

Available online at www.sciencedirect.com

SciVerse ScienceDirect

Journal of Economic Theory ••• (••••) ••••••

 JOURNAL OF
**Economic
 Theory**

www.elsevier.com/locate/jet

Dynamic sender–receiver games

Jérôme Renault ^a, Eilon Solan ^b, Nicolas Vieille ^{c,*}

^a TSE (GREMAQ, Université Toulouse 1), 21 allée de Brienne, 31000 Toulouse, France

^b School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel

^c Department of Economics and Decision Sciences, HEC Paris, 1, rue de la Libération, 78 351 Jouy-en-Josas, France

Received 20 February 2011; final version received 11 April 2012; accepted 31 July 2012

Abstract

We consider a dynamic version of sender–receiver games, where the sequence of states follows an irreducible Markov chain observed by the sender. Under mild assumptions, we provide a simple characterization of the limit set of equilibrium payoffs, as players become very patient. Under these assumptions, the limit set depends on the Markov chain only through its invariant measure. The (limit) equilibrium payoffs are the feasible payoffs that satisfy an individual rationality condition for the receiver, and an incentive compatibility condition for the sender.

© 2013 Elsevier Inc. All rights reserved.

JEL classification: C73; D83

Keywords: Sender–receiver; Dynamic games; Bayesian games; Communication; Repeated games

1. Introduction

Since [8], sender–receiver games, or cheap-talk games, have become a natural framework for studying issues of information transmission between a privately informed ‘expert’ and an uninformed decision maker, where the two parties have non-aligned interests.

[☆] This research was supported by the Israel Science Foundation (grant number 212/09), the Agence Nationale de la Recherche (grant ANR-10-BLAN 0112), the Fondation HEC, and by the Google Inter-university center for Electronic Markets and Auctions.

^{*} Corresponding author.

E-mail addresses: jerome.renault@tse-fr.eu (J. Renault), eilons@post.tau.ac.il (E. Solan), vieille@hec.fr (N. Vieille).

When the decision maker acts only *once*, the extent to which information can be shared at equilibrium has been studied extensively, when ‘talk’ takes place prior to the decision stage. While [8], see also [14], have focused on the case where communication is limited to a single costless and non-verifiable message from the sender to the receiver, more recent papers have shown that this restriction is not innocuous, and have characterized the equilibrium outcomes for general cheap-talk games, see [18,2]. This line of work has been motivated by numerous concrete situations. We refer to [19,11,23] for a discussion of these applications.

The present work is motivated by the following observation. Whether the sender is a financial advisor who provides advice to a client, an expert who is consulted on a project, or a referee on a project/person, the situation often calls for a dynamic approach. Indeed, the financial advisor provides advice on a series of investments, and the expert and the referee may be consulted on successive, related projects.

Such a situation is analyzed in [13]. They assume that the sender repeatedly sends messages, the receiver repeatedly makes decisions, while the state of the world remains *fixed* throughout. Within the framework of [8] (continuum of states/messages), they show that, for some specifications on the initial distribution on states, (necessarily complex) equilibria exist, that achieve full revelation of the state of the world in finite time.

We here deal with situations in which the state of the world may change through time. Specifically, we assume that the successive states form an irreducible Markov chain over some finite set. In every stage, the sender issues a message/recommendation, and the receiver makes a decision. States are only known to the sender, and payoffs only depend on the current state and on the receiver’s decision, but not on the message sent by the sender.

Since states are autocorrelated, any information disclosed in stage n provides valuable information in later stages as well, as in [13]. Yet, since the Markov chain is irreducible, this information becomes eventually valueless.

Intuitively, the inter-temporal situation puts some restrictions on the players’ behavior. As an illustration, the opinion of an expert who systematically provides laudatory reports will eventually come to be discounted, if not ignored, since the decision maker is aware of the fact that the time-average report of the quality of people/projects should reflect the invariant measure of the states of the world. On the other hand, an expert who genuinely provides accurate information to promote efficiency, but sees that the decision maker only acts in his interests, may become wary and may stop to provide valuable information to the decision maker. As is well known from repeated games, the sender may indeed provide powerful incentives by conditioning his future communication policy on the behavior of the decision maker. Similar insights already appear in the literature on dynamic contracting, see [5,6,4].

