above a stationary minmax value can be approximated arbitrarily closely in a perfect Bayesian equilibrium as the discount factor goes to one. As an intermediate step we construct an approximately efficient dynamic mechanism for long finite horizons without assuming transferable utility.

1. INTRODUCTION

Repeated Bayesian games, also known as repeated games of adverse selection, provide a model of long-run relationships where the parties have asymmetric information about their objectives.¹ It is well known that if each player's payoff-relevant private information, or *type*, is independently and identically distributed (iid) over time, repeated play can facilitate cooperation beyond what is achievable in a one-shot interaction.² In particular, the folk theorem of Fudenberg, Levine, and Maskin (1994) implies that first-best efficiency can be approximately achieved as the discount factor tends to one.

However, the iid assumption on types appears restrictive in many applications. For example, firms in an oligopoly or bidders in a series of procurement auctions may have private information about production costs, which tend to be autocorrelated. Furthermore, under the iid assumption, there is asymmetric information only about current payoffs. In contrast, when types are serially dependent, the players also have private information about the distribution of future payoffs. This

Date: January 11, 2013. We thank Manuel Amador, Kyle Bagwell, Aaron Bodoh-Creed, Matt Elliot, Johannes Hörner, Matt Jackson, Carlos Lever, Jon Levin, Romans Pancs, Andy Skrzypacz, and especially Ilya Segal for useful discussions. We are grateful to the editor and four anonymous referees for comments. Escobar acknowledges the support of Fondecyt Project 11090166 and the Millennium Institute on Complex Engineering Systems. Toikka acknowledges financial support from the Yrjö Jahnsson Foundation. An earlier version of this work was circulated as "A Folk Theorem with Markovian Private Information." Contact: jescobar@dii.uchile.cl or toikka@mit.edu.

¹For economic applications, see, e.g., Athey, Bagwell, and Sanchirico (2004) on oligopolistic competition, Levin (2003) on relational contracting, Maggi and Morelli (2006) on voting in international organizations, or Skrzypacz and Hopenhayn (2004) on repeated auctions.

²See, e.g., Mailath and Samuelson (2006, Chapter 11)

introduces the possibility of signaling, which presents a new challenge for cooperation as a player may be tempted to alter his behavior to influence the other players' beliefs about his type.

In this paper, we study the problem of sustaining cooperation among patient players in repeated Bayesian games with serially dependent types. More specifically, we assume that the type profile follows an autonomous irreducible Markov chain where the evolution of types is independent across players. Focusing on the case of private values (also known as the known-own-payoffs case) and observable actions, we define the stationary minmax value as the lowest payoff that can be imposed on a patient player by having the other players play a constant pure-action profile. Our main result then shows that with cheap-talk communication, generically, any payoff profile v , which lies in (an appropriately defined convex superset of) the Pareto frontier and dominates the stationary minmax profile, can be approximately attained in a perfect Bayesian equilibrium if the players are sufficiently patient. Furthermore, the equilibrium can be taken to be stationary in the sense of the players' expected continuation payoffs remaining close to v along the equilibrium path.

The key step in our proof of this limit-efficiency result is the derivation of an analogous result for auxiliary reporting games where players communicate as in the original game, but actions are selected by a mechanism. We introduce the *credible reporting mechanism* for which payoffs can be bounded uniformly across equilibria. This allows us to assert the existence of equilibria with the desired payoffs without solving the players' best-response problems or tracking beliefs.³ The rest of the proof then extends a reporting-game equilibrium into an equilibrium of the original game by establishing the existence of player-specific punishments analogous to the stick-and-carrot schemes of Fudenberg and Maskin (1986).

Our credible reporting mechanism is of independent interest in that it gives an approximately efficient dynamic mechanism for patient players without assuming transferable utility.⁴ It uses a statistical test to assess whether the players' reports about their types are sufficiently likely to have resulted from truthful reporting. (If a player fails the test, his reports are henceforth replaced with appropriately chosen random messages.) The construction is inspired by the linking mechanism of Jackson and Sonnenschein (2007), who use message budgets to force the long-run distribution of each player's reports to match his true type distribution.⁵

³While our proof is virtually free of beliefs, the equilibria we identify are in general not “belief free” or “ex post” in the sense of Hörner and Lovo (2009), Hörner, Lovo, and Tomala (2011), or Fudenberg and Yamamoto (2011b).

⁴When utility is transferable, surplus-maximizing decision rules can be implemented using a dynamic VCG scheme—see Athey and Segal (2007) and Bergemann and Välimäki (2010).

⁵In the single-agent version the idea goes back at least to Radner (1981) and Townsend (1982). It is used in the context of a repeated Bayesian game with iid types by Hörner and Jamison (2007). Independently of our work, Renault, Solan, and Vieille (2011) use the linking mechanism

With iid types, their mechanism approximately implements efficient choice rules given a long enough horizon and sufficient patience. However, with Markov types, a player may be able to use his opponents' past reports to predict their future types. This gives rise to contingent deviations, which may undermine the linking mechanism. Our mechanism rules out such deviations by testing for the convergence of appropriately chosen conditional distributions.

Games with Markovian types were introduced in the reputation literature under the assumption that the player is replaced upon type change (see, e.g., Cole, Dow, and English, 1995, Mailath and Samuelson, 2001, or Phelan, 2006). Few papers consider Markovian types without replacements. The first results are due to Athey and Bagwell (2008) who analyze collusion in a Bertrand oligopoly with privately-known costs. They provide an example of a symmetric duopoly with irreducible binary costs where first-best collusion can be exactly achieved given sufficiently little discounting. Athey and Segal (2007) prove an efficiency result for a class of ergodic Markov games with transfers by showing that their Balanced Team Mechanism can be made self-enforcing if the players are sufficiently patient and there exists a "static" punishment equilibrium.

For the special case of iid types equilibrium payoffs can be characterized for a fixed discount factor using the methods of Abreu, Pearce, and Stacchetti (1990).⁶ A version of the folk theorem and a characterization of the limit equilibrium payoff set is provided by Fudenberg, Levine, and Maskin (1994) and Fudenberg and Levine (1994). At the other extreme, the results of Myerson and Satterthwaite (1983) imply that in the limiting case of perfectly persistent types there are games where equilibrium payoffs are bounded away from the first best even when players are arbitrarily patient.⁷ Together with our results this shows that the set of equilibrium payoffs may be discontinuous in the joint limit when the discount factor tends to one and the chain becomes perfectly persistent.

Finally, our game can be viewed as a stochastic game with asymmetric information about the state. Dutta (1995), Fudenberg and Yamamoto (2011a), and Hörner, Sugaya, Takahashi, and Vieille (2011) prove increasingly general versions of the folk theorem for stochastic games with a public irreducible state.

to characterize the limit set of equilibrium payoffs in sender-receiver games where the sender's type follows an ergodic Markov chain. The linking mechanism suffices in their case despite serial dependence because private information is one-sided. Our results are not stronger nor weaker than theirs since we assume private values whereas their game has interdependent values.

⁶Cole and Kocherlakota (2001) extend the approach to a class of games that includes ours (see also Athey and Bagwell, 2008). They operate on pairs of type-dependent payoff profiles and beliefs. Due to the inclusion of beliefs the characterization is difficult to put to work.

⁷Starting with the work of Aumann and Maschler (1995) on zero-sum games, there is a sizable literature on perfectly persistent types (e.g., Athey and Bagwell, 2008, Fudenberg and Yamamoto, 2011b, Hörner and Lovo, 2009, Peski, 2008, or Watson, 2002). Such models are also used in the reputation literature (e.g., Kreps and Wilson, 1982 and Milgrom and Roberts, 1982).

The next section illustrates our argument in the context of a simple Bertrand game. Section 3 introduces the model. The main result is presented in Section 4. The proof is developed in Sections 5 and 6 with 5 devoted to the mechanism design problem and 6 to the construction of equilibria. Section 7 concludes.

2. AN EXAMPLE

Consider repeated price competition between two firms, 1 and 2, whose privately known marginal costs are $\theta_1 \in \{L, H\}$ and $\theta_2 \in \{M, V\}$, respectively, with $L < M < H < V$ (for “low, medium, high, and very high”). Firm i 's ($i = 1, 2$) cost follows a Markov chain in which with probability $p \in]0, 1[$ the cost in period $t + 1$ is the same as in period t . The processes are independent across firms. In each period there is one buyer with unit demand and reservation value $r > V$. Having privately observed their current costs, the firms send reports to each other, quote prices, and the one with the lower price serves the buyer, provided its price does not exceed the buyer's reservation value.

This duopoly is a special case of the game introduced in Section 3, and thus our efficiency result (Theorem 4.1 and Corollary 4.1) applies. To illustrate the proof, we sketch the argument showing that, given sufficiently little discounting, there are equilibria with profits arbitrarily close to the first-best collusive scheme where in each period the firm with the lowest cost makes a sale at the monopoly price r .

2.1. A Mechanism Design Problem. Assume first that the horizon T is large but finite, and firms do not discount profits. Assume further that the firms only send cost reports and some mechanism automatically sets the price r for the firm that reported the lowest cost and $r + 1$ for the other firm.

If both firms report their costs truthfully, then, for T large, firm 2 makes a sale in approximately $\frac{T}{4}$ periods. The resulting (average) profits are approximately $v_1 = \frac{r-L}{2} + \frac{r-H}{4}$ for firm 1, and $v_2 = \frac{r-M}{4}$ for firm 2.

Note that if firm i is truthful and firm $j \neq i$ reports as if it were truthful, then firm i 's profit is still approximately v_i because the true cost of firm j does not directly enter i 's payoff (i.e., we have private values). With this motivation, consider a mechanism which tests in real time whether the firms are sufficiently likely to have been truthful. If a firm fails the test, then from the next period onwards its messages are replaced with random messages generated by simulating the firm's cost process. In particular, suppose that (with high probability):

- (1) A truthful firm passes the test regardless of the other firm's strategy.
- (2) The distribution of (simulated) reports faced by a truthful firm is the same as under mutual truth-telling regardless of the other firm's strategy.

Then each firm i can secure an expected profit close to v_i , say $v_i - \varepsilon$, simply by reporting truthfully, and hence in any equilibrium its expected profit is at least

	M	V
L	$.5(1-p)$	$.5p$
H	$.5p$	$.5(1-p)$

FIGURE 1. Distribution of reports when firm 1 reports truthfully and firm 2 reports M iff firm 1 reported H in the previous period.

$v_i - \varepsilon$. As the profile $v = (v_1, v_2)$ is Pareto efficient, this implies that the set of expected equilibrium profits in the mechanism is concentrated near v .

This observation is established for the general model in Theorem 5.1 by considering *credible reporting mechanisms*, which satisfy appropriate formalizations of (1) and (2). To motivate the construction, it is instructive to start with iid costs.

Jackson and Sonnenschein (2007) show that when costs are iid (i.e., $p = \frac{1}{2}$), a mechanism that satisfies (1) and (2) can be obtained by assigning to each firm a budget of reports with $\frac{T}{2}$ reports of each type. This *linking mechanism* approximately satisfies (1) for T large enough by the law of large numbers. As for (2), it is useful to decompose the condition into the following requirements:

(2a) The marginal distribution of each firm's reports matches the truth.

(2b) The distribution faced by a truthful firm is the product of the marginals.

Then (2a) follows by construction of the budgets. For (2b), suppose that, say, firm 1 reports truthfully. Then its reports are a sequence of iid draws, generated independently of firm 2's costs. Therefore, it is impossible for firm 2 to systematically correlate its reports with those of firm 1. This establishes (2b).

When costs are autocorrelated (i.e., $p \neq \frac{1}{2}$), the budgets in the linking mechanism still give (1) and (2a). However, (2b) fails as a firm can use its competitor's past reports to predict its current report. For example, suppose firm 1 reports truthfully but firm 2 deviates and reports M if and only if firm 1 reported H in the previous period. This leads to the distribution depicted in Figure 1, whereas under mutual truth-telling each report profile has probability $\frac{1}{4}$ for all p .⁸

In order to rule out the above problem, we require each firm's conditional reporting frequencies to mirror the true conditional probabilities of the cost process. (Note that simply augmenting the linking mechanism by testing for independence of reports across firms yields a mechanism that fails property (1).) In particular, our credible reporting mechanism tests, for every fixed profile of previous period reports (θ_1, θ_2) and a current report θ_j^t for firm j , whether the reports of firm $i \neq j$ are sufficiently likely to have resulted from truthful reporting. For example, along the random subsequence of periods in which firm 1 reported H and the previous-period reports were (H, V) , under the null hypothesis of truth-telling, firm 2's

⁸Note that the deviation leads to firm 2 making the sale in approximately $p\frac{T}{2}$ periods, earning an expected profit $p\frac{r-M}{4} + p\frac{r-V}{4}$. Thus the deviation is profitable for p large enough, and hence approximate truth-telling is in general not an equilibrium of the linking mechanism for $p \neq \frac{1}{2}$.

reports are iid with the probability of L being $1 - p$. Thus the test requires that along the subsequence, firm 2 reports L with a frequency converging to $1 - p$ at a pre-specified rate. This rules out the deviation contemplated above.⁹

Lemma 5.1 shows that the credible reporting mechanism indeed satisfies (1) and (2). The proof takes some work because of the need to consider arbitrary strategies for non-truthful players, but the result is conceptually straightforward: (2) follows as a firm either passes the test and hence satisfies (2), or fails the test and has its reports replaced by simulated reports which satisfy (2) by construction. As for (1), note that firm i passes the test if the marginal distribution of its reported cost transitions converges to the truth, and these transitions appear sufficiently independent of those of firm $j \neq i$. If i is truthful, then the first part is immediate, and the second follows as the test conditions on a previous-period report profile so that, by the Strong Markov property, the argument reduces to the iid case.

2.2. From the Mechanism to Game Equilibria. In order to construct equilibria of the original pricing game, we need to introduce discounting, extend the result to an infinite horizon, and make the mechanism self-enforcing.

Discounting can be introduced simply by continuity since the mechanism design problem has a finite horizon.

We cover the infinite horizon by having the firms repeatedly play the credible reporting mechanism over T -period blocks. This serves to guarantee that continuation profits are close to the target at all histories. It is worth noting that because of autocorrelation in costs, it is not possible to treat adjacent blocks independently of each other. However, the lower bound on profits from truthful reporting applies to each block and provides a bound on the continuation profits in all equilibria of the block mechanism (Corollary 5.1).

Finally, an equilibrium of the block mechanism is extended to an equilibrium of the original game by constructing off-path punishments that take the form of stick-and-carrot schemes. For example, if firm 1 deviates by quoting a price different from what the mechanism would have chosen, then firm 2 prices at L during the stick while firm 1 best responds. The carrot is analogous to the cooperative phase, but has firm 2 selling—still at the monopoly price—more frequently than on the equilibrium path to reward it for the losses incurred during the stick.

⁹With two players, it actually suffices to condition only on the previous-period report profile. However, conditioning on the other players' contemporaneous reports is needed in general: Suppose there are three firms with costs drawn iid from $\{H, L\}$ with both realizations equally likely. Suppose firms 1 and 2 report H in periods $2t - 1$ and $2t$ for t odd, and they report L for t even. Then their joint reports form the sequence

$$(H, H), (H, H), (L, L), (L, L), (H, H), (H, H), (L, L), (L, L), \dots$$

If firm 3 reports truthfully, then conditional on any previous-period reports $(\theta_1, \theta_2, \theta_3)$, each player reports H with the correct frequency of .5, but the joint distribution fails to converge to the truth. (An analogous problem arises in a static model—see Jackson and Sonnenschein, 2007).

Formally, the punishment equilibria are obtained by bounding payoffs uniformly across the equilibria of a *punishment mechanism* which appends a minmax phase to the beginning of the credible reporting mechanism (Lemma 6.1). Checking incentives is then analogous to Fudenberg and Maskin (1986) (see Section 6.2).

3. THE MODEL

3.1. The Stage Game. The *stage game* is a finite game of incomplete information

$$u : A \times \Theta \rightarrow \mathbb{R}^n,$$

where $A = \prod_{i=1}^n A_i$ and $\Theta = \prod_{i=1}^n \Theta_i$ for some finite sets $A_i, \Theta_i, i = 1, \dots, n$. The interpretation is that each player $i = 1, \dots, n$ has a privately known *type* $\theta_i \in \Theta_i$ and chooses an *action* $a_i \in A_i$. As usual, u is extended to $\Delta(A \times \Theta)$ by expected utility.¹⁰ In the proofs we assume without loss that $u(A \times \Theta) \subset [0, 1]$.

The players are assumed to know their own payoffs, stated formally as follows.

Assumption 3.1 (Private Values). For all $a \in A, \theta \in \Theta, \theta' \in \Theta$, and $i = 1, \dots, n$,

$$\theta_i = \theta'_i \quad \Rightarrow \quad u_i(a, \theta) = u_i(a, \theta').$$

Given the assumption, we write $u_i(a, \theta_i)$ for $u_i(a, \theta)$.

3.2. The Dynamic Game. The *dynamic game* has the stage game u played with communication in each period $t = 1, 2, \dots$. The extensive form corresponds to a multi-stage game with observable actions and Markovian incomplete information where each period t is divided into the following substages:

- t.1 Each player i privately learns his type $\theta_i^t \in \Theta_i$.
- t.2 The players simultaneously send public messages $m_i^t \in \Theta_i$.
- t.3 A public randomization device generates $\omega^t \in [0, 1]$.
- t.4 The stage game u is played with actions $a_i^t \in A_i$ perfectly monitored.

The public randomizations are iid draws from the uniform distribution on $[0, 1]$ and independent of the players' types.¹¹

Player i 's type θ_i^t evolves according to an autonomous Markov chain (λ_i, P_i) on Θ_i , where λ_i is the initial distribution and P_i is the transition matrix. We make two assumptions about the players' type processes.

Assumption 3.2 (Independent Types). $(\lambda_i, P_i), i = 1, \dots, n$, are independent.¹²

¹⁰For any finite set X , we write $\Delta(X)$ for the set of all probability measures on X .

¹¹Since we allow for communication, there is a sense in which public randomizations are redundant: If the set of possible messages is large enough, the players can conduct jointly-controlled lotteries to generate such randomizations endogenously (see Aumann and Maschler, 1995).

¹²Independence of transitions is crucial for the argument. However, independence of initial distributions is only for convenience. Together they imply that the equilibrium beliefs are public, which simplifies the description of the equilibrium. All results hold verbatim for an arbitrary λ .

Let (λ, P) denote the joint type process on Θ .

Assumption 3.3 (Irreducible Types). P is irreducible.¹³

Irreducibility of P implies that the dynamic game is stationary, or repetitive, and there exists a unique invariant distribution denoted π . Independence across players implies that the invariant distribution takes the form $\pi = \pi_1 \times \cdots \times \pi_n$, where π_i is the invariant distribution for P_i .

Given stage-game payoffs $(v_i^t)_{t=1}^\infty$, player i 's dynamic game payoff is

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} v_i^t,$$

where the discount factor $\delta \in [0, 1[$ is common for all players.

3.3. Histories, Assessments, and Equilibria. A public history consists of past messages, realizations of the randomization device, and actions. The set of all public histories in period $t \geq 1$ is

$$H^t = (\Theta \times [0, 1] \times A)^{t-1} \cup ((\Theta \times [0, 1] \times A)^{t-1} \times \Theta \times [0, 1]).$$

The first set in the union consists of all public histories the players may face at stage $t.2$ when they are about to send messages; the second consists of all feasible public histories at $t.4$ when the players are about to choose actions. Let $H = \cup_{t \geq 1} H^t$.

A private history of player i in period t consists of the sequence of types the player has observed up to and including period t . The set of all such histories is denoted $H_i^t = \Theta_i^t$. Let $H_i = \cup_{t \geq 1} H_i^t$.

A (*behavioral*) *strategy for player i* is a sequence $\sigma_i = (\sigma_i^t)_{t \geq 1}$ of functions $\sigma_i^t: H^t \times H_i^t \rightarrow \Delta(A_i) \cup \Delta(\Theta_i)$ with $\sigma_i^t(\cdot \mid h^t, h_i^t) \in \Delta(\Theta_i)$ if $h^t \in (\Theta \times [0, 1] \times A)^{t-1}$, and $\sigma_i^t(\cdot \mid h^t, h_i^t) \in \Delta(A_i)$ if $h^t \in (\Theta \times [0, 1] \times A)^{t-1} \times \Theta \times [0, 1]$. A *strategy profile* $\sigma = (\sigma_i)_{i=1}^n$ and the type process (λ, P) induce a probability distribution over histories in the obvious way.

As types are independent across players and actions are observable, we assume that the players beliefs about private histories satisfy the standard restrictions imposed by perfect Bayesian equilibrium in multi-stage games with observable actions and incomplete information: (i) players $-i$ have common beliefs about player i , (ii) types are believed to be independent across players, and (iii) “players cannot signal what they don’t know” (see Fudenberg and Tirole, 1991).

Formally, public histories h^t and \hat{h}^t in H are *i -indistinguishable* if either

- (i) $h^t = (h^{t-1}, a)$ and $\hat{h}^t = (h^{t-1}, \hat{a})$ for $h^{t-1} \in H$, $a, \hat{a} \in A$ such that $a_i = \hat{a}_i$,

or

¹³Under Assumption 3.2, a sufficient (but not necessary) condition for P to be irreducible is that each P_i is irreducible and aperiodic (i.e., that each P_i is ergodic).

- (ii) $h^t = (\tilde{h}^t, m, \omega)$ and $\hat{h}^t = (\tilde{h}^t, \hat{m}, \hat{\omega})$ for $\tilde{h}^t \in H$, $(m, \omega), (\hat{m}, \hat{\omega}) \in \Theta \times [0, 1]$ such that $m_i = \hat{m}_i$.

The common (public) beliefs about player i are given by a sequence $\mu_i = (\mu_i^t)_{t \geq 1}$ of functions $\mu_i^t: H^t \rightarrow \Delta(H_i^t)$ such that $\mu_i^t(h^t) = \mu_i^t(\hat{h}^t)$ whenever h^t and \hat{h}^t are i -indistinguishable. A profile $\mu = (\mu_i)_{i=1}^n$ is called a *belief system*. Given a belief system μ , player j 's belief about the private histories of players $-j$ at a public history h^t is the product measure $\prod_{i \neq j} \mu_i^t(h^t) \in \prod_{i \neq j} \Delta(H_i^t)$.

An *assessment* is a pair (σ, μ) where σ is a strategy profile and μ is a belief system. Given an assessment (σ, μ) , let $u_i^\mu(\sigma | h^t, h_i^t)$ denote player i 's expected continuation payoff at history (h^t, h_i^t) , i.e., the expected discounted average payoff of player i from history (h^t, h_i^t) onwards (period t inclusive) when the expectation over i 's rivals' private histories is taken according to μ and play evolves according to σ . An assessment (σ, μ) is *sequentially rational* if for any player i , any history (h^t, h_i^t) and any strategy σ'_i , $u_i^\mu(\sigma | h^t, h_i^t) \geq u_i^\mu(\sigma'_i, \sigma_{-i} | h^t, h_i^t)$. An assessment (σ, μ) is a *perfect Bayesian equilibrium* (or a *PBE* for short) if it is sequentially rational and μ is computed using Bayes rule given σ wherever possible (both on and off the path of play).

3.4. Feasible Payoffs. Write $(f^t)_{t \geq 1}$ for a sequence of decision rules $f^t: \Theta^t \times [0, 1]^t \rightarrow A$ mapping histories of types and public randomizations into actions. The set of feasible payoffs in the dynamic game with discount factor δ is then

$$V(\delta) = \left\{ v \in \mathbb{R}^n \mid \exists (f^t)_{t \geq 1} \text{ s.t. } v = (1 - \delta) \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} u(f^t(\theta^1, \omega^1, \dots, \theta^t, \omega^t), \theta^t) \right] \right\}.$$

(Unless otherwise noted, all expectations are with respect to the joint distribution of the type process (λ, P) and the public randomizations.) Consider the set of payoffs attainable using a pure decision rule in a one-shot interaction in which types are drawn from the invariant distribution π , or

$$V^p = \left\{ v \in \mathbb{R}^n \mid \exists f: \Theta \rightarrow A \text{ s.t. } v = \mathbb{E}_\pi [u(f(\theta), \theta)] \right\}.$$

Let $V = \mathbf{co}(V^p)$ denote the convex hull of V^p .

Lemma 3.1 (Dutta, 1995). *As $\delta \rightarrow 1$, $V(\delta) \rightarrow V$ in the Hausdorff metric.*

Heuristically, the result follows from noting that in a stationary environment, (randomized) stationary decision rules are enough to generate all feasible payoffs, and for δ close to 1 the expected payoff from such a rule depends essentially only on the invariant distribution.

In what follows we focus on the limit feasible set V , keeping in mind that it is an arbitrarily good approximation to $V(\delta)$ when players are patient.

3.5. Minmax Values. Define player i 's *stationary (pure-action) minmax value* as

$$\underline{v}_i = \min_{a_{-i} \in A_{-i}} \mathbb{E}_{\pi_i} \left[\max_{a_i \in A_i} u_i((a_i, a_{-i}), \theta_i) \right],$$

and let $\underline{v} = (\underline{v}_1, \dots, \underline{v}_n)$. This can be interpreted as the pure-action minmax value in a one-shot game where types are distributed according to the invariant distribution π . Thus in the special case of iid types, \underline{v}_i is simply the standard pure-action minmax value. Since this is a novel concept and responsible for a limitation of our results, some remarks are in order.

Our motivation for the definition is pragmatic: \underline{v}_i is approximately the lowest payoff that can be imposed on a patient player i by having players $-i$ play a fixed pure-action profile for a large number of periods while player i best responds knowing his current type. This facilitates constructing stick-and-carrot punishments that generate payoffs close to \underline{v}_i during the stick phase. For example, in the Bertrand game studied by Athey and Bagwell (2008) (and considered in Section 2) this actually yields a tight lower bound on the set of individually-rational payoffs:

Example 3.1 (Bertrand competition/First-price auction). *Each player i is a firm setting a price $a_i \in A_i \subset \mathbb{R}_+$. There is a buyer who demands one unit of the good each period, has a reservation value $r \in]0, 1[$, and buys from the firm with the lowest price (randomizing uniformly in case of ties). Assume $\{0, r, 1\} \subset A_i$ for all i . Firm i 's privately-known marginal cost is $\theta_i \in \Theta_i \subset \mathbb{R}_+$, and thus its profit is $u_i(a, \theta_i) = (a_i - \theta_i) \mathbf{1}_{\{a_i = \min\{a_1, \dots, a_n, r\}\}} \frac{1}{|\{k | a_k = \min_j a_j\}|}$. Then $\underline{v}_i = 0$, since the profit is nonpositive at $a_{-i} = 0$, whereas firm i can guarantee a zero profit with $a_i = 1$.*

In general there may be equilibria generating payoffs strictly below \underline{v}_i .¹⁴ In particular, the stationary minmax value entails two possible limitations:

First, defining \underline{v}_i using pure actions is obviously restrictive in general even with iid types. Indeed, as complete information games are a special case of our model, we may take u to be the standard matching pennies game to see that in the worst case we may even have $\underline{v}_i = \max_{a, \theta_i} u_i(a, \theta_i)$ with the vector \underline{v} lying above the feasible set V . We have nothing to add to this well-known observation.¹⁵

¹⁴The proper minmax value for our dynamic game is given by a problem of the form

$$\min_{\sigma_{-i}} \max_{\sigma_i} \mathbb{E}_{\lambda_i, \sigma} \left[(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t, \theta_i^t) \right],$$

which can be interpreted as the value of a discounted zero-sum game with Markovian private information on one side (see Renault (2006) or Neyman (2008) for the undiscounted case). This definition is difficult to put to work as little is known about the optimal strategy of the uninformed player (here, players $-i$) when the type is irreducible. However, Hörner and Lovo (2009) show that it can be used in games with perfectly persistent types as the strategy of the uninformed player is then given by the approachability result of Blackwell (1956).

¹⁵Fudenberg and Maskin (1986) proved a mixed minmax folk theorem for complete information games by adjusting continuation payoffs in the carrot phase to make the punishers indifferent over all actions in the support of a mixed minmax profile. Extending the argument to our setting

	L	R
U	θ_1	0
D	0	$1 - \theta_1$

FIGURE 2. The payoff of player 1 for $\theta_1 \in \{0, 1\}$ in Example 3.2.

Second, and more pertinent to the current setting, players $-i$ should in general tailor the punishment to the information they learn about player i 's type during the punishment. This is illustrated by the next example.

Example 3.2 (Renault (2006)). *Consider a two-player game where $A_1 = \{U, D\}$, $A_2 = \{L, R\}$, $\Theta_1 = \{0, 1\}$, and $\Theta_2 = \{0\}$. The stage-game payoffs of player 1 are depicted in Figure 2. Player 1's type follows a Markov chain where $\theta_1^t = \theta_1^{t-1}$ with probability $p \in [\frac{1}{2}, 1[$. By symmetry, $\pi_1(\theta_1) = \frac{1}{2}$ for $\theta_1 \in \Theta_1$. The stationary minmax value is $\underline{v}_i = \frac{1}{2}$ (even if mixed strategies were allowed in the definition). Hörner, Rosenberg, Solan, and Vieille (2010) show that in the limit as $\delta \rightarrow 1$ the proper minmax value of player 1 is $v_p = \frac{p}{4p-1}$ for $p \in [\frac{1}{2}, \frac{2}{3}]$. Thus for patient players the two coincide in the iid case (i.e., $v_{\frac{1}{2}} = \underline{v}_1$), but differ whenever there is serial correlation (i.e., $v_p < \underline{v}_i$ for all $p \in]\frac{1}{2}, \frac{2}{3}]$). Heuristically, the reason is that player 1's myopic best response reveals his type. This is harmless in the iid case, but allows player 2 to tailor the punishment when types are correlated.*

The above limitations notwithstanding, there is an important class of dynamic games in which \underline{v}_i is in fact a tight lower bound on player i 's individually rational payoffs. These are games where the stage game u is such that for every player i and all $\theta_i \in \Theta_i$, the mixed minmax value $\min_{\alpha_{-i} \in \Delta(A_{-i})} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}, \theta_i)$ is achieved in pure strategies, and

$$\emptyset \neq \bigcap_{\theta_i \in \Theta_i} \arg \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a, \theta_i).$$

Then the intersection on the right contains a pure action profile $a_{-i} \in A_{-i}$ such that for all types of player i , a_{-i} is the harshest possible punishment. While this property is arguably special, it is true for many familiar games including the Bertrand game of Example 3.1 and all of the examples introduced in the next section.

3.6. Examples. The following examples illustrate the model and some of the definitions already introduced. They also provide instances of economic applications to which our main results (Theorem 4.1 and Corollary 4.1) apply.

is not straightforward as the variation in payoffs during the stick phase is private information and our methods characterize continuation values only “up to an ε .” Alternatively, Gossner (1995) uses a statistical test to assess whether the players are mixing with the right distributions. Extending his approach seems promising, but beyond the scope of this paper.

	C	D
C	$\beta - \theta_1, \beta - \theta_2$	$1 - \theta_1, 1$
D	$1, 1 - \theta_2$	$0, 0$

FIGURE 3. The payoff function for the public good game of Example 3.4.

Example 3.3 (Cournot competition). *Each player i is a firm that chooses a quantity $a_i \in A_i \subset \mathbb{R}_+$, where $0 \in A_i$. The market price is given by the positive and decreasing inverse demand $p(\sum_i a_i)$. Firm i 's cost function takes the form $c_i(a_i, \theta_i) \geq 0$, where $c_i(0, \theta_i) = 0$ and $c_i(a_i, \theta_i)$ is nondecreasing in a_i for all $\theta_i \in \Theta_i$. Thus its profit is $u_i(a, \theta_i) = p(\sum_j a_j)a_i - c_i(a_i, \theta_i)$. Assuming that for all i there exists $\bar{a}_{-i} \in A_{-i}$ such that $p(\sum_{j \neq i} \bar{a}_j) = 0$, we deduce that $\underline{v}_i = 0$ for all i since i 's rivals can flood the market by setting $a_{-i} = \bar{a}_{-i}$, whereas firm i can guarantee a zero profit by setting $a_i = 0$.*

Example 3.4 (Public good provision/Partnership). *In the binary-contribution game of Palfrey and Rosenthal (1988) (see also Fudenberg and Tirole, 1991), each of two players chooses whether to contribute (C) or not (D) to a joint project. Player i 's cost of contributing is $\theta_i \in \Theta_i \subset \mathbb{R}_+$. The project yields a benefit 1 or $\beta \geq 1$ depending on whether one or two players contributed. The payoff function is depicted in Figure 3. Then $\underline{v}_i = \mathbb{E}_{\pi_i}[\max\{0, 1 - \theta_i\}]$. Note that the best response of player i depends on his type.*

Example 3.5 (Informal risk sharing). *Consider an n -player version of the insurance problem of Wang (1995), which is an incomplete-information variant of the model by Kocherlakota (1996). Each player i is an agent with a random endowment $\theta_i \in \Theta_i \subset \mathbb{R}_+$. Agent i chooses transfers $a_i = (a_{i,1}, \dots, a_{i,n})$ in $A_i(\theta_i) \subset \left\{ a_i \in \mathbb{R}_+^n \mid \sum_j a_{i,j} \leq \theta_i \right\}$, $0 \in A_i(\theta_i)$.¹⁶ His utility of consumption is given by $u_i(a, \theta_i) = \bar{u}_i(\theta_i - \sum_j (a_{i,j} - a_{j,i}))$, where \bar{u}_i is nondecreasing and concave. Then $\underline{v}_i = \mathbb{E}_{\pi_i}[\bar{u}_i(\theta_i)]$, since agents $-i$ can opt to make no transfers to agent i , whereas agent i can guarantee this autarky payoff by consuming his endowment.*

4. THE MAIN RESULT

Let $\mathcal{P}(V)$ denote the (strong) Pareto frontier of the limit feasible set V , and let V^c denote the set of payoffs from constant decision rules, i.e.,

$$V^c = \left\{ v \in \mathbb{R}^n \mid \exists a \in A \text{ s.t. } v = \mathbb{E}_{\pi}[u(a, \theta)] \right\}.$$

Let $W = \mathbf{co}(\mathcal{P}(V) \cup V^c)$ denote the convex hull of $\mathcal{P}(V) \cup V^c$. Consider the set

$$W^* = \left\{ v \in W \mid v_i > \underline{v}_i, \ i = 1, \dots, n \right\},$$

¹⁶In this game the set of feasible actions $A_i(\theta_i)$ depends on the realized type θ_i . However, by letting $A_i = \cup_{\theta_i \in \Theta_i} A_i(\theta_i)$ and defining $u_i(a, \theta_i)$ to be some very large negative number if $a_i \notin A_i(\theta_i)$, we recover an essentially equivalent game where A_i is independent of θ_i .




FIGURE 4. Illustration of Theorem 4.1 for $n = 2$. (a) The shaded area is $\text{co}(V^c)$. The thick line segments are $\mathcal{P}(V)$. (b) The shaded area is W^* , which is contained in the set of limit PBE payoffs. w^1 and w^2 are used to construct player-specific punishments for v .

which consists of all vectors in W that are strictly above the stationary minmax profile \underline{v} . See Figure 4 for a schematic illustration for $n = 2$.

Definition 4.1. A vector $v \in \mathbb{R}^n$ allows player-specific punishments in W^* if there exists a collection of payoff profiles $\{w^i\}_{i=1}^n \subset W^*$ such that for all i ,

$$v_i > w_i^i,$$

and for all $j \neq i$,

$$w_i^j > w_i^i.$$

Note that while we are mainly interested in the Pareto frontier $\mathcal{P}(V)$, including the set V^c in the definition of W expands the set from which punishment profiles can be chosen.