Our paper relates to the recent and growing literature on incomplete information games, in which the uncertainty evolves according to a stochastic process. Examples of such papers include, e.g., [1,9,20–22,24,15] and, especially, [10]. We provide a characterization of the limit set of sequential equilibrium payoffs, when players are very patient. Our main findings are the following. We first show (Theorem 1) that a feasible payoff vector is a (limit) equilibrium payoff as soon as the following two conditions are met. On the one hand, the payoff of the receiver should be at least his babbling equilibrium payoff. This condition is an individual rationality condition. Indeed, the latter payoff is equal to the receiver’s minmax payoff in the dynamic game since the receiver has the option to ignore the announcements of the sender. On the other hand, the sender’s payoff should satisfy an incentive compatibility condition, which reflects the fact that the sender has the option of substituting artificially generated states to the

true ones when playing, as long as the artificial states are statistically indistinguishable from the true ones. As it turns out, this incentive constraint takes the form of finitely many linear inequalities.

In the corresponding equilibria, with high probability the sender truthfully reports the current state most of the time while the receiver responds in a stationary manner to the announcements of the sender, and checks that the distribution of these announcements is consistent with the invariant measure of the underlying Markov chain.¹

We next show (Theorem 2) that the converse inclusion holds under some additional condition on the Markov chain, which we call Assumption A: any limit equilibrium payoff must satisfy the individual rational condition and the specific version of the incentive compatibility requirement of Theorem 1.

A noteworthy consequence is that, under Assumption A, the limit set of equilibrium payoffs does not depend on *how* successive states are correlated, nor on fine details of the sequence of states, but only on the invariant measure. It is also irrelevant whether the sender learns some, or even all, of the realization of the future states in advance. In particular, the set of equilibrium payoffs can be computed *as if* successive states were independent.

Our results are valid for a large (open) class of payoff functions for the static game, but not for all of them. More precisely, we prove that for generic payoff functions (and under Assumption A), either our results hold, or all equilibria of the repeated game are payoff-equivalent to babbling equilibria.

The paper is organized as follows. The model is described in Section 2. In Section 3 we explain most insights by means of an example. The main results appear in Section 4, together with an illustration. Proofs are discussed in Section 5 and Appendices A, B, C, and D. Additional results and comments are provided in Section 6. Appendices A, B, C, and D contain all proof details.

2. Model

We study dynamic sender–receiver games, in which the state of the world changes through time. At each stage $n \geq 1$, the sender (player 1) observes the current state of the world $s_n \in S$, and makes an announcement $a_n \in A$. Upon observing a_n , the receiver (player 2) chooses an action $b_n \in B$. The current action b_n , together with the current state s_n , determines the utility vector $u(s_n, b_n) \in \mathbf{R}^2$ at stage n . Only the action b_n is then publicly disclosed. We thus maintain the assumption that payoffs are not observed. The two players share a common discount factor δ .

We assume throughout that the set of states S , the set of messages A , and the set of actions B are finite. We also assume that there are at least as many messages as states. This assumption ensures that the only motives for concealing the state are strategic. We thus leave aside situations in which, due to capacity constraints, the sender might be forced to choose which feature of the state to reveal. For simplicity, we will actually assume throughout that the set A of messages coincides with the set S of states. (As will be seen, this assumption is without loss of generality in our setup.)

¹ While this is reminiscent of the revelation principle, we stress that no revelation principle applies in our setting.

l	r
$c, 2$	$2, 1$
State L	

l	r
$1, -1$	$2, 1$
State R	

Fig. 1. The payoffs of the two players.

We assume that the states (s_n) follow a Markov chain over S , with transition function $p(\cdot | \cdot)$, which is irreducible and aperiodic.² The Markov chain therefore admits a unique invariant measure $m \in \Delta(S)$. For convenience, we assume that the first state, s_1 , is drawn according to m . This ensures that the law of s_n is equal to m , for every $n \geq 1$.

In this setup, a strategy of the sender maps past and current realized states and past play into a mixed message, and is thus a function $\sigma : \bigcup_{n \geq 0} (S \times A \times B)^n \times S \rightarrow \Delta(A)$, while a strategy of the receiver is a function $\tau : \bigcup_{n \geq 0} (A \times B)^n \rightarrow \Delta(B)$. A *stationary* strategy of the receiver is a function $y : A \rightarrow \Delta(B)$ with the interpretation that the receiver chooses his action according to $y(\cdot | a) \in \Delta(B)$ whenever told $a \in A$.