The following is our main result.

Theorem 4.1. *Let $v \in W^*$ allow player-specific punishments in W^* . Then, for all $\varepsilon > 0$, there exists $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$, there exists a perfect Bayesian equilibrium where the expected continuation payoffs are within distance ε of v at all on-path histories.*

Theorem 4.1 shows that player-specific punishments are sufficient for a payoff vector v , which is a convex combination of payoffs to Pareto-efficient and constant decision rules and strictly dominates the stationary minmax value \underline{v} , to be virtually attainable in a perfect Bayesian equilibrium when players are patient. Furthermore, the equilibrium can be taken to be stationary in the sense of continuation payoffs remaining close to v at all on-path histories.

The statement of Theorem 4.1 leaves open the question about the nonemptiness of the set W^* and the existence of player-specific punishments. To address these, say that $v \in \mathbb{R}^n$ is a *limit equilibrium payoff* if it satisfies the conclusion of Theorem 4.1. Observe that every v in the interior of W^* allows player-specific punishments in W^* and hence is a limit equilibrium payoff (cf. Fudenberg and Maskin, 1986). The following limit-efficiency result then obtains by noting that W^* has nonempty interior if and only if W is full-dimensional and there exists $v \in V$ such that $v > \underline{v}$.

Corollary 4.1. *If W is full-dimensional, then every $v \in \mathcal{P}(V)$ such that $v > \underline{v}$ is a limit equilibrium payoff. If $n = 2$, then the full-dimensionality of W can be dispensed with.*

See Appendix A for the proof of the two-player result.

More generally, if W is full-dimensional, then every $v \in W^*$ is a limit equilibrium payoff.¹⁷ It is thus worth noting that the full-dimensionality of W can be naturally viewed as a generic property: Given the type process (λ, P) , it holds for an open set of full Lebesgue measure in the space of stage games with private values (i.e., in $\mathbb{R}^{|A| \sum_i |\Theta_i|}$) provided that each player has at least two actions.¹⁸ However, there still remains the possibility that W^* may be empty, since—as discussed in Section 3.5—the stationary minmax value may lie outside the feasible set V . While this is a robust problem in general, it does not appear to be particularly limiting for economic applications. For instance, W^* is nonempty and full-dimensional in all of the examples of Sections 3.5 and 3.6 (save for Example 3.2 where player 2’s payoffs were left unspecified).

5. CREDIBLE REPORTING MECHANISMS

In this section we consider auxiliary games where there is a mechanism which automatically selects a history-dependent action profile in each period so that the game reduces to one where the players just send messages.

A (direct) T -period mechanism, $T \in \mathbb{N} \cup \{\infty\}$, is a collection $(f^t)_{t=1}^T$ of decision rules $f^t : \Theta^t \times [0, 1]^t \rightarrow A$ mapping histories of messages and public randomizations into action profiles. Each mechanism induces a T -period reporting game that is obtained from the dynamic game defined in Subsection 3.2 by replacing stage $t.4$ with

t.4' The mechanism selects the action profile $f^t(m^1, \omega^1, \dots, m^t, \omega^t) \in A$,

¹⁷In fact, it suffices that the players’ expected utilities from efficient or constant decision rules satisfy the non-equivalent-utilities condition of Abreu, Dutta, and Smith (1994).

¹⁸To see this, note that $\mathbf{co}(V^c)$ (and hence W) is full-dimensional if for all vectors $v(a) = (\pi_i \cdot u_i(a, \cdot))_{i=1}^n$, with $a \in A$, the system $\sum_{a \in A} \beta_a v(a) = 0$ only has the trivial solution $\beta \equiv 0$. Since $\pi_i \in \mathbb{R}^n$ is a probability measure for all i , the Transversality theorem can be used to show that this holds for all $(u_i(a, \theta_i))_{i,a,\theta_i} \in \mathbb{R}^{|A| \sum_{i=1}^n |\Theta_i|}$ in an open set of full measure. The details are standard and hence omitted.

and truncating the game after period T . A strategy ρ_i for player i is simply the restriction of some dynamic game strategy σ_i to the appropriate histories. Write ρ_i^* for the truthful strategy.¹⁹ We abuse terminology by using ‘mechanism’ to refer both to a collection $(f^t)_{t=1}^T$ as well as to the game it induces.

In what follows we introduce the (class of) *credible reporting mechanism(s)*, or *CRM* for short. A CRM tests in every period whether each player’s past messages are sufficiently likely to have resulted from truthful reporting. If a player fails the test, then the mechanism ignores his messages from the next period onwards and substitutes appropriately generated random messages for them. The CRM maps current messages to actions according to some stationary decision rule.

Formally, given a sequence $(x^1, \dots, x^t) \in \Theta^t$, $t \geq 1$, let

$$\tau^t(\theta, \theta') = |\{2 \leq s \leq t \mid (x^{s-1}, x^s) = (\theta, \theta')\}|$$

and

$$\tau_i^t(\theta, \theta'_{-i}) = \sum_{\theta'_i} \tau^t(\theta, \theta')$$

for $(\theta, \theta') \in \Theta \times \Theta$ and $i = 1, \dots, n$. Define the empirical frequency

$$P_i^t(\theta'_i \mid \theta, \theta'_{-i}) = \frac{\tau^t(\theta, \theta')}{\tau_i^t(\theta, \theta'_{-i})},$$

where $\frac{0}{0} = 0$ by convention. A *test* is a sequence (b_k) such that $b_k \rightarrow 0$ (i.e., a null sequence). We say that player i *passes the test* (b_k) *at* (x^1, \dots, x^t) if

$$(5.1) \quad \sup_{\theta'_i} |P_i(\theta'_i \mid \theta_i) - P_i^t(\theta'_i \mid \theta, \theta'_{-i})| < b_{\tau_i^t(\theta, \theta'_{-i})} \quad \forall (\theta, \theta'_{-i}) \in \Theta \times \Theta_{-i}.$$

That is, i passes the test if and only if, for all (θ, θ'_{-i}) , the distribution of x_i^s over periods s such that $(x^{s-1}, x^s) = (\theta, \theta'_{-i})$ is within $b_{\tau_i^t(\theta, \theta'_{-i})}$ of player i ’s true conditional distribution $P_i(\cdot \mid \theta_i)$ in the sup-norm.²⁰ Note that if the sequence (x^1, \dots, x^t) is generated by the true type process (λ, P) , then player i ’s types over the said periods are in fact iid draws from $P_i(\cdot \mid \theta_i)$ by the Strong Markov property and Assumption 3.2. Since the left-hand side of (5.1) is the Kolmogorov-Smirnov statistic for testing the hypothesis that the sample $P_i^t(\cdot \mid \theta, \theta'_{-i})$ is generated by independent draws from $P_i(\cdot \mid \theta_i)$, this implies that a test (b_k) can be chosen such that under the true process (λ, P) player i passes with probability arbitrarily close to 1 even as $t \rightarrow \infty$.

¹⁹I.e., $\rho_i^{*t}(\theta_i^t \mid h^t, (h_i^{t-1}, \theta_i^t)) = 1$ for all t , $h^t \in \Theta^{t-1} \times [0, 1]^{t-1} \times A^{t-1}$, and $(h_i^{t-1}, \theta_i^t) \in H_i^t$.

²⁰Our test is related to, but different from, standard statistical methods in Markov chains (see, e.g., Amemiya, 1985, Chapter 11). Indeed, the transition count is the maximum likelihood estimator for a first-order Markov model. But our objective is unconventional as we want i to pass the test given arbitrary strategies (and hence arbitrary processes) for $-i$ as long as i ’s transitions converge to the truth and appear sufficiently independent from those of $-i$.

A CRM is a triple $(f, (b_k), T)$, where $f : \Theta \rightarrow A$ is a decision rule, (b_k) is a test, and $T < \infty$ denotes the time horizon. Let $\xi : \Theta \times [0, 1] \rightarrow \Theta$ be a function such that if ω is drawn from the uniform distribution on $[0, 1]$, then for all $\theta \in \Theta$, $\xi(\theta, \omega)$ is distributed on Θ according to $P(\cdot \mid \theta)$. The CRM $(f, (b_k), T)$ selects actions according to the following recursion:

For all $1 \leq t \leq T$ and every player i , put

$$x_i^t = \begin{cases} m_i^t & \text{if } i \text{ passes } (b_k) \text{ at } (x^1, \dots, x^s) \text{ for all } 1 \leq s < t, \\ \xi_i(x^{t-1}, \omega^t) & \text{otherwise,} \end{cases}$$

and let $a^t = f(x^t)$.

Remark 5.1. We note for future reference that the above recursion implicitly defines functions $\chi = (\chi^t)_{t=1}^T$, $\chi^t : \Theta^t \times [0, 1]^t \rightarrow \Theta$, such that for all i and t , $x_i^t = \chi_i^t(m^1, \dots, m^{t-1}, m_i^t, \omega^1, \dots, \omega^t)$. A CRM is thus a mechanism $(f^t)_{t=1}^T$ where $f^t = f \circ \chi^t$.

The next lemma shows that there exists a test lenient enough for a truthful player to be likely to pass it, yet stringent enough so that the empirical distribution of (x^1, \dots, x^T) converges to the invariant distribution π of the true type process as $T \rightarrow \infty$ regardless of the players' strategies.

Lemma 5.1. *Let $\varepsilon > 0$. There exists a test (b_k) satisfying the following conditions:*

- (1) *In every CRM $(f, (b_k), T)$, for all i , all ρ_{-i} , and all λ ,*

$$\mathbb{P}_{\rho_i^*, \rho_{-i}}[i \text{ passes } (b_k) \text{ at } (x^1, \dots, x^t) \text{ for all } t] \geq 1 - \varepsilon.$$

- (2) *$\exists \bar{T} < \infty$ such that in every CRM $(f, (b_k), T)$ with $T > \bar{T}$, for all ρ and all λ , the empirical distribution of (x^1, \dots, x^T) , denoted π^T , satisfies*

$$\mathbb{P}_\rho[\|\pi^T - \pi\| < \varepsilon] \geq 1 - \varepsilon.$$

When all players are truthful (i.e., if $\rho = \rho^*$), the existence of the test can be shown by using the convergence properties of Markov chains. The lemma extends the result uniformly to arbitrary strategies. The proof, presented in Appendix B, relies on the independence of transitions and our formulation of the test. In particular, for player i to pass the test requires that (i) the marginal distribution of his (reported) type transitions converges to the truth, and (ii) these transitions appear sufficiently independent of the other players' transitions (as the test is done separately for each (θ, θ'_{-i})). If player i reports truthfully, (i) is immediate and (ii) follows since, by independence, it is impossible for players $-i$ to systematically correlate their transitions with those of player i when the test conditions on the previous period type profile. This explains the first part of the lemma.

For the second part, suppose player i plays any strategy ρ_i . Then either he passes the test in all periods, in which case (i) and (ii) follow by definition of the

test, or he fails in some period t , after which the CRM generates x_i^s , $s > t$, by simulating the true type process for which (i) and (ii) follow by the argument in the first part of the lemma. An accounting argument then establishes the convergence of the long-run distribution of (x^1, \dots, x^T) .

We say that *player i can truthfully secure \bar{v}_i in the CRM $(f, (b_k), T)$* if

$$\min_{\rho_{-i}} \mathbb{E}_{(\rho_i^*, \rho_{-i})} \left[\frac{1 - \delta}{1 - \delta^T} \sum_{t=1}^T \delta^{t-1} u_i(f(x^t), \theta_i^t) \right] \geq \bar{v}_i,$$

where the expectation is with respect to the distribution induced by the strategy profile (ρ_i^*, ρ_{-i}) . That is, truthful reporting secures expected payoff \bar{v}_i to player i if, regardless of the reporting strategies of the other players, player i 's expected payoff from truthful reporting is at least \bar{v}_i .

For every $v \in V^p$, take $f^v : \Theta \rightarrow A$ such that $v = \mathbb{E}_\pi[u(f^v(\theta), \theta)]$. Our interest in the CRMs stems from the following security-payoff property:

Theorem 5.1. *Let $\varepsilon > 0$. There exist a test (b_k) and a time \bar{T} such that for all $T > \bar{T}$, there exists a discount factor $\bar{\delta} < 1$ such that for all $v \in V^p$, all $\delta > \bar{\delta}$, and all λ , every player i can truthfully secure $v_i - \varepsilon$ in the CRM $(f^v, (b_k), T)$.*

Since we have private values (Assumption 3.1), player i 's payoff depends only on the joint distribution of his own true types $(\theta_i^1, \dots, \theta_i^T)$ and the “edited reports” (x^1, \dots, x^T) . Thus the result follows essentially immediately from Lemma 5.1. See Appendix B for the details.

Remark 5.2. By picking v on the Pareto frontier of V^p , Theorem 5.1 implies the existence of T -period CRMs whose all Nash equilibria (and hence any refinements thereof) are approximately efficient if T is large and the players are sufficiently patient. The proof consists of bounding payoffs from below by Theorem 5.1 and from above by feasibility. (This also establishes that truthful reporting forms an ε -equilibrium of the mechanism.) See the proof of Corollary 5.1 for an analogous argument with an infinite horizon.

We now extend CRMs to an infinite horizon by constructing “block mechanisms” in which the players repeatedly play randomly chosen T -period CRMs. Let $(\phi, (b_k), T)$ denote a *random CRM*, i.e., a CRM where the decision rule f is determined once and for all by an initial public randomization $\phi \in \Delta(A^\Theta)$. For concreteness, we use the period-1 public randomization so that, with slight abuse of notation, $f = \phi(\cdot, \omega_1)$ for some $\phi : \Theta \times [0, 1] \rightarrow A$.²¹

A *block CRM* $(\phi, (b_k), T)^\infty$ is an infinite-horizon mechanism where the random CRM $(\phi, (b_k), T)$ is applied to each T -period block $(k - 1)T + 1, \dots, kT$, $k \in \mathbb{N}$.

²¹It is clear from the proof that the players not knowing which decision rule will be chosen before sending their period-1 messages does not effect the security-payoff result of Theorem 5.1.

Note that, by construction, the action profile selected by the mechanism depends only on the messages and public randomizations in the current block.

A result of Fudenberg and Levine (1983, Theorem 6.1) implies that a perfect Bayesian equilibrium exists in a block CRM.²² Together with the following corollary to Theorem 5.1 this shows that every payoff profile v which is a convex combination of payoffs to Pareto-efficient or constant decision rules can be approximated by a block CRM where in every perfect Bayesian equilibrium, the continuation-payoff profile is close to v at all histories.

For every $v \in V$, take $\phi^v \in \Delta(A^\Theta)$ such that $v = E_{\phi^v} E_\pi[u(f(\theta), \theta)]$.

Corollary 5.1. *Let $\varepsilon > 0$. There exist a test (b_k) , a time T , and a discount factor $\bar{\delta} < 1$ such that for all $v \in W$, all $\delta > \bar{\delta}$, and all λ , the expected continuation payoffs are within distance ε of v at all histories in all perfect Bayesian equilibria of the block CRM $(\phi^v, (b_k), T)^\infty$.*

The idea for the proof, presented in Appendix B, can be sketched as follows. Note first that if we replace W with V^c in the statement of the Corollary, then the claim is obvious since, by definition, any $v \in V^c$ can be generated with a constant decision rule under which actions are independent of the players' reports. On the other hand, if we replace W with $\mathcal{P}(V)$, we can apply the security-payoff result of Theorem 5.1 to each block to bound equilibrium payoffs from below, and then use feasibility to bound them from above. In particular, a lower bound follows by observing that at any history each player can revert to truthful reporting for the rest of the game, which guarantees the security payoff from all future blocks. With sufficiently little discounting, this is essentially all that matters for continuation payoffs. We then obtain the result for the entire set $W = \mathbf{co}(\mathcal{P}(V) \cup V^c)$ by randomizing over decision rules at the start of the block. While characterizing actual equilibrium behavior would be difficult, we note in passing that truthful reporting forms an ε -equilibrium.

6. CONSTRUCTING GAME EQUILIBRIA

In this section we finally construct game equilibria to deduce Theorem 4.1. To this end, let $v \in W^*$ allow player-specific punishments in W^* . Without loss of generality, take $\varepsilon > 0$ small enough so that the collection of punishment profiles $\{w^i\}_{i=1}^n \subset W^*$ satisfies, for all i ,

$$v_i > w_i^i + 2\varepsilon,$$

²²Fudenberg and Levine (1983) assume that all players including Nature have finitely many actions at each stage whereas we have a continuous public randomization in each period. However, the range of ξ is finite, and hence a finite randomization device suffices for any given CRM. Since block CRMs employ only finitely many CRMs, the same holds for them. An equilibrium of such a coarsened game remains an equilibrium in the game with a continuous randomization device.

and for all $j \neq i$,

$$w_i^j > w_i^i + 2\varepsilon.$$

(See Figure 4.) Assume further that $\varepsilon > 0$ is small enough so that there exists $\gamma \in]0, 1[$ such that for all $i \neq j$,

$$(6.1) \quad \gamma > \frac{2\varepsilon}{w_i^i - \underline{v}_i}$$

and

$$(6.2) \quad \gamma(\underline{v}_j + \varepsilon) + (1 - \gamma)(w_j^j - w_j^i + 2\varepsilon) < 0.$$

(These inequalities ensure that ε is small enough given the below choice of stick-phase length $L(\delta)$, which needs to be varied with δ due to the ‘‘approximation error’’ ε in payoffs.)

Corollary 5.1 yields block CRMs $(\phi^0, (b_k), T)^\infty, (\phi^1, (b_k), T)^\infty, \dots, (\phi^n, (b_k), T)^\infty$ and a discount factor $\delta_0 < 1$ such that for all $\delta > \delta_0$, the expected continuation payoffs are within distance ε of the corresponding target payoffs v, w^1, \dots, w^n at all histories in all perfect Bayesian equilibria of the block CRMs. Along the equilibrium path the players are simply assumed to mimic an equilibrium of the block CRM $(\phi^0, (b_k), T)$. However, as the players are now able to choose actions, they have to be enforced by means of suitable punishments. The other block CRMs are used to establish the existence of such punishments as follows.

6.1. Player-Specific Punishments. For each player i take a minmaxing profile

$$a_{-i}^i \in \arg \min_{a_{-i} \in A_{-i}} \mathbb{E}_{\pi_i} \left[\max_{a_i \in A_i} u_i(a, \theta_i) \right].$$

Given $L \in \mathbb{N}$ and the block CRM $(\phi^i, (b_k), T)^\infty$, we construct an auxiliary game $(L, (\phi^i, (b_k), T)^\infty)$ that runs over $t = 1, 2, \dots$ as follows. At each $t \in \{1, \dots, L\}$, the game proceeds exactly as the dynamic game defined in Section 3.2, except that every player $j \neq i$ is forced to take the action $a_j^t = a_j^i$. At $t = L + 1$, the block CRM $(\phi^i, (b_k), T)^\infty$ starts and runs over all subsequent periods. (Note that the construction of the block mechanism starting at $L + 1$ does not depend on how play transpires during the first L periods.) The evolution of types in the game $(L, (\phi^i, (b_k), T)^\infty)$ is identical to that in the dynamic game. However, it will be useful to have separate notation for the initial distribution of types, which will be denoted $B \in \Delta(\Theta)$ with $B(\theta) = \prod_{i=1}^n B_i(\theta_i)$.

We refer to the game $(L, (\phi^i, (b_k), T)^\infty)$ defined above as the *punishment mechanism against i* . It starts with players $j \neq i$ being restricted to minmax player i for L periods while player i can best respond with any actions $a_i^t \in A_i$, $t = 1, \dots, L$. The block CRM $(\phi^i, (b_k), T)^\infty$ then ensues. For reasons that will become clear in

the proof, we choose L as

$$L = L(\delta) = \max \left\{ n \in \mathbb{N} \mid n \leq \frac{\ln(1-\gamma)}{\ln(\delta)} \right\}.$$

Note that as $\delta \rightarrow 1$, we have $L(\delta) \rightarrow \infty$ and $\delta^{L(\delta)} \rightarrow 1 - \gamma$.

The following result provides bounds on the players' equilibrium payoffs in the punishment mechanism against i .

Lemma 6.1. *There exists $\delta_1 \geq \delta_0$ such that the following hold:*

- (1) *For all i , all $\delta > \delta_1$, all initial beliefs B , and all perfect Bayesian equilibria of the punishment mechanism $(L(\delta), (\phi^i, (b_k), T)^\infty)$, the expected continuation payoffs are within distance ε of w^i at all period- t histories for $t > L(\delta)$.*
- (2) *For all i , all $\delta > \delta_1$, and all $\theta_i \in \Theta_i$,*

$$\frac{1-\delta}{1-\delta^{L(\delta)}} \sum_{t=1}^{L(\delta)} \delta^{t-1} \mathbb{E} \left[\max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta_i^t) \mid \theta_i^1 = \theta_i \right] \leq \underline{v}_i + \varepsilon.$$

The proof is presented in Appendix C. The first part follows immediately from Corollary 5.1, since in periods $t > L(\delta)$ the players are playing the block CRM $(\phi^i, (b_k), T)^\infty$. To interpret the second part, observe that the left-hand side is an upper bound on player i 's discounted average payoff from the periods in which he is being minmaxed (i.e., from the first $L(\delta)$ periods of $(L(\delta), (\phi^i, (b_k), T)^\infty)$) given initial type θ_i .

The existence of a perfect Bayesian equilibrium in the punishment mechanisms follows by Fudenberg and Levine (1983, Theorem 6.1).

6.2. Phases and Equilibrium Strategies. An equilibrium yielding continuation payoffs within ε of v can now be informally described as follows (see Appendix D for a formal description of the equilibrium strategies and beliefs). The equilibrium starts in the *cooperative phase* where play mimics an equilibrium of the block CRM $(\phi^0, (b_k), T)^\infty$. That is, the players send messages according to the equilibrium of the mechanism and play the action profile the mechanism would have chosen given the history of messages and public randomizations. As long as there has never been a period where some player deviated from the action prescribed by the mechanism, play remains in the cooperative phase. This is the case everywhere on the equilibrium path.

A deviation by player i from the prescribed action in the cooperative phase triggers the *punishment phase against i* where play mimics an equilibrium of the punishment mechanism $(L, (\phi^i, (b_k), T)^\infty)$. This consists of the *stick subphase*—in which the deviator i is minmaxed for L periods—followed by the *carrot subphase*, which builds on the block CRM $(\phi^i, (b_k), T)^\infty$. In the stick subphase, all players send messages and player i chooses actions as in the equilibrium of the punishment mechanism, and players $-i$ play the minmax profile prescribed by the mechanism.

In the carrot subphase, the players continue to send messages according to the equilibrium, and all players play the action profile prescribed by the mechanism. As long as there has never been a period where a player deviated from the action prescribed by the punishment mechanism (either some player $k \neq i$ during the stick, or any player during the carrot), play remains in the punishment phase against i . A deviation by player j from the prescribed action triggers the punishment phase against j .

Within each phase beliefs evolve as in the equilibrium being mimicked. If some player i triggers a change of phase by deviating from the action prescribed by the mechanism, then the initial beliefs for the punishment mechanism against i are determined by the current public beliefs about the players' types in the period where the deviation occurred.²³

Remark 6.1. It is worth emphasizing that a change in phase is triggered only when a player deviates from the action prescribed by the mechanism. A player may also deviate from the equilibrium of the mechanism by sending a different message, and this deviation may be observable (e.g., suppose that the equilibrium of the mechanism has the player reporting truthfully, and type transitions do not have full support). Similarly, the punishment mechanism against i does not prescribe actions for player i when he is being minmaxed so that he may have an observable deviation there. However, both of these deviations result in a history which is feasible in the mechanism, and therefore the equilibrium that is being mimicked prescribes some continuation strategies and beliefs following the deviation. Accordingly, we assume that the players simply continue to mimic the equilibrium.

As the strategies described above have the players mimic an equilibrium of the block CRM $(\phi^0, (b_k), T)^\infty$ on the path of play, they result in continuation payoffs that are within distance ε of v after all on-path histories by Corollary 5.1. Thus to complete the proof of Theorem 4.1, it suffices to show that the strategies are sequentially rational. To this end, note that within each phase play corresponds to an equilibrium of some mechanism, and hence deviations that do not lead to a change of phase are unprofitable a priori. Therefore, it is enough to verify that at every history, no player gains by triggering a change of phase by deviating in action.²⁴ We do this by showing that there exists $\bar{\delta} < 1$ such that, regardless of the history, triggering a change of phase cannot be optimal for a player when $\delta > \bar{\delta}$.

²³More precisely, if the punishment phase against i starts in period t , then the initial beliefs in the punishment mechanism against i are given by the current public beliefs over the entire private histories Θ^t . However, the continuation equilibrium of the punishment mechanism depends only on the beliefs about the period- t type profile.

²⁴Note that such deviations include “double deviations,” where a player first deviates within a phase, and only then deviates in a way that triggers the punishment.

Cooperative-phase histories: Suppose play is in the cooperative phase. Then player j 's expected continuation payoff is at least $v_j - \varepsilon$. A one-stage deviation in action triggers the punishment phase against j and, by Lemma 6.1, for $\delta \geq \delta_1$ yields at most

$$(1 - \delta) + (\delta - \delta^{L(\delta)+1})(v_j + \varepsilon) + \delta^{L(\delta)+1}(w_j^j + \varepsilon) \leq (1 - \delta) + \delta(w_j^j + \varepsilon).$$

At $\delta = 1$, the right-hand side is strictly less than $v_j - \varepsilon$. Therefore, we can find $\delta_2 \geq \delta_1$ such that for all $\delta > \delta_2$ the deviation is unprofitable.

Stick-subphase histories: Consider a history at which player i should be min-maxed. It is enough to show that, for $j \neq i$, it is in player j 's interest to choose a_j^i (see Remark 6.1). By conforming, j 's payoff is at least $\delta^{L(\delta)}(w_j^i - \varepsilon)$, whereas a one-stage deviation in action triggers the punishment phase against j . By Lemma 6.1, for $\delta \geq \delta_1$ the incentive constraint takes the form

$$(1 - \delta) + (\delta - \delta^{L(\delta)+1})(v_j + \varepsilon) + \delta^{L(\delta)+1}(w_j^j + \varepsilon) \leq \delta^{L(\delta)}(w_j^i - \varepsilon).$$

As $\delta \rightarrow 1$, the inequality becomes $\gamma(v_j + \varepsilon) + (1 - \gamma)(w_j^j - w_j^i + 2\varepsilon) \leq 0$, which holds strictly by (6.2). Thus, there exists $\delta_3 \geq \delta_2$ such that for all $\delta > \delta_3$, the deviation is unprofitable.

Carrot-subphase histories: Suppose play is in the carrot subphase of the punishment phase against i . Then player j 's continuation payoff is at least $w_j^i - \varepsilon$. A one-stage deviation in action triggers the punishment phase against j and we use Lemma 6.1 to write the incentive constraint preventing the deviation for $\delta \geq \delta_1$ as

$$(1 - \delta) + (\delta - \delta^{L+1})(v_j + \varepsilon) + \delta^{L+1}(w_j^j + \varepsilon) \leq w_j^i - \varepsilon.$$

As $\delta \rightarrow 1$, the inequality becomes $\gamma(v_j + \varepsilon) + (1 - \gamma)(w_j^j + \varepsilon) \leq w_j^i - \varepsilon$. It is enough to check this inequality when $j = i$ since $w_j^i > w_j^j$ for $j \neq i$. But by our choice of γ and ε , (6.1) holds for all j and hence the limit incentive constraint holds with strict inequality. We conclude that there exists $\delta_4 \geq \delta_3$ such that for all $\delta > \delta_4$, the deviation is unprofitable.

Put $\bar{\delta} = \delta_4$. Then all one-stage deviations that trigger a change of phase are unprofitable for all $\delta > \bar{\delta}$, and hence the strategies are sequentially rational. Theorem 4.1 follows.

7. CONCLUDING REMARKS

Our main result (Theorem 4.1) shows that repeated interaction may allow individuals to overcome the problems of self-interested behavior and asymmetric information even if private information is persistent. The proof suggests that, given enough patience, cooperation can be supported with behavior that amounts to a form of mental accounting: If the history of a player's reports about his private state appears credible when evaluated against the reports of the other, then that

player's reports are taken at face value when deciding on actions. If a player loses credibility, which may happen with some probability, then his reports no longer matter for the choice of actions. However, as long as the player complies with the chosen actions, this is his only punishment, and he will regain credibility after a while. It is only in the event of a deviation in actions that a harsher punishment is used, but this is off the equilibrium path.

7.1. On the Limit Equilibrium Payoff Set. Let $E(\delta) \subset \mathbb{R}^n$ be the set of PBE payoffs in our dynamic Bayesian game. Theorem 4.1 provides an inner bound for the set $E(\delta)$ in the limit as $\delta \rightarrow 1$, but a sharper characterization of this set is an important open question. In some special cases, however, our results provide a tight estimate of $E(\delta)$ as players become arbitrarily patient.

Observe first that in the special case of a game of complete information, $V = V^c$. Thus under the full-dimensionality assumption, we recover an approximate pure-minmax version of the folk theorem by Fudenberg and Maskin (1986).

More interestingly, in a Bertrand game with privately-known costs our inner bound is actually tight:

Proposition 7.1. *Consider the Bertrand game of Example 3.1 and assume that $\max\{\theta_i \mid i = 1, \dots, n, \theta_i \in \Theta_i\} < r$. Then,*

$$\lim_{\delta \rightarrow 1} E(\delta) = V \cap \mathbb{R}_+^n = W \cap \mathbb{R}_+^n.$$

This result yields a folk theorem for dynamic Bayesian Bertrand games as feasibility and individual rationality imply $E(\delta) \subset V(\delta) \cap \mathbb{R}_+^n$ for all δ . From Proposition 7.1 and Lemma 3.1 we then deduce that feasibility and individual rationality are the only restrictions that can be imposed on equilibrium payoffs as $\delta \rightarrow 1$.

7.2. Robustness. There are various directions in which the robustness of our results could be explored:

We assume that the type process is autonomous. Extending the results to decision-controlled processes studied in the literature on stochastic games (see, e.g., Dutta, 1995) appears feasible but notationally involved.

The main restrictions we impose on the players' information are private values (Assumption 3.1), independence across players (Assumption 3.2), and irreducibility (Assumption 3.3). As mentioned in the Introduction, when types are perfectly persistent, the limit equilibrium payoff set may be bounded away from efficiency. Similarly, with interdependent values, efficiency need not be achievable even with transfers (see Jehiel and Moldovanu, 2001). Thus private values and irreducibility are necessary for general efficiency results. In contrast, correlation of types typically expands the set of implementable outcomes in a mechanism design setting

(see Cremer and McLean, 1988). Therefore we conjecture that our results extend to correlated transitions even if our proof does not.²⁵

Finally, the assumption about communication in every period cannot in general be dispensed with without affecting the set of achievable payoffs. To see this, it suffices to consider the Cournot game of Example 3.3 in the special case of iid types. The collusive scheme where the firm with the lowest cost always produces the monopoly output is infeasible without communication. In contrast, by Corollary 4.1 the resulting profits are limit equilibrium payoffs when cheap talk is allowed.

APPENDIX A. A PROOF FOR SECTION 4

Proof of Corollary 4.1. If W is full-dimensional, the result follows by the argument given in the text. Thus it suffices to take $n = 2$ and $0 \leq \dim W \leq 1$ (since $W \neq \emptyset$). Observe that if $\dim \mathcal{P}(V) = 1$, then every $v \in \mathcal{P}(V)$, $v > \underline{v}$, allows player-specific punishments in W^* (they can be chosen in $\mathcal{P}(V)$). So it remains to consider the case where $\mathcal{P}(V)$ is a singleton. Then there exists a decision rule $f: \Theta \rightarrow A$ such that for all $\theta \in \Theta$ and all i ,

$$f(\theta) \in \arg \max \{u_i(a, \theta_i) \mid a \in A\}.$$

Consider the following pure strategy for player i . For all $h^t \in H^t$ and $h_i^t \in \Theta_i^t$, if $h^t = (m^1, \omega^1, a^1, \dots, m^{t-1}, \omega^{t-1}, a^{t-1})$, then

$$\sigma_i^t(h^t, h_i^t) = \theta_i^t,$$

whereas if $h^t = (m^1, \omega^1, a^1, \dots, m^{t-1}, \omega^{t-1}, a^{t-1}, m^t, \omega^t)$, then

$$\sigma_i^t(h^t, h_i^t) \begin{cases} = f_i(m^t) & \text{if } m_i^t = \theta_i^t, \\ \in \arg \max \{u_i(a_i, f_{-i}(m^t), \theta_i^t) \mid a_i \in A_i\} & \text{otherwise.} \end{cases}$$

At any public history, the belief μ_i about player i puts probability 1 to the private history that coincides with the history of i 's messages. To verify that strategies are sequentially rational, observe that for any $t \geq 1$, continuation play from period $t+1$ onward does not depend on the outcome in period t . Thus, a player cannot gain by deviating in messages as truthful messages maximize his current payoff. Following (truthful or non-truthful) reports in period t , a player cannot gain by choosing a different action either as the equilibrium action maximizes his current payoff given his type. The corollary now follows by noting that the expected equilibrium payoff converges to the efficient profile $\mathbb{E}_\pi[u(f(\theta), \theta)]$ as $\delta \rightarrow 1$ (uniformly in λ). \square

²⁵Our proof does extend to the case where the correlation is due to a public Markov state but transitions are independent conditional on the state.

APPENDIX B. PROOFS FOR SECTION 5

B.1. Preliminaries. We state and prove two convergence results that are used in the proof of Lemma 5.1. The first is a corollary of Massart's (1990) result about the rate of convergence in the Glivenko-Cantelli theorem. Throughout $\|\cdot\|$ denotes the sup-norm.