Our goal is to study to what extent the dynamic structure of the game affects the equilibrium outcomes. Formally, we aim at providing a characterization of the limit set of sequential equilibrium payoffs, and at understanding equilibrium behavior when players are very patient.

3. An example

We here illustrate our main results by means of a simple example. There are two states, $S = \{L, R\}$, and two actions for the receiver, l and r . Successive states are independent and equally likely. Payoffs are given by the two tables in Fig. 1, where c is a fixed parameter with $c \in (1, 2)$. The sender and the receiver are respectively players 1 and 2.

The one-shot information transmission game has a unique equilibrium, in which the receiver plays r with probability 1. To see this, note that the sender strictly prefers action r over action l , no matter what the state is. Thus, at equilibrium, all messages that are sent with positive probability induce the same mixed action by the receiver. This constant mixed action, being always *ex post* optimal for the receiver, is therefore also *ex ante* optimal. It must thus assign probability one to action r .

All equilibria in the one-shot game are therefore babbling equilibria.³ Plainly, the dynamic game admits a babbling equilibrium, in which the sender repeatedly makes the same announcement, the receiver treats the announcements as being non-informative, and plays r in every stage. Note that the receiver can always choose to ignore the announcements of the sender and to play r in every stage, thereby getting 1. As a result, the babbling equilibrium is the *worst* equilibrium for the receiver, in both the one-shot game and the dynamic game.

We claim that the dynamic game has equilibrium payoffs that are arbitrarily close to $(\frac{2+c}{2}, \frac{3}{2})$, which is the best possible outcome for the receiver. In particular, and in contrast with the receiver, there are equilibrium payoffs for the sender that are below the babbling equilibrium payoff. Here is the intuition. The sender announces the true state at every stage. The receiver listens to the announcements of the sender, and plays l when told L , and r when told R . To prevent the sender from announcing R in every stage, the receiver monitors the announcements of the sender, and

² That is, for any two states $s, t \in S$, and for every $N \in \mathbf{N}$ large enough, the probability of moving from s to t in exactly N stages is positive.

³ In the sense that the action of the receiver is independent of the message sent by the sender.

stops listening if there is an obvious bias (towards either L or R). Under the constraint that he should announce both states equally often, the expected payoff of the sender is highest when he reports truthfully.

While this intuition is simple, formalizing it into an equilibrium of the discounted game is not straightforward. Indeed, because payoffs are discounted, the sender may have a preference to send at first the message R more frequently.

We start with a simple construction that yields an equilibrium payoff distinct from $(2, 1)$. Assume that the discount factor satisfies $\delta > \frac{4-2c}{3-c}$, and consider the following strategy profile.

- At odd stages, the sender announces truthfully the current stage, and the receiver plays l if told L , and r if told R .
- At even stages, the sender announces a constant message, and the receiver plays the action that he did *not* play in the previous stage.
- If the receiver deviates, both players switch to the babbling equilibrium forever.

Under this strategy profile, expected payoffs are equal to $\frac{1}{1+\delta}(\frac{2+c}{2} + \delta\frac{5+c}{4})$ for the sender and $\frac{1}{1+\delta}(\frac{3}{2} + \frac{3}{4}\delta)$ for the receiver. Because a deviation of the receiver is followed by the babbling equilibrium, which yields 1 to the receiver, and because the (conditional) expected payoff of the receiver is at least 1 in every stage, no deviation of the receiver is profitable. Regarding the sender, it is sufficient to show that he cannot profit by deviating in any block of two stages. In such a block the sender has two possible deviations: to announce L in the first stage of the block when the true state is R , and to announce R in the first stage of the block when the true state is L . In the former case, he gets 1 at the first stage and 2 in the second (instead of 2 at the first stage and $\frac{1+c}{2}$ at the second stage if he announces truthfully). In the latter case, he gets 2 at the first stage and $\frac{1+c}{2}$ at the second stage (instead of c at the first stage and 2 at the second stage if he announces truthfully). The choice of δ ensures that none of these deviations is profitable.

To get payoffs closer to $(\frac{2+c}{2}, \frac{3}{2})$, we will be relying on a slightly more complex construction. We let the size $2N$ of a block be large enough, so that a law of large numbers will apply. Once N is fixed, we let the discount factor δ be high enough, so that the contribution of any individual block to the overall discounted payoff is very small.