Lemma B.1. *Let Θ be a finite set, and let $g \in \Delta(\Theta)$. Given an infinite sequence of independent draws from g , let g^k denote the empirical measure obtained by observing the first k draws. (I.e., for all $k \in \mathbb{N}$ and all $\theta \in \Theta$, $g^k(\theta) = \frac{1}{k} \sum_{t=1}^k 1_{\{\theta^t = \theta\}}$.) For all $\alpha > 0$, there exists a null sequence $(b_k)_{k \in \mathbb{N}}$ such that*

$$\mathbb{P}[\forall k \in \mathbb{N} \|g^k - g\| \leq b_k] \geq 1 - \alpha.$$

Proof. Fix $\alpha > 0$ and define the sequence $(b_k)_{k \in \mathbb{N}}$ by

$$b_k = \sqrt{\frac{2}{k} \log \frac{\pi^2 k^2}{3\alpha}}.$$

(In this proof only, π denotes the ratio of a circle's circumference to its diameter, not the invariant distribution.) Clearly $b_k \rightarrow 0$ so that $(b_k)_{k \in \mathbb{N}}$ is a null sequence. Without loss, label the elements of Θ from 1 to $|\Theta|$. Define the cdf G from g by setting $G(l) = \sum_{j=1}^l g(j)$. The empirical cdf's G^k are defined analogously from g^k . For all k , all l ,

$$|g^k(l) - g(l)| \leq |G^k(l) - G(l)| + |G^k(l-1) - G(l-1)|,$$

so that $\|g^k - g\| \leq 2\|G^k - G\|$. Defining the events $B_k = \{\|g^k - g\| \leq b_k\}$ we then have $\{\|G^k - G\| \leq \frac{b_k}{2}\} \subset B_k$. Thus,

$$\mathbb{P}[B_k] \geq \mathbb{P}[\|G^k - G\| \leq \frac{b_k}{2}] \geq 1 - 2e^{-2k(\frac{b_k}{2})^2} = 1 - \frac{6\alpha}{\pi^2 k^2},$$

where the second inequality is by Massart (1990) and the equality is by definition of b_k . The lemma now follows by observing that

$$\mathbb{P}[\bigcap_{k \in \mathbb{N}} B_k] = 1 - \mathbb{P}[\bigcup_{k \in \mathbb{N}} B_k^C] \geq 1 - \sum_{k \in \mathbb{N}} \mathbb{P}[B_k^C] \geq 1 - \sum_{k \in \mathbb{N}} \frac{6\alpha}{\pi^2 k^2} = 1 - \alpha,$$

where the last equality follows since $\sum_{k=1}^{\infty} \frac{1}{k^2} = \frac{\pi^2}{6}$. \square

Lemma B.2. *Let P be an irreducible stochastic matrix on a finite set Θ , and let π denote the unique invariant distribution for P . Let $(\theta^t)_{t \in \mathbb{N}}$ be a sequence in Θ . For all t , define the empirical matrix P^t by setting*

$$P^t(\theta' | \theta) = \frac{|\{s \in \{1, \dots, t-1\} : (\theta^s, \theta^{s+1}) = (\theta, \theta')\}|}{|\{s \in \{1, \dots, t-1\} : \theta^s = \theta\}|},$$

and define the empirical distribution π^t by setting

$$\pi_\theta^t = \frac{|\{s \in \{1, \dots, t\} : \theta^s = \theta\}|}{t}.$$

For all $\varepsilon > 0$ there exists $T < \infty$ and $\eta > 0$ such that for all $t \geq T$,

$$\|P^t - P\| < \eta \quad \Rightarrow \quad \|\pi^t - \pi\| < \varepsilon.$$

P^t is an empirical transition matrix that records for each state θ the empirical conditional frequencies of transitions $\theta \rightarrow \theta'$ in $(\theta^s)_{s=1}^t$. Similarly, π^t is an empirical measure that records the frequencies of states in $(\theta^s)_{s=1}^t$. So the lemma states roughly that if the conditional transition frequencies converge to those in P , then the empirical distribution converges to the invariant distribution for P .

Proof. Fix $\theta' \in \Theta$ and $t \in \mathbb{N}$. Note that $t\pi_{\theta'}^t$ is the number of visits to θ' in $(\theta^s)_{s=1}^t$. Since each visit to θ' is either in period 1 or preceded by some state θ , we have

$$t\pi_{\theta'}^t \leq 1 + \sum_{\theta \in \Theta} |\{s < t : \theta^s = \theta\}| P^t(\theta' | \theta) \leq |\Theta| + \sum_{\theta \in \Theta} t\pi_\theta^t P^t(\theta' | \theta).$$

On the other hand,

$$t\pi_{\theta'}^t \geq \sum_{\theta \in \Theta} |\{s < t : \theta^s = \theta\}| P^t(\theta' | \theta) \geq \sum_{\theta \in \Theta} t\pi_\theta^t P^t(\theta' | \theta) - |\Theta|,$$

where the second inequality follows, since $|\{s < t : \theta^s = \theta\}| \geq t\pi_\theta^t - 1$ and $\sum_{\theta} P^t(\theta' | \theta) \leq |\Theta|$. Putting together the above inequalities gives

$$-\frac{|\Theta|}{t} \leq \pi_{\theta'}^t - \sum_{\theta \in \Theta} \pi_\theta^t P^t(\theta' | \theta) \leq \frac{|\Theta|}{t}.$$

Since θ' was arbitrary, we have in vector notation

$$-\frac{|\Theta|}{t} \mathbf{1} \leq \pi^t(I - P^t) \leq \frac{|\Theta|}{t} \mathbf{1},$$

where I is the identity matrix and $\mathbf{1}$ denotes a $|\Theta|$ -vector of ones. This implies that for all t , there exists $e^t \in \mathbb{R}^{|\Theta|}$ such that $\|e^t\| \leq \frac{|\Theta|}{t}$ and $\pi^t(I - P^t) = e^t$. Let E be a $|\Theta| \times |\Theta|$ -matrix of ones. Then

$$\pi^t(I - P^t + E) = \mathbf{1} + e^t \quad \text{and} \quad \pi(I - P + E) = \mathbf{1}.$$

It is straightforward to verify that the matrix $I - P + E$ is invertible when P is irreducible (see, e.g., Norris, 1997, Exercise 1.7.5). The set of invertible matrices is open, so there exists $\eta_1 > 0$ such that $I - P^t + E$ is invertible if $\|P^t - P\| < \eta_1$. Furthermore, the mapping $Q \mapsto (I - Q + E)^{-1}$ is continuous at P , so there exists $\eta_2 > 0$ such that $\|(I - P^t + E)^{-1} - (I - P + E)^{-1}\| < \frac{\varepsilon}{4|\Theta|}$ if $\|P^t - P\| < \eta_2$. Put $\eta = \min\{\eta_1, \eta_2\}$ and put

$$T = \frac{2|\Theta|^2 \|(I - P + E)^{-1}\|}{\varepsilon}.$$

If $t \geq T$ and $\|P^t - P\| < \eta$, then

$$\begin{aligned} \|\pi^t - \pi\| &= \|(\mathbf{1} + e^t)(I - P^t + E)^{-1} - \mathbf{1}(I - P + E)^{-1}\| \\ &\leq \|(\mathbf{1} + e^t)[(I - P^t + E)^{-1} - (I - P + E)^{-1}]\| + \|e^t(I - P + E)^{-1}\| \\ &\leq 2|\Theta| \|(I - P^t + E)^{-1} - (I - P + E)^{-1}\| + \frac{|\Theta|^2}{t} \|(I - P + E)^{-1}\| \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2}. \end{aligned}$$

The lemma follows. \square

B.2. Proof of Lemma 5.1. Fix $\varepsilon > 0$ once and for all. By Lemma B.1, there exists a test (b_k) such that for all probability measures g with finite support, $\mathbb{P}[\forall k \in \mathbb{N} \|g^k - g\| \leq b_k] \geq 1 - \frac{\varepsilon}{|\Theta|^2 n}$, where g^k denotes the empirical measure of the first k observations of an infinite sequence of iid draws from g . Fix some such test (b_k) . We claim that it satisfies conditions (1) and (2) of Lemma 5.1.

Consider first condition (1). Fix T , i , λ , and ρ_{-i} . It is without loss to assume that λ is degenerate and ρ_{-i} is a pure strategy profile as the general case then follows by taking expectations. Similarly, it obviously suffices to show the claim conditional on an arbitrary realization $(\theta_{-i}^t, \omega^t)_{t=1}^T$. To this end, note that by Assumption 3.2, for any degenerate λ , the Markov chain (λ_i, P_i) is a version of the conditional distribution of $(\theta_i^t)_{t=1}^T$ given $(\theta_{-i}^t, \omega^t)_{t=1}^T$. Furthermore, for a fixed realization $(\theta_{-i}^t, \omega^t)_{t=1}^T$, the vectors x_{-i}^t ($t = 1, \dots, T$) are generated as deterministic functions of i 's truthful messages according to some ‘‘pure strategy’’ $r = (r^t)_{t=1}^T$, $r^t : \Theta_i^{t-1} \rightarrow \Theta_{-i}$, induced by ρ_{-i} and χ , where χ is the mapping defined in Remark 5.1. Thus it suffices to establish the following:

Lemma B.3. *For all $r = (r^t)_{t=1}^T$, $r^t : \Theta_i^{t-1} \rightarrow \Theta_{-i}$, if θ_i^t follows (λ_i, P_i) and $x_{-i}^t = r^t(\theta_i^1, \dots, \theta_i^{t-1})$, then*

$$\mathbb{P}[i \text{ passes } (b_k) \text{ at } ((\theta_i^1, x_{-i}^1), \dots, (\theta_i^t, x_{-i}^t)) \text{ for all } t] \geq 1 - \frac{\varepsilon}{n}.$$

Proof. We start by introducing a collection of auxiliary random variables, which are used to generate player i 's types. (The construction that follows is inspired by Billingsley, 1961.) Let $([0, 1], \mathcal{B}, \hat{\mathbb{P}})$ be a probability space. On this space define a countably infinite collection of *independent* random variables

$$\tilde{\psi}_{\theta, \theta'_{-i}}^s : [0, 1] \rightarrow \Theta_i, \quad (\theta, \theta'_{-i}) \in \Theta \times \Theta_{-i}, \quad s \in \mathbb{N},$$

where

$$\hat{\mathbb{P}}[\tilde{\psi}_{(\theta_i, \theta_{-i}), \theta'_{-i}}^s = \theta'_i] = P_i(\theta'_i | \theta_i).$$

That is, for any fixed $\theta = (\theta_i, \theta_{-i})$ and θ'_{-i} , the variables $\tilde{\psi}_{\theta, \theta'_{-i}}^s$, $s = 1, 2, \dots$, form a sequence of iid draws from $P_i(\cdot | \theta_i)$.

Given any $r = (r^t)_{t=1}^T$, we generate the path $(\theta_i^t, x_{-i}^t)_{t=1}^T$ of player i 's types (which equal his messages) and the vectors $x_{-i}^t \in \Theta_{-i}$ recursively as follows:

(θ_i^1, x_{-i}^1) is a constant given by the degenerate initial distribution λ_i and r^1 . For $1 < t \leq T$, suppose $(\theta_i^\tau, x_{-i}^\tau)_{\tau=1}^{t-1}$ have been generated. Then let

$$x_{-i}^t = r^t(\theta_i^1, \dots, \theta_i^{t-1}) \quad \text{and} \quad \theta_i^t = \psi_{((\theta_i^{t-1}, x_{-i}^{t-1}), x_{-i}^t)}^\tau,$$

where $\tau = |\{2 \leq s \leq t \mid ((\theta_i^{s-1}, x_{-i}^{s-1}), x_{-i}^s) = ((\theta_i^{t-1}, x_{-i}^{t-1}), x_{-i}^t)\}|$. That is, r^t determines x_{-i}^t , and then θ_i^t is found by sampling the first unsampled element in the sequence $(\psi_{((\theta_i^{t-1}, x_{-i}^{t-1}), x_{-i}^t)}^s)_{s \in \mathbb{N}}$.²⁶

Denote by $E_{\theta, \theta'_{-i}} \in \mathcal{B}$ the event where, for all $k \in \mathbb{N}$, the empirical distribution of the first k variables in the sequence $(\tilde{\psi}_{(\theta, \theta'_{-i})}^s)_{s \in \mathbb{N}}$ is within b_k of the true distribution $P_i(\cdot \mid \theta_i)$ in the sup-norm. Let $E = \bigcap_{\theta \in \Theta} \bigcap_{\theta'_{-i} \in \Theta_{-i}} E_{\theta, \theta'_{-i}}$. By definition of (b_k) , $\hat{\mathbb{P}}(E_{\theta, \theta'_{-i}}) \geq 1 - \frac{\varepsilon}{|\Theta|^2 n}$ for all θ, θ'_{-i} , and hence $\hat{\mathbb{P}}(E) \geq 1 - \frac{\varepsilon}{n}$.

To complete the proof, note that conditional on E , player i passes the test (b_k) at $((\theta_i^1, x_{-i}^1), \dots, (\theta_i^t, x_{-i}^t))$ for all $t = 1, \dots, T$. Indeed, by construction, for all t and all $(\theta, \theta'_{-i}) \in \Theta \times \Theta_{-i}$, $P_i(\cdot \mid \theta, \theta'_{-i})$ is the empirical distribution of the first k variables in the sequence $(\tilde{\psi}_{(\theta, \theta'_{-i})}^s)_{s \in \mathbb{N}}$ for some $k \in \mathbb{N}$. But conditional on E , this distribution is within b_k of $P_i(\cdot \mid \theta_i)$ by definition of E . \square

Having established condition (1) of Lemma 5.1, we now turn to condition (2). By Lemma B.2, it suffices to show that the test (b_k) satisfies the following:

- (2') For all $\eta > 0$, there exists $\bar{T} < \infty$ such that in every CRM $(f, (b_k), T)$ with $T > \bar{T}$, for all ρ and all λ ,

$$\mathbb{P}_\rho[\|P^T - P\| < \eta] \geq 1 - \varepsilon,$$

where P^T is the empirical matrix defined for each $(x^1, \dots, x^T) \in \Theta^T$ by

$$P^T(\theta' \mid \theta) = \frac{|\{s \in \{1, \dots, T-1\} : (x^s, x^{s+1}) = (\theta, \theta')\}|}{|\{s \in \{1, \dots, T-1\} : x^s = \theta\}|}.$$

It is useful to define the sequence (c_k) from (b_k) by

$$c_k = 2 \max_{1 \leq j \leq k} \frac{j}{k} b_j + \frac{1}{k}.$$

²⁶To see that this generates the right process, fix a path $(\theta_i^t, x_{-i}^t)_{t=1}^T$. It has positive probability only if $x_{-i}^t = r^t(\theta_i^1, \dots, \theta_i^{t-1})$ for all t , in which case its probability under (λ_i, P_i) is simply

$$\lambda_i(\theta_i^1) P_i(\theta_i^2 \mid \theta_i^1) \cdots P_i(\theta_i^T \mid \theta_i^{T-1}).$$

On the other hand, our auxiliary construction assigns it probability

$$\lambda_i(\theta_i^1) \hat{\mathbb{P}}[\tilde{\psi}_{(\theta_i^1, x_{-i}^1), x_{-i}^2}^1 = \theta_i^2] \cdots \hat{\mathbb{P}}[\tilde{\psi}_{(\theta_i^{T-1}, x_{-i}^{T-1}), x_{-i}^T}^\tau = \theta_i^T],$$

where $\tau = |\{2 \leq s \leq T \mid ((\theta_i^{s-1}, x_{-i}^{s-1}), x_{-i}^s) = ((\theta_i^{T-1}, x_{-i}^{T-1}), x_{-i}^T)\}|$, and where we have used independence of the $\tilde{\psi}_{\theta, \theta'_{-i}}^s$ to write the joint probability as a product. But by construction,

$$\hat{\mathbb{P}}[\tilde{\psi}_{(\theta_i^1, x_{-i}^1), x_{-i}^2}^1 = \theta_i^2] = P_i(\theta_i^2 \mid \theta_i^1) \quad \text{and} \quad \hat{\mathbb{P}}[\tilde{\psi}_{(\theta_i^{T-1}, x_{-i}^{T-1}), x_{-i}^T}^\tau = \theta_i^T] = P_i(\theta_i^T \mid \theta_i^{T-1}),$$

(and similarly for elements not written out) so both methods assign the path the same probability.

Note that (c_k) is a relaxed version of the test (b_k) : clearly $c_k \geq b_k$ for all k , whereas $c_k \rightarrow 0$ follows by the following observation, stated as a lemma for future reference:

Lemma B.4. *For every test (d_k) , $\lim_{k \rightarrow \infty} \max_{1 \leq j \leq k} \frac{j}{k} d_j = 0$.*

Proof. If $\max_k d_k = 0$, we are done. Otherwise, let $\eta > 0$ and put $\alpha = \frac{\eta}{\max_k d_k}$. Take \bar{k} such that $d_k < \eta$ for all $k \geq \alpha \bar{k}$. Let j_k be a maximizer for k . Then for any $k \geq \bar{k}$ we have $\frac{j_k}{k} d_{j_k} \leq \min\{d_{j_k}, \frac{j_k}{k} \max_k d_k\} < \eta$, where the second inequality follows by noting that if $\frac{j_k}{k} \max_k d_k \geq \eta$, then $j_k \geq \alpha k \geq \alpha \bar{k}$, and thus $d_{j_k} < \eta$. \square

Lemma B.5. *In every CRM $(f, (b_k), T)$, for all ρ and λ ,*

$$\mathbb{P}_\rho[\text{every } i \text{ passes } (c_k) \text{ at } (x^1, \dots, x^T)] \geq 1 - \varepsilon.$$

That is, the sequence (x^1, \dots, x^T) , which the mechanism uses to determine actions, has every player i passing the relaxed test (c_k) at the end of the CRM with high probability irrespective of the players' strategies. (The formula for (c_k) is of little interest. In what follows we only use uniformity in ρ .)

Proof. Fix a CRM $(f, (b_k), T)$, ρ , λ , and i . Let $1 < s < T$ and consider a history where player i fails the test (b_k) at (x^1, \dots, x^s) . The CRM then generates x_i^{s+1}, \dots, x_i^T by simulating i 's true type process. Thus continuation play is isomorphic to a situation where i reports the rest of his types truthfully. Hence Lemma B.3 implies that, conditional on failing (b_k) at (x^1, \dots, x^s) , i passes (b_k) at (x^{s+1}, \dots, x^T) with probability at least $1 - \frac{\varepsilon}{n}$. In this event, for any (θ, θ'_{-i}) , we can decompose the empirical frequency of x_i^t 's as

$$P_i^T(\cdot \mid \theta, \theta'_{-i}) = \frac{k_1}{k} \Phi_1 + \frac{k_2}{k} \Phi_2 + \frac{k - k_1 - k_2}{k} \Phi_3,$$

where

- $k_1 = \tau_i^{s-1}(\theta, \theta'_{-i})$ for (x^1, \dots, x^{s-1}) and $\Phi_1 = P_i^{s-1}(\cdot \mid \theta, \theta'_{-i})$,
- $k_2 = \tau_i^{T-s}(\theta, \theta'_{-i})$ for (x^{s+1}, \dots, x^T) and Φ_2 is the corresponding empirical frequency,
- Φ_3 is the empirical frequency of i 's reports in period τ ,
- $k = k_1 + k_2 + 1$ iff $(x^{\tau-1}, x^{\tau}_{-i}) = (\theta, \theta'_{-i})$; otherwise $k = k_1 + k_2$.

By definition of the Φ_i 's we then have

$$\begin{aligned} & \|P_i^T(\cdot \mid \theta, \theta'_{-i}) - P_i(\cdot \mid \theta_i)\| \\ & \leq \frac{k_1}{k} \|\Phi_1 - P_i(\cdot \mid \theta_i)\| + \frac{k_2}{k} \|\Phi_2 - P_i(\cdot \mid \theta_i)\| + \frac{k - k_1 - k_2}{k} \|\Phi_3 - P_i(\cdot \mid \theta_i)\| \\ & \leq \frac{k_1}{k} b_{k_1} + \frac{k_2}{k} b_{k_2} + \frac{1}{k} \leq c_k. \end{aligned}$$

We conclude that conditional on failing (b_k) at some (x^1, \dots, x^τ) , i passes (c_k) at (x^1, \dots, x^T) with probability at least $1 - \frac{\varepsilon}{n}$. On the other hand, if i never fails (b_k) ,

he passes (c_k) a priori as $c_k \geq b_k$. Thus $\mathbb{P}_\rho[i \text{ passes } (c_k) \text{ at } (x^1, \dots, x^T)] \geq 1 - \frac{\varepsilon}{n}$, which implies the result. \square

For all $T \in \mathbb{N}$, let

$$\Xi^T = \{(x^1, \dots, x^T) \in \Theta^T \mid \text{every } i \text{ passes } (c_k) \text{ at } (x^1, \dots, x^T)\}.$$

By Lemma B.5 the set Ξ^T has the desired probability given any ρ and λ . Hence for condition (2') it suffices to establish that for every sequence of sequences $(x^1, \dots, x^T) \in \Xi^T$, $T \in \mathbb{N}$, we have $P^T \rightarrow P$, and that the rate of convergence is uniform across all such sequences. As the first step in this accounting exercise, the next lemma shows that all transitions that have positive probability under the true process appear infinitely often in $(x^1, \dots, x^T) \in \Xi^T$ as $T \rightarrow \infty$.

Lemma B.6. *There exists a map $\kappa : \mathbb{N} \rightarrow \mathbb{R}$ with $\kappa(T) \rightarrow \infty$ such that for all $(x^1, \dots, x^T) \in \Xi^T$, $T \in \mathbb{N}$, and all $(\theta, \theta') \in \Theta^2$, if $P(\theta' \mid \theta) > 0$, then*

$$\tau^T(\theta, \theta') \geq \kappa(T).$$

Proof. Let $p = \min \{P_i(\theta'_i \mid \theta_i) \mid P_i(\theta'_i \mid \theta_i) > 0, (\theta_i, \theta'_i) \in \Theta_i^2, i = 1, \dots, n\}$. We ignore integer constraints throughout the proof to simplify notation.

Claim B.1. *Let $(x^1, \dots, x^T) \in \Xi^T$, $T \in \mathbb{N}$, and $(\theta, \theta') \in \Theta^2$. If $\tau^T(\theta, \theta') \geq k$, then for all i and all $\theta''_i \in \text{supp } P_i(\cdot \mid \theta_i)$,*

$$\tau^T(\theta, (\theta''_i, \theta'_{-i})) \geq k(p - c_k).$$

Proof of Claim. Since $(x^1, \dots, x^T) \in \Xi^T$, we have

$$|P_i^T(\theta''_i \mid \theta, \theta'_{-i}) - P_i(\theta''_i \mid \theta_i)| < c_{\tau_i^T(\theta, \theta'_{-i})} \leq c_{\tau^T(\theta, \theta')} \leq c_k.$$

Thus

$$\tau^T(\theta, (\theta''_i, \theta'_{-i})) = \tau_i^T(\theta, \theta'_{-i}) P_i^T(\theta''_i \mid \theta, \theta'_{-i}) \geq k(P_i(\theta''_i \mid \theta_i) - c_k) \geq k(p - c_k),$$

proving the Claim. \square

Now fix $(x^1, \dots, x^T) \in \Xi^T$. There exists $(\theta^0, \theta) \in \Theta^2$ such that

$$\tau^T(\theta^0, \theta) \geq \frac{T-1}{|\Theta^2|} \geq \frac{\frac{T}{|\Theta|} - 1}{|\Theta|} =: k_0^1.$$

We claim that for all $\theta^1 \in \text{supp } P(\cdot \mid \theta^0)$, we have $\tau^T(\theta^0, \theta^1) \geq k_n^1$, where k_n^1 is determined from k_0^1 by setting $l = 1$ in the recursion

$$(B.1) \quad k_i^l = k_{i-1}^l (p - c_{k_{i-1}^l}), \quad i = 1, \dots, n.$$

Indeed, any $\theta^1 \in \text{supp } P(\cdot \mid \theta^0)$ can be obtained from θ in n steps through the chain

$$(\theta_1, \dots, \theta_n), (\theta_1^1, \theta_2, \dots, \theta_n), (\theta_1^1, \theta_2^1, \theta_3, \dots, \theta_n), \dots, (\theta_1^1, \dots, \theta_n^1),$$

where $\theta_i^1 \in \text{supp } P_i(\cdot | \theta_i^0)$ for all i . Hence the bound follows by applying Claim B.2 n times.

We note then that for every $\theta^1 \in \text{supp } P(\cdot | \theta^0)$, there exists $\theta \in \Theta$ such that

$$\tau^T(\theta^1, \theta) \geq \frac{k_n^1 - 1}{|\Theta|} =: k_0^2.$$

Thus applying Claim B.2 again n times allows us to deduce that $\tau^T(\theta^1, \theta^2) \geq k_n^2$ for all $\theta^2 \in \text{supp } P(\cdot | \theta^1)$, where k_n^2 is determined from k_0^2 by setting $l = 2$ in (B.1). We observe that $k_n^2 < k_n^1$.

Continuing in this manner defines a decreasing sequence $(k_n^1, k_n^2, \dots, k_n^l, \dots)$ such that for each l , k_n^l is given by the n -step recursion (B.1) with initial condition

$$(B.2) \quad k_0^l = \frac{k_n^{l-1} - 1}{|\Theta|}, \quad k_n^0 = \frac{T}{|\Theta|}.$$

By construction, for any sequence $(\theta^0, \theta^1, \dots, \theta^L)$ such that $\prod_{l=1}^L P(\theta^l | \theta^{l-1}) > 0$, we have $\tau^T(\theta^{l-1}, \theta^l) \geq k_n^l \geq k_n^L$ for all l . Since P is irreducible, there exists $L < \infty$ such that every pair $(\theta', \theta'') \in \Theta^2$ with $P(\theta'' | \theta') > 0$ is along some such sequence starting from any θ^0 . For this L we have $\tau^T(\theta', \theta'') \geq k_n^L$ for all $(\theta', \theta'') \in \Theta^2$ such that $P(\theta'' | \theta') > 0$. Furthermore, the bound k_n^L is independent of (x^1, \dots, x^T) by inspection of (B.1) and (B.2).

It remains to argue that $k_n^L \rightarrow \infty$ as $T \rightarrow \infty$. But since L is finite, this follows by noting that k_0^L grows linearly in T by (B.2), and for each $l = 1, \dots, L$, $k_n^l \rightarrow \infty$ as $k_0^l \rightarrow \infty$ by (B.1). We may thus put $\kappa(T) = k_n^L$ to conclude the proof. \square

Now fix a sequence $(x^1, \dots, x^T) \in \Xi^T$ for each $T \in \mathbb{N}$. Let P^T denote the empirical matrix for (x^1, \dots, x^T) as defined in condition (2'). Consider $P^T(\theta' | \theta)$ for some $(\theta, \theta') \in \Theta^2$ such that $P(\theta' | \theta) > 0$. Write

$$P^T(\theta' | \theta) = \prod_{i=1}^n P_i^T(\theta'_i | \theta, \theta'_{\{i+1, \dots, n\}}),$$

where $\theta'_{\{i+1, \dots, n\}} = (\theta'_{i+1}, \dots, \theta'_n)$ and we have defined

$$P_i^T(\theta'_i | \theta, \theta'_{\{i+1, \dots, n\}}) = \sum_{y \in \prod_{j=1}^{i-1} \Theta_j} w(y) P_i^T(\theta'_i | \theta, (y, \theta'_{\{i+1, \dots, n\}})),$$

where $(y, \theta'_{\{i+1, \dots, n\}}) = (y_1, \dots, y_{i-1}, \theta'_{i+1}, \dots, \theta'_n)$ and the weight $w(y)$ is given by

$$w(y) = \frac{\tau_i^T(\theta, (y, \theta'_{\{i+1, \dots, n\}}))}{\sum_{z \in \prod_{j=1}^{i-1} \Theta_j} \tau_i^T(\theta, (z, \theta'_{\{i+1, \dots, n\}}))}.$$

Since $(x^1, \dots, x^T) \in \Xi^T$, we have for all $y \in \prod_{j=1}^{i-1} \Theta_j$,

$$\left| P_i^T(\theta'_i | \theta, (y, \theta'_{\{i+1, \dots, n\}})) - P_i(\theta'_i | \theta_i) \right| \leq c_{\tau_i^T(\theta, (y, \theta'_{\{i+1, \dots, n\}}))},$$

and thus

$$\left| P_i^T(\theta'_i \mid \theta, \theta'_{\{i+1, \dots, n\}}) - P_i(\theta'_i \mid \theta_i) \right| \leq \sum_{y \in \prod_{j=1}^{i-1} \Theta_j} w(y) c_{\tau_i^T(\theta, (y, \theta'_{\{i+1, \dots, n\}}))}.$$

Let $K(T) := \sum_{y \in \prod_{j=1}^{i-1} \Theta_j} \tau_i^T(\theta, (y, \theta'_{\{i+1, \dots, n\}}))$, and note that

$$K(T) \geq \tau_i^T(\theta, \theta'_{-i}) \geq \tau^T(\theta, \theta') \geq \kappa(T),$$

where the last inequality is by Lemma B.6 since $P(\theta' \mid \theta) > 0$ by assumption.

Thus

$$\sum_{y \in \prod_{j=1}^{i-1} \Theta_j} w(y) c_{\tau_i^T(\theta, (y, \theta'_{\{i+1, \dots, n\}}))} \leq |\Theta| \max_{1 \leq j \leq K(T)} \frac{j}{K(T)} c_j.$$

Since $K(T) \geq \kappa(T)$ and $\kappa(T) \rightarrow \infty$, we have $K(T) \rightarrow \infty$. Lemma B.4 then implies

$$\left| P_i^T(\theta'_i \mid \theta, \theta'_{\{i+1, \dots, n\}}) - P_i(\theta'_i \mid \theta_i) \right| \leq |\Theta| \max_{1 \leq j \leq K(T)} \frac{j}{K(T)} c_j \rightarrow 0 \quad \text{as } T \rightarrow \infty.$$

Therefore,

$$P^T(\theta' \mid \theta) \rightarrow \prod_i P_i(\theta'_i \mid \theta_i) = P(\theta' \mid \theta)$$

for all $(\theta, \theta') \in \Theta^2$ such that $P(\theta' \mid \theta) > 0$. Furthermore, $\kappa(T)$ is independent of the sequences (x^1, \dots, x^T) , $T \in \mathbb{N}$, by Lemma B.6, and hence convergence is uniform as desired.

To finish the proof, we observe that

$$1 - \sum_{\theta'' \notin \text{supp } P(\cdot \mid \theta)} P^T(\theta'' \mid \theta) = \sum_{\theta' \in \text{supp } P(\cdot \mid \theta)} P^T(\theta' \mid \theta) \rightarrow \sum_{\theta' \in \text{supp } P(\cdot \mid \theta)} P(\theta' \mid \theta) = 1,$$

which implies that $P^T(\theta' \mid \theta) \rightarrow 0$ for all $(\theta, \theta') \in \Theta^2$ such that $P(\theta' \mid \theta) = 0$. This completes the proof of condition (2') and that of Lemma 5.1.

B.3. Proof of Theorem 5.1. Fix $\varepsilon > 0$ once and for all. We first find the cutoff discount factor $\bar{\delta} < 1$ for every $T < \infty$. To this end, for $T < \infty$ and $\delta \in [0, 1]$, consider the problem

$$d_{\delta, T} = \sup_{u \in [0, 1]^{nT}} \left\| \frac{1}{T} \sum_{t=1}^T u^t - \frac{1 - \delta}{1 - \delta^T} \sum_{t=1}^T \delta^{t-1} u^t \right\|.$$

The objective function is continuous in (δ, u) on $[0, 1] \times [0, 1]^{nT}$ and $d_{1, T} = 0$. Thus the Maximum Theorem implies that for all $T < \infty$, there exists $\bar{\delta}(T) < 1$ such that for all $\delta > \bar{\delta}(T)$, $d_{\delta, T} = d_{\delta, T} - d_{1, T} \leq \frac{\varepsilon}{2}$.

Let $\eta = \frac{\varepsilon}{4|\Theta|}$. By Lemma 5.1 there exists a test (b_k) and a time $\bar{T} < \infty$ such that for all i , ρ_{-i} , λ , and $T > \bar{T}$, we have both $\theta_i^t = x_i^t$ for all t and $\|\pi^T - \pi\| < \eta$ with $\mathbb{P}_{\rho_i^*, \rho_{-i}}$ -probability at least $1 - \eta$. Therefore, for all v , i , ρ_{-i} , λ , $T > \bar{T}$, and

$\delta > \bar{\delta}(T)$, player i 's payoff from ρ_i^* in the CRM $(f^v, (b_k), T)$ satisfies

$$\begin{aligned} \left| v_i - \mathbb{E}_{\rho_i^*, \rho_{-i}} \left[\frac{1-\delta}{1-\delta^T} \sum_{t=1}^T \delta^t u_i(f^v(x^t), \theta_i^t) \right] \right| &\leq \left| v_i - \mathbb{E}_{\rho_i^*, \rho_{-i}} \left[\frac{1}{T} \sum_{t=1}^T u_i(f^v(x^t), \theta_i^t) \right] \right| + \frac{\varepsilon}{2} \\ &\leq \eta + (1-\eta) \sum_{\theta \in \Theta} |\pi(\theta) - \pi^T(\theta)| + \frac{\varepsilon}{2} \\ &< 2\eta |\Theta| + \frac{\varepsilon}{2} = \varepsilon. \end{aligned}$$

Thus player i can truthfully secure $v_i - \varepsilon$ as desired.

B.4. Proof of Corollary 5.1. We need two preliminary results. Let $v \in \mathbb{R}^n$. Take $p \in \mathbb{R}_+^n \setminus \{0\}$ such that for all $w \in V$, $p \cdot w \leq p \cdot v$. Let $\kappa > 0$ and define the set

$$\mathbf{Tr}(\kappa, v) = \{w \in \mathbb{R}^n \mid p \cdot w \leq p \cdot v, w_i \geq v_i - \kappa \ i = 1, \dots, n\}.$$

Lemma B.7. *Assume that $p \gg 0$ and $\sum_i p_i = 1$. Then, for all $w \in \mathbf{Tr}(\kappa, v)$, $\|w - v\| \leq \kappa \max\{\frac{1}{p_i} \mid i = 1, \dots, n\}$.*

Proof. Consider the problem $\max\{\|w - v\| \mid w \in \mathbf{Tr}(\kappa, v)\}$. This is a problem of maximizing a convex function on a convex and compact set. Corollary 32.3.2 in Rockafellar (1970) implies that the maximum is attained at extreme points of $\mathbf{Tr}(\kappa, v)$. Now, observe that $\mathbf{Tr}(\kappa, v)$ is a polytope that can be written as the intersection of $n + 1$ linear inequalities

$$\mathbf{Tr}(\kappa, v) = \bigcap_{i=0}^n \{w \in \mathbb{R}^n \mid w \cdot \beta_i \leq \alpha_i\},$$

where $\beta_0 = p$, $\alpha_0 = p \cdot v$, and for $i = 1, \dots, n$, $\beta_i \in \mathbb{R}^n$ is minus the unit vector having 1 in the i -th component and $\alpha_i = \kappa - v_i$.

The polytope $\mathbf{Tr}(\kappa, v)$ has $n + 1$ extreme points. To see this, note first that if one of the $n + 1$ linear inequalities defining $\mathbf{Tr}(\kappa, v)$ does not bind at some extreme point w , then all the other n linear inequalities must bind for otherwise we could obtain w as a convex combination of vectors in $\mathbf{Tr}(\kappa, v)$. It then follows that the set of extreme points of $\mathbf{Tr}(\kappa, v)$ equals $\{w^0, w^1, \dots, w^n\}$ where $w^0 = v - \kappa(1, \dots, 1)^T$ and for $i = 1, \dots, n$ $w_i^i = v_i + \frac{\kappa}{p_i} \sum_{j \neq i} p_j$ and $w_j^i = v_j - \kappa$ for $i \neq j$. We deduce that

$$\max\{\|w - v\| \mid w \in \mathbf{Tr}(\kappa, v)\} \leq \max\{\kappa, \max_{i=1, \dots, n} \frac{\kappa}{p_i}\} \leq \kappa \max\{\frac{1}{p_i} \mid i = 1, \dots, n\},$$

which proves the lemma. \square

Lemma B.8. *Fix $\bar{\varepsilon} > 0$. There exists \bar{T} such that for all $T \geq \bar{T}$ there exists $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$ the following hold:*

(1) For any $f: \Theta \rightarrow \Delta(A)$

$$\left\| \mathbb{E}_\pi[u(f(\theta), \theta)] - \frac{1-\delta}{1-\delta^T} \sum_{t=1}^T \mathbb{E}[u(f(\theta^t), \theta^t)] \right\| < \bar{\varepsilon}.$$

(2) For any payoff vector \bar{v} obtained as

$$\bar{v} = \frac{1-\delta}{1-\delta^T} \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} u(f^t(\theta^1, \dots, \theta^t), \theta^t) \right]$$

with decision rules $f^t: \Theta^t \rightarrow \Delta(A)$, $t = 1, \dots, T$, there exists $w \in V$ within distance $\bar{\varepsilon}$ of \bar{v} .