We first describe a pure strategy τ of the receiver. In each block (unless if the receiver has deviated earlier), the receiver listens to the sender's announcements, plays l if told L , r if told R , until the number of announcements of either L or R exceeds N . When this is the case, the receiver stops listening to the sender's announcements, and repeats the least frequent action until the end of the current block.⁴ In a sense, the sender is restricted to announcing both states equally often in any given block of $2N$ stages. As such, the intuition here is similar to some extent to the one behind the *linking mechanism* of [17] and, even more, to the analysis in [10].⁵

If indeed the sender reports truthfully the current state, there is a high probability that the receiver will be listening to the sender most of the time, and the expected payoff is therefore close to $(\frac{2+c}{2}, \frac{3}{2})$.

In contrast with the situation examined above, it need not be optimal for the sender to report truthfully when facing τ . However, a crucial insight is that *any* best reply of the sender to τ must

⁴ An alternative construction, that we adopt in the general case, is for the receiver to generate a specific sequence of fictitious announcements, and continue as if the sender's announcements were equal to the fictitious ones.

⁵ The present analysis and the one in [10] were developed independently.

be reporting truthfully most of the time, with high probability. To see why, observe that any best reply achieves a payoff of at least, say, $\frac{2+c}{2} - \varepsilon$. But since the receiver plays both actions l and r equally likely on each block, this implies that with high probability the action of the receiver matches the state, most of the time.

We let σ be any pure best reply of the sender to τ . On the equilibrium path, we let players play according to σ and τ . By construction, the equilibrium property holds for the sender. To deter the receiver from deviating, both players switch forever to the babbling equilibrium once a deviation of the receiver is detected. Since the discount factor is high, the expected continuation payoff of the receiver is close to $\frac{3}{2}$ following any history, while the receiver gets a payoff of 1 (or close to 1) if he deviates.

4. Main results

We here state and discuss two results on the limit set of equilibrium payoffs. Loosely speaking, according to [Theorem 1](#) all payoff vectors that are individually rational for the receiver and incentive compatible for the sender are (asymptotically) equilibrium payoffs. [Theorem 2](#) proves the converse inclusion. Further results are provided in [Section 6](#).

4.1. Theorem 1

We start with some notations. We denote by $\mathcal{M} \subset \Delta(S \times A)$ the set of *copulas* based on m ; that is, the set of distributions μ over $S \times A$ whose marginals on S and on A are both equal⁶ to m . The set \mathcal{M} is defined by a finite number of linear inequalities, hence it is a compact convex polyhedron, and in particular it has finitely many extreme points.

We denote by $\mu_0 \in \mathcal{M}$ the specific distribution defined as $\mu_0(s, s) = m(s)$ for each $s \in S$, and $\mu_0(s, a) = 0$ if $s \neq a$. Under μ_0 , the messages and the states coincide *a.s.* Thus, the distribution μ_0 is the long-run average distribution of the sequence $(s_n, a_n)_n$ when the sender reports truthfully the current state.

Given a copula $\mu \in \mathcal{M}$ and a stationary strategy $y : A \rightarrow \Delta(B)$ we set

$$U(\mu, y) := \sum_{s \in S, a \in A} \mu(s, a) u(s, y(\cdot | a)) \in \mathbf{R}^2,$$

where $u(s, y(\cdot | a)) = \sum_{b \in B} u(s, b) y(b | a)$. This is the expected payoff vector when the sender's report is drawn according to $\mu(\cdot | s)$, and the receiver plays y .

We denote by

$$v^2 := \max_{b \in B} \sum_{s \in S} m(s) u^2(s, b) \quad (1)$$

the babbling equilibrium payoff for the receiver.

Definition 1. Let $E(\mathcal{M})$ denote the set of payoff vectors $U(\mu_0, y)$, where $y : A \rightarrow \Delta(B)$, that satisfy

- C1. $U^1(\mu_0, y) \geq U^1(\mu, y)$ for every $\mu \in \mathcal{M}$, and
 C2. $U^2(\mu_0, y) \geq v^2$.

⁶ Recall that the set A of messages is a copy of the set S of states.