Proof. Let us first prove (1). Given a sequence of types $(\theta^1, \dots, \theta^T)$ denote the empirical distribution by $\pi^T \in \Delta(\Theta)$ and note that there exists \bar{T} sufficiently large such that

$$\mathbb{E}[\|\pi^T - \pi\|] < \frac{\bar{\varepsilon}}{2|\Theta|}.$$

As in the proof of Theorem 5.1, for any $T \geq \bar{T}$ we can take $\bar{\delta}(T) < 1$ big enough such that for all $(u^t)_{t=1}^T \subset [0, 1]^n$ and all $\delta > \bar{\delta}(T)$,

$$\left\| \frac{1}{T} \sum_{t=1}^T u^t - \frac{1-\delta}{1-\delta^T} \sum_{t=1}^T \delta^{t-1} u^t \right\| < \frac{\bar{\varepsilon}}{2}.$$

Therefore

$$\begin{aligned} & \left\| \mathbb{E}_\pi[u(f(\theta), \theta)] - \frac{1-\delta}{1-\delta^T} \sum_{t=1}^T \mathbb{E}[u(f(\theta^t), \theta^t)] \right\| \\ & \leq \left\| \mathbb{E}_\pi[u(f(\theta), \theta)] - \frac{1}{T} \sum_{t=1}^T \mathbb{E}[u(f(\theta^t), \theta^t)] \right\| + \frac{\bar{\varepsilon}}{2} \leq \mathbb{E}[|\Theta| \|\pi^T - \pi\|] + \frac{\bar{\varepsilon}}{2} < \bar{\varepsilon}, \end{aligned}$$

proving the first part of the lemma.

To prove (2), note that the set

$$\begin{aligned} & V(\delta, T) \\ & = \left\{ v \in \mathbb{R}^n \mid v = \frac{1-\delta}{1-\delta^T} \mathbb{E} \left[\sum_{t=1}^T u(f^t(\theta^1, \dots, \theta^t), \theta^t) \right] \text{ with } f^t: \Theta^t \rightarrow \Delta(A) \text{ all } t \right\} \end{aligned}$$

is convex. Therefore, we can find a stationary rule $f: \Theta \rightarrow \Delta(A)$ such that

$$\bar{v} = \frac{1-\delta}{1-\delta^T} \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} u(f(\theta^t), \theta^t) \right].$$

Define $w = \mathbb{E}_\pi[u(f(\theta), \theta)] \in V$ and apply the first part of the lemma to deduce that $\|w - \bar{v}\| < \bar{\varepsilon}$. \square

Let us now prove the corollary. Fix $\varepsilon > 0$. Let $v^1, \dots, v^{\bar{Q}}$ be the set of extreme points of $\mathbf{co}(\mathcal{P}(V))$. Let $f^q: \Theta \rightarrow A$ be the rule giving expected payoffs v^q . For each extreme point v^q , let $p^q \in \mathbb{R}_{++}^n$ with $\sum_{i=1}^n p_i^q = 1$ and $p^q \cdot w \leq p^q \cdot v^q$ for all $w \in V$. Take $\bar{\varepsilon} > 0$ defined as

$$\bar{\varepsilon} \left(1 + 2 \max \left\{ \frac{1}{p_i^q} \mid q = 1, \dots, \bar{Q}, i = 1, \dots, n \right\} \right) = \frac{\varepsilon}{2}.$$

Apply Theorem 5.1 to find a test (b_k) and \bar{T} such that for all $T > \bar{T}$ there exists a discount factor $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$ and all initial distributions λ , each player i can secure $v_i^q - \bar{\varepsilon}$ in the CRM $(f^q, (b_k), T)$ for all $q = 1, \dots, \bar{Q}$.

We additionally restrict \bar{T} and $\bar{\delta}$ to be large enough so that Lemma B.8 applies given $\bar{\varepsilon}$. Finally, for any $T > \bar{T}$ we will take $\bar{\delta}$ big enough such that for all $\delta > \bar{\delta}$, $(1 - \delta^T) \leq \frac{\varepsilon}{2}$. In the sequel, (b_k) , $T \geq \bar{T}$, and $\delta > \bar{\delta}$ are fixed.

Let now $v \in W$ as in the statement of the corollary. Take $\phi^v \in \Delta(A^\Theta)$ giving expected payoffs equal to v and such that there exists a family of rules $\tilde{f}^1, \dots, \tilde{f}^{\bar{Q}} \in A^\Theta$ with $\phi^v(\{\tilde{f}^1, \dots, \tilde{f}^{\bar{Q}}\}) = 1$, $\tilde{v}^q = \mathbb{E}_\pi[u(\tilde{f}^q(\theta), \theta)] \in V^c \cup \{v^1, \dots, v^{\bar{Q}}\}$ and $v = \sum_{q=1}^{\bar{Q}} \phi^v(\tilde{f}^q) \tilde{v}^q$.

Consider the block CRM $(\phi^v, (b_k), T)^\infty$ and let (ρ, μ) be a PBE of the block CRM. Take any history h of length mT , with $m \in \mathbb{N}$, right after the rule applying during the ensuing T rounds is realized. Let \tilde{f}^q be such rule and denote the expected normalized sum of discounted payoffs over the T -period block by $v^{\delta, mT} \in V(\delta, T)$. We first argue that $\|v^{\delta, mT} - \tilde{v}^q\| \leq \varepsilon/2$. From the first part of Lemma B.8, this inequality is obviously true if $\tilde{v}^q \in V^c$ because $\bar{\varepsilon} < \varepsilon/2$. If $\tilde{v}^q \in \{v^1, \dots, v^{\bar{Q}}\}$, note that $v_i^{\delta, mT} \geq \tilde{v}_i^q - \bar{\varepsilon}$ for all $i = 1, \dots, n$. From the second part of Lemma B.8, there exists $w^{\delta, mT} \in V$ such that $\|w^{\delta, mT} - v^{\delta, mT}\| < \bar{\varepsilon}$. It follows that for all $i = 1, \dots, n$, $w_i^{\delta, mT} \geq \tilde{v}_i^q - 2\bar{\varepsilon}$ and therefore $w^{\delta, mT} \in \mathbf{Tr}(2\bar{\varepsilon}, v^q)$. Applying Lemma B.7,

$$\|w^{\delta, mT} - \tilde{v}^q\| \leq 2\bar{\varepsilon} \max \left\{ \frac{1}{p_i^q} \mid i = 1, \dots, n \right\}.$$

It follows that $\|v^{\delta, mT} - \tilde{v}^q\| \leq \|v^{\delta, mT} - w^{\delta, mT}\| + \|w^{\delta, mT} - \tilde{v}^q\| \leq \varepsilon/2$. In any case, $\|v^{\delta, mT} - \tilde{v}^q\| \leq \varepsilon/2$. We can then compute the expected payoff vector right before the rule applying during the ensuing T rounds is realized, denoted $\bar{v}^{\delta, T}$, and deduce that $\|\bar{v}^{\delta, T} - v\| \leq \frac{\varepsilon}{2}$.

Take now an arbitrary history $h \in (\prod_{i=1}^n H_i^t) \times H^t$ of realized types, reports, public randomizations, and actions up to period $t \geq 1$ in the block CRM. Write the discounted sum of continuation payoffs from period t on as

$$v(h) = (1 - \delta) \mathbb{E}_{\rho, \phi^v} \left[\sum_{t' \geq t} \delta^{t'-t} u(a^{t'}, \theta^{t'}) \mid h \right],$$

where the expectation is taken conditional on h , given the strategy profile ρ and the randomized rule ϕ^v . Let $m^* = \arg \min \{mT \mid mT \geq t\}$ and rewrite the sum

above as

$$v(h) = \mathbb{E}_{\rho, \phi^v} \left[\sum_{t'=t}^{m^*T} \delta^{t'-t} (1-\delta) u(a^{t'}, \theta^{t'}) \mid h \right] \\ + \mathbb{E}_{\rho, \phi^v} \left[\delta^{m^*T-t} (1-\delta^T) \sum_{m \geq m^*} \delta^{(m-m^*)T} \bar{v}^{\delta, mT} \mid h \right].$$

The corollary now follows by observing that

$$\|v(h) - v\| \leq (1 - \delta^{m^*T-t}) \|v\| + \delta^{m^*T-t} \frac{\varepsilon}{2} \leq (1 - \delta^T) \|v\| + \frac{\varepsilon}{2} \leq \varepsilon.$$

APPENDIX C. A PROOF FOR SECTION 6

Proof of Lemma 6.1. The first part is immediate, given the construction of the block CRM $(\phi^i, (b_k), T)$.

To see the second part, fix a player i and the initial state $\theta_i^1 = \theta_i$. Let $P^{(t)}(\theta'_i \mid \theta_i) = \mathbb{P}[\theta_i^t = \theta'_i \mid \theta_i^1 = \theta_i]$. From Theorem 1.8.4 in Norris (1997), there exists a partition $(C_r^i)_{r=1}^{d^i}$ of Θ_i such that $P_i^{(n)}(\theta'_i \mid \theta_i) > 0$ only if $\theta_i \in C_r^i$ and $\theta'_i \in C_{r+n}^i$ for some $r \in \{1, \dots, d^i\}$, where we write $C_{nd^i+r}^i = C_r^i$. Observe that, without loss, we can assume that the initial state is such that $\theta_i \in C_1^i$.

From Theorem 1.8.5 in Norris (1997), there exists $N = N(\theta_i) \in \mathbb{N}$ such that for all $n \geq N$ and all $\theta'_i \in C_r^i$, $|P^{(nd^i+r)}(\theta'_i \mid \theta_i) - d^i \pi_i(\theta'_i)| \leq \frac{\varepsilon}{4|\Theta_i|}$. Note that for any such $n \geq N$,

$$\left| \sum_{r=1}^{d^i} \sum_{\theta'_i \in \Theta_i} \max_{a_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(nd^i+r)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \\ = \left| \sum_{r=1}^{d^i} \sum_{\theta'_i \in C_r^i} \max_{a_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(nd^i+r)}(\theta'_i \mid \theta_i) - d^i \pi_i(\theta'_i)) \right| \\ \leq \sum_{r=1}^{d^i} \sum_{\theta'_i \in C_r^i} \frac{\varepsilon}{4|\Theta_i|} \leq \frac{\varepsilon}{4}.$$

Now, note that for any δ and any $L \geq (N+1)d^i + 1$,

$$\left| \frac{1-\delta}{1-\delta^L} \sum_{t=1}^L \delta^{t-1} \mathbb{E}[\max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta_i^t) \mid \theta_i] - v_i \right| \\ \leq \frac{1-\delta^{Nd^i}}{1-\delta^L} 2 + \left| \frac{1-\delta}{1-\delta^L} \sum_{t=Nd^i+1}^L \delta^{t-1} \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(t)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right|.$$

To bound the second term, take $\bar{L} = \max\{nd^i \mid nd^i \leq L\} (\geq Nd^i + 1)$ and note that

$$\begin{aligned}
& \left| \sum_{t=Nd^i+1}^{\bar{L}} \delta^{t-1} \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(t)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \\
& \leq \sum_{n=N}^{\bar{L}/d^i-1} \delta^{nd^i-1} \left| \sum_{r=1}^{d^i} \delta^r \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(nd^i+r)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \\
& \leq \sum_{n=N}^{\bar{L}/d^i-1} \delta^{nd^i-1} \left\{ \left| \sum_{r=1}^{d^i} (1-\delta^r) \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(nd^i+r)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \right. \\
& \quad \left. + \left| \sum_{r=1}^{d^i} \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(nd^i+r)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \right\} \\
& \leq \sum_{n=N}^{\bar{L}/d^i-1} \delta^{nd^i-1} \left\{ (1-\delta^{d^i}) 2d^i |\Theta_i| + \frac{\varepsilon}{4} \right\} \\
& = \frac{\delta^{d^i N-1} - \delta^{\bar{L}-1}}{1-\delta^{d^i}} \left\{ (1-\delta^{d^i}) 2d^i |\Theta_i| + \frac{\varepsilon}{4} \right\},
\end{aligned}$$

and thus

$$\begin{aligned}
& \left| \frac{1-\delta}{1-\delta^{\bar{L}}} \sum_{t=Nd^i+1}^{\bar{L}} \delta^{t-1} \sum_{\theta'_i \in \Theta_i} \max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta'_i) (P^{(t)}(\theta'_i \mid \theta_i) - \pi_i(\theta'_i)) \right| \\
& \leq \frac{1-\delta}{1-\delta^{d^i}} \frac{\delta^{d^i N-1} - \delta^{\bar{L}-1}}{1-\delta^{\bar{L}}} \left\{ (1-\delta^{d^i}) 2d^i |\Theta_i| + \frac{\varepsilon}{4} \right\} \\
& \leq \left\{ (1-\delta^{d^i}) 2d^i |\Theta_i| + \frac{\varepsilon}{4} \right\} \leq \frac{\varepsilon}{2},
\end{aligned}$$

if δ is big enough (uniformly in L). Let $\delta(i) \in]0, 1[$ be such that the last inequality holds for all $\delta \geq \delta(i)$.

Now, let δ_{θ_i} be such that for all $\delta \geq \delta_{\theta_i}$, $L(\delta) \geq (N(\theta_i) + 1)d^i + 1$ and

$$\frac{1-\delta^{Nd^i}}{1-\delta^{L(\delta)}} 2 + \frac{1-\delta}{1-\delta^{L(\delta)}} 2d^i < \frac{\varepsilon}{2}.$$

Defining $\delta_{i,\theta_i} = \max\{\delta_{\theta_i}, \delta(i)\}$, it then follows that for all $\delta \geq \delta_{i,\theta_i}$,

$$\left| \frac{1-\delta}{1-\delta^{L(\delta)}} \sum_{t=1}^{L(\delta)} \delta^{t-1} \mathbb{E}[\max_{a_i \in A_i} u_i(a_i, a_{-i}^i, \theta_i^t) \mid \theta_i] - \underline{v}_i \right| < \varepsilon.$$

Finally, taking $\delta_1 = \max\{\delta_0, \max\{\delta_{i,\theta_i} \mid i = 1, \dots, n, \theta_i \in \Theta_i\}\}$ gives the result. \square

APPENDIX D. FORMAL DESCRIPTION OF STRATEGIES AND BELIEFS

In this appendix we present a formal description of the equilibrium strategies and beliefs in terms of an automaton.

It will be useful to introduce some notation to describe how actions are selected in each block CRM $(\phi^i, (b_k), T)^\infty$. For an arbitrary sequence $y = (y^1, \dots, y^t)$, let $y_{\geq n}$ be the sequence (y^n, \dots, y^t) . Let $\bar{T}(t) = \max\{nT + 1 \mid n \geq 0, nT + 1 \leq t\}$. Given $(m^1, \omega^1, \dots, m^t, \omega^t)$, the block CRM selects the action

$$a^t = \phi^i(\chi^{t-\bar{T}(t)+1}((m^1, \omega^1, \dots, m^t, \omega^t)_{\geq \bar{T}(t)}), \omega^{\bar{T}(t)}) \in A,$$

where χ was defined in Remark 5.1.

Let $H(0) \subset H$ denote the set of feasible public histories in the block CRM $(\phi^0, (b_k), T)^\infty$, the restriction being the requirement that actions be the ones the mechanisms would select. For $i = 1, \dots, n$, let $H(i) \subset H$ denote the set of feasible public histories in the punishment mechanism $(L, (\phi^i, (b_k), T)^\infty)$. The histories in $H(i)$ consist of stick-subphase histories (periods $t = 1, \dots, L$) and carrot-subphase histories (periods $t \geq L + 1$). For the stick-subphase histories, the restriction is that the action played by each player $j \neq i$ coincide with a_j^i for $t = 1, \dots, L$. Denote by $H(i, s)$ the set of all such histories. Feasible carrot-subphase histories are such that their first L periods coincide with an element of $H(i, s)$, and for $t > L$ the actions are the ones the mechanism would select. Denote by $H(i, c)$ the set of all such histories. By definition, $H(i) = H(i, s) \cup H(i, c)$. Note that the null history \emptyset is an element of each of the sets $H(0), \dots, H(n)$.

Take $h \in H(i)$, for some $i = 0, 1, \dots, n$, and $(m, \omega, a) \in \Theta \times [0, 1] \times A$ such that $(h, (m, \omega, a)) \notin H(i)$. Then, by construction, there exists some j whose action a_j does not match the action the corresponding mechanism would have selected given the history.

Our construction of the automaton distinguishes between two different stages $r \in \{0, 1\}$ within each period $t \geq 1$. The idea is that $r = 0$ corresponds to the reporting stage (i.e., $t.2$), and $r = 1$ to the action stage (i.e., $t.4$). Thus, the index used to describe the evolution of the automaton is the pair (t, r) endowed with the lexicographic order $(t, 0) < (t, 1) < (t + 1, 0)$.

Let $S = H(0) \cup H(1) \cup \dots \cup H(n)$, and define

$$\mathcal{B} = \bigcup_{t \geq 1} \left(\prod_{i=1}^n \Delta(\Theta_i^t) \right).$$

The state space of the automaton is the product

$$\{0, 1, \dots, n\} \times S \times \mathcal{B} \times \mathcal{B},$$

and we write (ι, s, \bar{B}, B) for a generic element. The first component ι indicates the mechanism players are mimicking, while s indicates the current history in

the mechanism. The third component \bar{B} indicates the (public) beliefs the players entertained about the private histories of types when the mechanism ι was triggered. The fourth component B indicates the players' current beliefs about the whole history of private types. For any $s \in H^t$, let $T(s) = t$.