Let $\widehat{E}(\mathcal{M})$ be the set of payoff vectors $U(\mu_0, y) \in E(\mathcal{M})$ where the inequalities in C1 and C2 are strict. That is, $\widehat{E}(\mathcal{M})$ is the set of vectors $U(\mu_0, y), y : A \rightarrow \Delta(B)$, such that

- D1. $U^1(\mu_0, y) > U^1(\mu, y)$ for every $\mu \in \mathcal{M}, \mu \neq \mu_0$, and
 D2. $U^2(\mu_0, y) > v^2$.

Note that condition D1 holds as soon as the inequality $U^1(\mu_0, y) > U^1(\mu, y)$ is satisfied for each of the finitely many extreme points $\mu \neq \mu_0$ of \mathcal{M} .

We denote by SE_δ the set of sequential equilibrium payoffs of the game with discount factor δ . Our first main result, [Theorem 1](#), shows that all payoffs in $E(\mathcal{M})$ can be obtained as equilibrium payoffs, provided that players are sufficiently patient.

Theorem 1. *Suppose that there exists a public randomizing device, which outputs a (uniformly distributed) number in $[0, 1]$ in every stage, after the announcement of the sender. If $\widehat{E}(\mathcal{M}) \neq \emptyset$ then*

$$E(\mathcal{M}) \subseteq \liminf_{\delta \rightarrow 1} SE_\delta.$$

[Theorem 1](#) means that for every $\gamma \in E(\mathcal{M})$ and for every $\varepsilon > 0$, there exists $\delta_0 < 1$ such that, for every $\delta \geq \delta_0$, the δ -discounted game has a sequential equilibrium payoff within ε of γ . The proof⁷ of [Theorem 1](#) is provided in [Section 5.1](#).

Few comments are in order. The babbling payoff v^2 is equal to the minmax value of the receiver in the dynamic game. Hence condition C2 in [Definition 1](#) reads as an individual rationality condition. Condition C1 is akin to an incentive compatibility condition: under the constraint that the distribution of messages is equal to the distribution of states, truth-telling is optimal for the sender. According to [Theorem 1](#), any payoff vector $U(\mu_0, y)$ that is incentive compatible for the sender and individually rational for the receiver is an equilibrium payoff for δ large.

Our construction will have the somewhat surprising feature that the sender reports truthfully, at least most of the time and with high probability. The intuition is reminiscent of the revelation principle in mechanism design. Let an equilibrium (σ, τ) be given. Consider the strategy profile where the sender reports truthfully, and the receiver first computes the message that the strategy σ would have sent, and next plays what τ would have played given this message. We argue loosely that this new profile (when supplemented with threats) is an equilibrium. The key to the argument is twofold. On the one hand, the sender can check that the receiver does indeed play as prescribed, and does not use the additional information provided by the knowledge of the true state. On the other hand, the threat of switching to the babbling equilibrium is effective because the knowledge of the state at a given stage becomes eventually valueless in predicting distant stages, because of the irreducibility property of the sequence of states. However, we stress that no revelation principle applies in our setup and our equilibrium construction relies on the threat that the sender will stop providing information following a deviation.

[Theorem 1](#) relies on two assumptions. The public randomizing device can easily be dispensed with, provided one slightly extends the communication options offered to the players. To be

⁷ We will actually prove the stronger statement that δ_0 can be chosen to be independent of γ :

$$\lim_{\delta \rightarrow 1} \sup_{\gamma \in E(\mathcal{M})} d(\gamma, SE_\delta) = 0.$$

specific, assume that the players are allowed to exchange simultaneous ‘messages’, after the sender has reported a state. Under such an assumption, players can implement *jointly controlled lotteries* as in [3], which can substitute the randomization device. Details are standard and omitted. However, when instead communication is restricted to a single message sent by the sender, then the existence of a public randomizing device is not without loss of generality; see Section 6.2 for an example.

Theorem 1 also requires $\widehat{E}(\mathcal{M})$ to be non-empty. This is similar to the non-empty interior type conditions which appear in Folk Theorems. Yet, we stress that our assumption is somewhat stronger, since $\widehat{E}(\mathcal{M})$ need not be equal to the relative interior of $E(\mathcal{M})$, and the condition that $\widehat{E}(\mathcal{M}) \neq \emptyset$ is not generically satisfied. We provide a robust example where $\widehat{E}(\mathcal{M}) = \emptyset$ and elaborate further on this issue in Section 6.1.

It is not an easy task to rely on **Definition 1** to check whether a given payoff vector $U(\mu, y)$ belongs to $E(\mathcal{M})$. Fortunately, it turns out that conditions C1 and D1 are equivalent to much simpler conditions.

Lemma 1. *Let $y : A \rightarrow \Delta(B)$ be given. Conditions C1 and D1 are respectively equivalent to conditions C'1 and D'1 below.*

- C'1. $\sum_{s \in S} u^1(s, y(\cdot | s)) \geq \sum_{s \in S} u^1(s, y(\cdot | \phi(s)))$ for every permutation ϕ over S .
 D'1. $\sum_{s \in S} u^1(s, y(\cdot | s)) > \sum_{s \in S} u^1(s, y(\cdot | \phi(s)))$ for every permutation ϕ over S that is not the identity mapping.

The proof of **Lemma 1** appears in **Appendix A**. Interestingly, conditions C'1 and D'1 do *not* involve the invariant distribution m . The intuition is best explained in the case of two states, s_0 and s_1 . Assume that the sender is considering mis-reporting the state, under the constraint that the distribution of reports matches the invariant distribution m of the state. The only way to do this is to report s_1 instead of s_0 , *as often* as to report s_0 instead of s_1 . Whether such a deviation is profitable is equivalent to asking how the *unweighted* sum of the payoffs obtained in s_0 when reporting s_1 and in s_1 when reporting s_0 compares to the *unweighted* sum of the payoffs obtained in the two states when reporting truthfully.

4.2. Theorem 2

Our second main result provides the converse inclusion to that in **Theorem 1**. It requires one substantive assumption on the behavior of the state, **Assumption A** below.

Assumption A. There exist non-negative numbers $(\alpha_s)_{s \in S}$, with $\sum_{s \in S \setminus \{\bar{s}\}} \alpha_s \leq 1$ (for every $\bar{s} \in S$), such that $p(s' | s) = \alpha_{s'}$ whenever $s' \neq s$.

Assumption A is restrictive, yet it holds, e.g., in the following cases.

Assume first that changes in the state are due to shocks, which occur at random times. Once drawn, the state remains constant until a shock occurs. The state is then drawn anew, according to m . The inter-arrival times of the successive shocks are i.i.d. and follow a geometric distribution. In that case, **Assumption A** is met. Indeed, it suffices to set $\alpha_t = \pi \times m(t)$ for every $t \in S$, where π is the per-stage probability of a shock. The parameter π is here a measure of the state persistence. When π increases from 0 to 1, the situation evolves from one in which the state remains constant through time, to a situation in which successive states are independent.

When $\pi = 1$, the successive states are independent, and identically distributed according to m . Thus, **Assumption A** holds in the case of i.i.d. states.

Assumption A also holds in the benchmark case where there are only two possible states. Indeed, denoting the two states by s_1 and s_2 , it suffices to set $\alpha_1 = p(s_1 | s_2)$ and $\alpha_2 = p(s_2 | s_1)$. In particular, it is satisfied in the models in [1,21,24].

As a further simple illustration, consider a symmetric random walk on three states. That is, whenever in a state, the chain moves to each of the two other states with probability $\frac{1}{2}$. Again, **Assumption A** is met, with $\alpha_s = \frac{1}{2}$ for each $s \in S$.

Note that **Assumption A** implies that $m(s) = \alpha_s (\sum_{t \in S} \alpha_t)$ for every $s \in S$. This implies in particular that the Markov chain is reversible: $m(s)p(s' | s) = m(s')p(s | s')$ for all states $s, s' \in S$, and therefore $\mathbf{P}(s_n = s, s_{n+1} = s') = \mathbf{P}(s_n = s', s_{n+1} = s)$ for each stage $n \in \mathbf{N}$.

We denote by NE_δ the set of (Nash) equilibrium payoffs in the game with discount factor δ .

Theorem 2. *Suppose that **Assumption A** holds. Then, for every $\delta < 1$, one has*

$$NE_\delta \subseteq E(\mathcal{M}).$$

Provided that $\widehat{E}(\mathcal{M}) \neq \emptyset$ and that **Assumption A** is met, **Theorems 1 and 2** thus imply that the set of sequential equilibrium payoffs SE_δ converges to the set $E(\mathcal{M})$ (as soon as a randomizing device is available).