Before describing the evolution of the automaton, we choose equilibria for the mechanisms as follows: Let $(\rho^{0,\lambda}, \mu^{0,\lambda})$ be a PBE assesment for $(\phi^0, (b_k), T)^\infty$ given initial beliefs λ . For all $i = 1, \dots, n$, all beliefs $B \in \mathcal{B}$ such that $B \in \prod_{i=1}^n \Delta(\Theta_i^t)$, and all punishment mechanisms $(L, (\phi^i, (b_k), T)^\infty)$, take a PBE assesment $(\rho^{i,B}, \mu^{i,B})$ such that the strategy $\rho^{i,B}$ depends on B only through the marginal distribution of the period- t profile θ^t . For completeness, we extend the belief system $\mu^{i,B}$ so that, for each $k \geq 0$, and each public history $\tilde{h} \in H(i)$ of the form $(m^t, \omega^t, a^t, \dots, m^{t+k}, \omega^{t+k}, a^{t+k})$ or

$$(m^t, \omega^t, a^t, \dots, m^{t+k}, \omega^{t+k}, a^{t+k}, m^{t+k+1}, \omega^{t+k+1}),$$

it gives a distribution over the entire private histories up to period $t + k + 1$, denoted

$$\mu^{i,B}((\theta^1, \dots, \theta^t, \theta^{t+1}, \dots, \theta^{t+k+1}) \mid \tilde{h}).$$

This extended belief system is computed using Bayes rule given the prior $B \in \prod_{i=1}^n \Delta(\Theta_i^t)$ and is assumed to satisfy the requirements imposed on PBE beliefs.

The automaton evolves as follows: Let $\iota^{1,0} = 0$, $s^{1,0} = \emptyset$, $\bar{B}^{1,0} = B^{1,0} = \lambda$. For any $t \geq 1$, define $(\iota^{t,1}, s^{t,1}, \bar{B}^{t,1}, B^{t,1})$ as follows. Let $\iota^{t,1} = \iota^{t,0}$, $s^{t,1} = (s^{t,0}, (m^t, \omega^t))$, $\bar{B}^{t,1} = \bar{B}^{t,0}$ and

$$B^{t,1}(\theta^1, \dots, \theta^t) = \mu^{\iota^{t,1}, \bar{B}^{t,1}}((\theta^1, \dots, \theta^t) \mid s^{t,1}).$$

For $t \geq 2$, define $(\iota^{t,0}, s^{t,0}, \bar{B}^{t,0}, B^{t,0})$ as follows. Let

$$B^{t,0}(\theta^1, \dots, \theta^t) = \mu^{\iota^{t-1,1}, \bar{B}^{t-1,1}}((\theta^1, \dots, \theta^t) \mid s^{t-1,1}).$$

If $(s^{t-1,1}, a^{t-1}) \in H(\iota^{t-1,1})$, then $\iota^{t,0} = \iota^{t-1,1}$, $s^{t,0} = (s^{t-1,1}, a^{t-1})$, $\bar{B}^{t,0} = \bar{B}^{t-1,1}$. If $(s^{t-1,1}, a^{t-1}) \notin H(\iota^{t-1,1})$, then take any j whose action a_j^{t-1} differed from what the mechanism mandated and define $\iota^{t,0} = j$, $\bar{B}^{t,0} = B^{t,0}$ and $s^{t,0} = \emptyset$.

The assesment (σ, μ) is constructed as follows. Fix a player i , a private history $h_i^t = (\theta_i^1, \dots, \theta_i^t) \in \Theta_i^t$ and a public history $h^t \in H^t$. If $h^t \in (\Theta \times [0, 1] \times A)^{t-1}$, let

$$\sigma_i(\cdot \mid h^t, h_i^t) = \rho_i^{\iota^t, \bar{B}^t}(\cdot \mid s^{t,0}, (\theta_i^{t-T(s^{t,0})}, \dots, \theta_i^t)).$$

If $h^t \in (\Theta \times [0, 1] \times A)^{t-1} \times \Theta \times [0, 1]$, let

$$\sigma_i(\cdot \mid h^t, h_i^t) = \begin{cases} a_i^{t,1} & \text{if } \iota^{t,1} \neq 0, i \neq \iota^{t,1} \\ & \text{and } T(s^{t,1}) \leq L-1, \\ \rho_i^{t,1, \bar{B}^{t,1}}(\cdot \mid s^{t,1}, (\theta_i^{t-T(s^{t,1})}, \dots, \theta_i^t)) & \text{if } \iota^{t,1} = i \\ & \text{and } T(s^{t,1}) \leq L-1, \\ \phi^i(\chi^{l-\bar{T}(l)+1}((\tilde{m}^1, \tilde{\omega}^1, \dots, \tilde{m}^l, \tilde{\omega}^l)_{\geq \bar{T}(l)}), \tilde{\omega}^{\bar{T}(l)}) & \text{otherwise,} \end{cases}$$

where $(\tilde{m}^1, \tilde{\omega}^1, \dots, \tilde{m}^l, \tilde{\omega}^l) = ((s^{t,1})_{\geq \min\{\iota^{t,1}\}_{L+1}}) \setminus A$ is the history of reports and public randomizations in the current mechanism, excluding those occurring during the stick subphase if $\iota^t \neq 0$. (Here, for any vector x having some components in A , $x \setminus A$ denotes all the components of x that are not in A .)

Public beliefs about player i are as follows. For any $h^t \in (\Theta \times [0, 1] \times A)^{t-1}$, let

$$\mu_i^t((\theta_i^1, \dots, \theta_i^t) \mid h^t) = B_i^{t,0}(\theta_i^1, \dots, \theta_i^t),$$

and for any $h^t \in (\Theta \times [0, 1] \times A)^{t-1} \times \Theta \times [0, 1]$, let

$$\mu_i^t((\theta_i^1, \dots, \theta_i^t) \mid h^t) = B_i^{t,1}(\theta_i^1, \dots, \theta_i^t).$$

Note that unless a deviation causes ι to change, play mimics some fixed PBE $(\rho^{i,B}, \mu^{i,B})$, $i = 0, 1, \dots, n$, $B \in \mathcal{B}$. Thus to show that the strategies are sequentially rational, it suffices to show that any one-stage deviations that triggers a change in ι is unprofitable. By inspection of the automaton, such deviations consist of one-stage deviations in actions at the following three classes of states:

- (1) $\iota^{t,1} = 0$;
- (2) $\iota^{t,1} = i \neq 0$, $T(s^{t,1}) \leq L-1$;
- (3) $\iota^{t,1} = i \neq 0$, $T(s^{t,1}) \geq L$.

These are informally described in Section 6.2 as cooperative-phase, stick-subphase, and carrot-subphase histories. Sequential rationality follows by the analysis there.

APPENDIX E. A PROOF FOR SECTION 7

Proof of Proposition 7.1. We prove first that $V \cap \mathbb{R}_+^n \subset W \cap \mathbb{R}_+^n$. More specifically, we show that for any $v \in V \cap \mathbb{R}_+^n$, there exists $\alpha \in [0, 1]$ and $w \in \mathcal{P}(V)$ such that $v = \alpha w$. Since $0 \in V^c$, this implies $v \in W \cap \mathbb{R}_+^n$. Define

$$Y = \left\{ v \in \mathbb{R}^n \mid \exists s: \Theta \rightarrow \mathbb{R}_+^n \text{ s.t. } v_i = \mathbb{E}_\pi[(r - \theta_i)s_i(\theta)] \forall i \text{ and } \sum_{i=1}^n s_i(\theta) \leq 1 \forall \theta \right\}.$$

We prove at the end of this proof that

$$(E.1) \quad V \cap \mathbb{R}_+^n = Y.$$

The set $Y \subset \mathbb{R}_+^n$ can be interpreted as the set of utility vectors in an exchange economy consisting of $|\Theta|$ goods and n agents, with total endowment of each good equal to 1, and with linear preferences over bundles $s_i(\cdot) \in \mathbb{R}^\Theta$ that are strongly monotone (because $r > \theta_i$ for all θ_i) for each i . Varian (1974) provides results about the Pareto frontiers of exchange economies, and Step 1 in the proof of his Theorem 2.4 shows that for any $v \in V \cap \mathbb{R}_{++}^n = Y \cap \mathbb{R}_{++}^n$, we can find $\alpha \leq 1$ and $w \in \mathcal{P}(Y) \subset \mathcal{P}(V)$ such that $v = \alpha w$. Therefore, $v \in W \cap \mathbb{R}_+^n$ and $V \cap \mathbb{R}_+^n \subset W \cap \mathbb{R}_+^n$.

From feasibility and individual rationality, $E(\delta) \subset V(\delta) \cap \mathbb{R}_+^n$. Therefore, $\limsup_{\delta \rightarrow 1} E(\delta) \subset V \cap \mathbb{R}_+^n$. The full-dimensionality assumption is satisfied in the Bertrand game. (Indeed, let $v^0 \in \mathbb{R}^n$ be the zero vector and for $i = 1, \dots, n$ let $\bar{v}^i \in \mathbb{R}^n$ be the vector which equals $\mathbb{E}_{\pi_i}[r - \theta_i]$ in the i th component and zero otherwise; clearly $\{v^i\}_{i=0}^n \subset W$ is a family of affinely independent vectors.) Theorem 4.1 then implies that $W \cap \mathbb{R}_+^n \subset \liminf_{\delta \rightarrow 1} E(\delta)$. Since $V \cap \mathbb{R}_+^n \subset W \cap \mathbb{R}_+^n$, it follows that $\lim_{\delta \rightarrow 1} E(\delta) = V \cap \mathbb{R}_+^n$.

It remains to prove (E.1). Take $v \in V \cap \mathbb{R}_+^n$ and find a pricing rule $p: \Theta \rightarrow \mathbb{R}_+$ and an allocation $\bar{s}: \Theta \rightarrow [0, 1]^n$ describing the expected sale of each firm $i = 1, \dots, n$, with $\sum_{i=1}^n \bar{s}_i(\theta) \leq 1$, such that $v_i = \mathbb{E}[(p(\theta) - \theta_i)\bar{s}_i(\theta)]$ for all i . Define a new feasible allocation $\bar{s}: \Theta \rightarrow [0, 1]^n$ as follows. For all $\theta \in \Theta$ and i with $(p(\theta) - \theta_i)s_i(\theta) > 0$, define $\bar{s}_i(\theta) = \frac{p(\theta) - \theta_i}{r - \theta_i} s_i(\theta) \leq s_i(\theta)$. Leave all the other components unchanged (i.e. $\bar{s}_i(\theta) = s_i(\theta)$ for θ and i such that $(p(\theta) - \theta_i)s_i(\theta) \leq 0$). Now, construct a rule s as follows. For each i let $N_i \subset \Theta$ be the set of all θ such that $(p(\theta) - \theta_i)\bar{s}_i(\theta) < 0$. For $\theta \in N_i$, let $s_i(\theta) = 0$. Now, let

$$k_i = \frac{v_i}{\sum_{\theta_i \notin N_i} (r - \theta_i)\bar{s}_i(\theta)\pi(\theta)} = \frac{v_i}{v_i - \sum_{\theta_i \in N_i} (p(\theta) - \theta_i)\bar{s}_i(\theta)\pi(\theta)} \leq 1,$$

and define $s_i(\theta) = k_i \bar{s}_i(\theta)$ for all $\theta \notin N_i$. It is obviously true that $\sum_{i=1}^n s_i(\theta) \leq 1$ and

$$\mathbb{E}_\pi[(r - \theta_i)s_i(\theta)] = \sum_{\theta \notin N_i} (r - \theta_i)s_i(\theta) = v_i.$$

Therefore, $v \in Y$ and $V \cap \mathbb{R}_+^n \subset Y$. The other inclusion is immediate. \square

REFERENCES

- ABREU, D., P. DUTTA, AND L. SMITH (1994): "The Folk Theorem for Repeated Games: A NEU Condition," *Econometrica*, 62(4), 939–948.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.
- AMEMIYA, T. (1985): *Advanced Econometrics*. Harvard University Press.
- ATHEY, S., AND K. BAGWELL (2008): "Collusion with Persistent Cost Shocks," *Econometrica*, 76(3), 493–540.

- ATHEY, S., K. BAGWELL, AND C. SANCHIRICO (2004): “Collusion and price rigidity,” *Review of Economic Studies*, 71(2), 317–349.
- ATHEY, S., AND I. SEGAL (2007): “An Efficient Dynamic Mechanism,” Discussion paper, Harvard University.
- AUMANN, R., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. The MIT press.
- BERGEMANN, D., AND J. VÄLIMÄKI (2010): “The Dynamic Pivot Mechanism,” *Econometrica*, 78(2), 771–789.
- BILLINGSLEY, P. (1961): “Statistical Methods in Markov Chains,” *The Annals of Mathematical Statistics*, 32, 12–40.
- BLACKWELL, D. (1956): “An Analog of the Minimax Theorem for Vector Payoffs,” *Pacific Journal of Mathematics*, 6(1), 1–8.
- COLE, H., J. DOW, AND W. ENGLISH (1995): “Default, Settlement, and Signalling: Lending Resumption in a Reputational Model of Sovereign Debt,” *International Economic Review*, 36(2), 365–385.
- COLE, H., AND N. KOCHERLAKOTA (2001): “Dynamic Games with Hidden Actions and Hidden States,” *Journal of Economic Theory*, 98(1), 114–126.
- CREMER, J., AND R. P. MCLEAN (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56(6), 1247–1257.
- DUTTA, P. (1995): “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66(1), 1–32.
- FUDENBERG, D., AND D. LEVINE (1983): “Subgame Perfect Equilibria of Finite and Infinite Horizon Games,” *Journal of Economic Theory*, 31(2), 251–268.
- (1994): “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, 62(1), 103–135.
- FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62(5), 997–997.
- FUDENBERG, D., AND E. MASKIN (1986): “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54(3), 533–554.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press.
- FUDENBERG, D., AND Y. YAMAMOTO (2011a): “The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring,” *Journal of Economic Theory*, 146(4).
- (2011b): “Learning from Private Information in Noisy Repeated Games,” *Journal of Economic Theory*, 148(5).
- GOSSNER, O. (1995): “The Folk Theorem for Finitely Repeated Games with Mixed Strategies,” *International Journal of Game Theory*, 24(1), 95–107.
- HÖRNER, J., AND J. JAMISON (2007): “Collusion with (Almost) No Information,” *The RAND Journal of Economics*, 38(3), 804–822.

- HÖRNER, J., AND S. LOVO (2009): “Belief-Free Equilibria in Games With Incomplete Information,” *Econometrica*, 77(2), 453–487.
- HÖRNER, J., S. LOVO, AND T. TOMALA (2011): “Belief-free Equilibria in Games with Incomplete Information: Characterization and Existence,” *Journal of Economic Theory*, 146(5), 1770–1795.
- HÖRNER, J., D. ROSENBERG, E. SOLAN, AND N. VIEILLE (2010): “On a Markov Game with One-Sided Incomplete Information,” *Operations Research*, 58, 1107–1115.
- HÖRNER, J., T. SUGAYA, S. TAKAHASHI, AND N. VIEILLE (2011): “Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem,” *Econometrica*, 79(4), 1277–1318.
- JACKSON, M. O., AND H. F. SONNENSCHNEIN (2007): “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75(1), 241–258.
- JEHIEL, P., AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69(5), 1237–59.
- KOCHERLAKOTA, N. (1996): “Implications of Efficient Risk Sharing without Commitment,” *The Review of Economic Studies*, pp. 595–609.
- KREPS, D., AND R. WILSON (1982): “Reputation and Imperfect Information,” *Journal of Economic Theory*, 27, 253–279.
- LEVIN, J. (2003): “Relational Incentive Contracts,” *American Economic Review*, 93(3), 835–857.
- MAGGI, G., AND M. MORELLI (2006): “Self-Enforcing Voting in International Organizations,” *American Economic Review*, 96(4), 1137–1158.
- MAILATH, G., AND L. SAMUELSON (2001): “Who Wants a Good Reputation?,” *Review of Economic Studies*, 68(2), 415–441.
- (2006): *Repeated Games and Reputations*. Oxford University Press, Oxford.
- MASSART, P. (1990): “The Tight Constant in The Dvoretzky-Kiefer-Wolfowitz Inequality,” *The Annals of Probability*, 18, 1269–1283.
- MILGROM, P., AND J. ROBERTS (1982): “Predation, Reputation, and Entry Deterrence,” *Journal of Economic Theory*, 27(2), 280–312.
- MYERSON, R. B., AND M. A. SATTERTHWAITTE (1983): “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 29(2), 265–281.
- NEYMAN, A. (2008): “Existence of Optimal Strategies in Markov games with incomplete information,” *International Journal of Game Theory*, 37, 581–596.
- NORRIS, J. (1997): *Markov Chains*. Cambridge University Press.
- PALFREY, T. R., AND H. ROSENTHAL (1988): “Private Incentives in Social Dilemmas : The Effects of Incomplete Information and Altruism,” *Journal of Public Economics*, 35(3), 309–332.

- PESKI, M. (2008): “Repeated Games with Incomplete Information on One Side,” *Theoretical Economics*, 3(1), 29–84.
- PHELAN, C. (2006): “Public Trust and Government Betrayal,” *Journal of Economic Theory*, 130(1), 27–43.
- RADNER, R. (1981): “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, 49(5), 1127–1148.
- RENAULT, J. (2006): “The Value of Markov Chain Games with Lack of Information on One Side,” *Mathematics of Operations Research*, 31, 490–512.
- RENAULT, J., E. SOLAN, AND N. VIEILLE (2011): “Dynamic Sender-Receiver Games,” Discussion paper, TSE.
- ROCKAFELLAR, R. T. (1970): *Convex Analysis*. Princeton University Press.
- SKRZYPACZ, A., AND H. HOPENHAYN (2004): “Tacit Collusion in Repeated Auctions,” *Journal of Economic Theory*, 114(1), 153–169.
- TOWNSEND, R. (1982): “Optimal Multi-Period Contracts and the Gain from Enduring Relationships under Private Information,” *Journal of Political Economy*, 90(6), 1166–1186.
- VARIAN, H. (1974): “Equity, Envy, and Efficiency,” *Journal of Economic Theory*, 9(1), 63–91.
- WANG, C. (1995): “Dynamic Insurance with Private Information and Balanced Budgets,” *The Review of Economic Studies*, 62(4), 577–595.
- WATSON, J. (2002): “Starting Small and Commitment,” *Games and Economic Behavior*, 38(1), 176–199.

CENTER FOR APPLIED ECONOMICS, DEPARTMENT OF INDUSTRIAL ENGINEERING, UNIVERSITY OF CHILE, SANTIAGO, CHILE

DEPARTMENT OF ECONOMICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA, USA