Note that the set \mathcal{M} of copulas only depends on the invariant measure m , and not on finer details of the transition function. A striking implication of the characterization is that, under **Assumption A**, the limit set of equilibrium payoffs only depends on the invariant measure m . In particular, the limit set of equilibrium payoffs is the same as when the states are drawn independently across stages. That is, the amount of state persistence is irrelevant for the determination of the limit set of equilibrium payoffs.

If the initial state s were to remain fixed throughout the play, the game would fall into the class of repeated games with incomplete information introduced by [3]. (This is the setup studied in [13].) In this case, the limit set of discounted equilibrium payoffs, when δ goes to 1, is typically not equal to $E(\mathcal{M})$. Hence, there is a discontinuity in the limit set of equilibrium payoffs when successive states become perfectly autocorrelated.⁸

By contrast, for a fixed discount factor, the set of equilibrium payoffs is upper hemi-continuous with respect to the transition function. The source of this apparent paradox can be traced back to the fact that, in loose terms, the convergence of the set SE_δ to $E(\mathcal{M})$ is slower, the more correlated successive states are.

The main insight to be derived from **Theorem 2** is the following. The incentive compatibility condition C2 is a very *strong* one. Indeed, it only requires from deviations that the distribution of announcements matches the invariant measure. In particular, according to **Theorem 2**, all equilibria are payoff-equivalent to equilibria in which the receiver only checks that the announcements frequencies are consistent with m . Yet, much more sophisticated tests are available to the receiver. The receiver might, e.g., check that the empirical distribution of two-letter words (s, s') matches the transition function p , as in [10], or look at the distribution of three-letter words, etc. This might potentially allow the receiver to impose *weaker* incentive constraints than the one in C1, and therefore, allow for equilibrium payoffs outside of $E(\mathcal{M})$. **Theorem 2** thus identifies one class of Markov chains for which this is not the case.

⁸ Our **Theorems 1 and 2** extend to cover the case of uniform equilibrium payoffs.

	$b = 0$	$b = 1$	$b = 2$	$b = 3$	$b = 4$
$s = 0$	1, 1	$c, 0$	0, 0	0, 0	0, 0
$s = 1$	$c, 0$	1, 1	0, 0	0, 0	0, 0
$s = 2$	0, 0	0, 0	1, 1	0, 0	0, 0
$s = 3$	0, 0	0, 0	0, 0	1, 1	0, 0
$s = 4$	0, 0	0, 0	0, 0	0, 0	1, 1

Fig. 2. The game in Example 1.

We provide below an example where the conclusion of Theorem 2 fails to hold if Assumption A is not satisfied. Thus, in general, the limit set of sequential equilibrium payoffs does not only depend on the invariant measure, but also on finer details of the transition function.

Example 1. Consider a game with 5 states $S := \{0, 1, 2, 3, 4\}$. The sequence of states follows a random walk on S . When in s , the chain moves either to $s + 1 \pmod{5}$ or to $s - 1 \pmod{5}$ with equal probabilities. The action set B of the receiver coincides with S , and the payoff function is described in Fig. 2, where $c > 1$.

Thus, both players receive a payoff 1 if the action matches the current state, and 0 otherwise, except when the receiver chooses action 1 in state 0, or action 0 in state 1.

The payoff vector $(1, 1)$ is not in $E(\mathcal{M})$ as soon as $c > 1$. Indeed, the stationary strategy $y : A \rightarrow \Delta(B)$ defined by $y(s | s) = 1$ is the only strategy such that $U(\mu_0, y) = (1, 1)$. But then, the sender profits by reporting $t = 1$ whenever $s = 0$, and $t = 0$ whenever $s = 1$. On the other hand, $(1, 1)$ is an equilibrium payoff, as soon as $c < \frac{3}{2}$, provided the players are patient enough. Indeed, consider the strategy of the receiver in which he matches the announcement of the sender, as long as $|a_{n+1} - a_n| = 1$ modulo 5, and switches forever to the babbling equilibrium (e.g., playing always $b = 4$) if $|a_{n+1} - a_n| \neq 1$ modulo 5 for some stage n . Provided c is not too large, the best response of the sender is to report the true state. If instead, say, the sender chooses to report $t = 1$ when in fact $s = 0$ in a given stage, he gains $c - 1$, but then in the next period, with probability $\frac{1}{2}$ the new state will be $s = 4$, and then he will either report $t \in \{0, 2\}$ and receive 0, or report $t \in \{1, 3, 4\}$ and be punished with the babbling equilibrium payoff $\frac{1}{5}$ forever. Provided the players are patient enough, such a deviation is not profitable.

In this equilibrium, the receiver checks that the one-step transitions between successive announcements are consistent with the transitions of the Markov chain. As it turns out, under Assumption A such a sophisticated statistical analysis of the announcements is not more powerful than a statistical analysis which is based only on the empirical frequencies of the different announcements. ♦

Theorems 1 and 2 hold as soon as the sender knows the current state. As will be clear from the proof, they continue to hold if the sender knows more. In particular, they hold in the extreme case where the sender learns the entire sequence of realized states in stage 1, or in any intermediate setup.

Note that we interpret the case $\delta \rightarrow 1$ as players being very patient. It is not possible here to interpret it as a situation in which players would interact more frequently. Indeed, a proper analysis of this latter case would take into account the impact on transitions: when players interact more frequently, states become more persistent between successive interactions.

	$b = L$	$b = M$	$b = R$
$s = s_0$	1, 1	0, 0	0, 0
$s = s_1$	0, 0	1, 0	0, 1
$s = s_2$	0, 0	0, 1	1, 0

Fig. 3. The game in the example in Section 4.3.

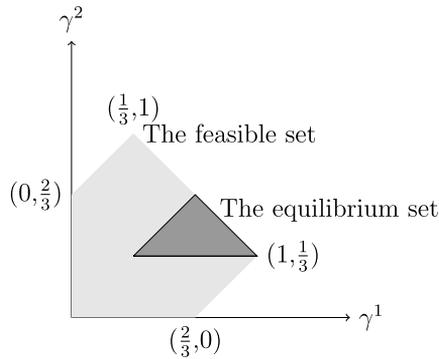


Fig. 4. The equilibrium set in the example in Section 4.3.

4.3. An example

We here analyze a simple, specific example to show how to pin down the set of equilibrium payoffs using our results. We let the set of states be $S = \{s_0, s_1, s_2\}$. Between any two stages, the state changes with probability one, and each of the two possible states is equally likely. Thus, $p(t | s) = \frac{1}{2}$ for every $s \neq t \in S$. Assumption A on the transition function holds and the invariant measure assigns probability $\frac{1}{3}$ to each state.

The receiver has three actions, denoted L, M and R , and the payoffs in the different states are given by the matrix; see Fig. 3.

All extreme points of the feasible set are obtained by having the sender report truthfully the state, and the receiver then play a pure state-dependent action. Thus, all extreme points are obtained by picking one entry in each row, and averaging. For instance, picking L (resp. M, R) in row s_0 (resp. s_1, s_2) and averaging over states leads to a payoff of $(1, \frac{1}{3})$. One checks that the feasible set is the convex hull of the five payoffs $(0, 0), (0, \frac{2}{3}), (\frac{1}{3}, 1), (\frac{2}{3}, 0)$ and $(1, \frac{1}{3})$; see Fig. 4.

Without any information on the state, all three actions of the receiver yield $\frac{1}{3}$, hence $v^2 = \frac{1}{3}$.

Let $\gamma = U(\mu_0, y) \in E(\mathcal{M})$ be a limit equilibrium payoff, and denote by μ the copula obtained when the sender exchanges the two states s_1 and s_2 when reporting. Thus, $\mu(s_1, s_2) = \mu(s_2, s_1) = \mu(s_0, s_0) = \frac{1}{3}$, and $\mu(s, t) = 0$ otherwise. Since the payoffs of the two players are exchanged in the two states s_1 and s_2 , one has $U^1(\mu, y) = U^2(\mu_0, y)$ and $U^1(\mu_0, y) = U^2(\mu, y)$. The incentive condition $U^1(\mu_0, y) \geq U^1(\mu, y)$ thus yields $\gamma^1 \geq \gamma^2$.

Note finally that the sum of the players' payoffs cannot exceed 2 in state s_0 , and 1 in states s_1 and s_2 . Thus, $\gamma^1 + \gamma^2 \leq \frac{1}{3}(2 + 1 + 1) = \frac{4}{3}$.

Hence, any equilibrium payoff (γ^1, γ^2) lies in the shaded triangle defined by the inequalities $\gamma^2 \geq \frac{1}{3}, \gamma^1 + \gamma^2 \leq \frac{4}{3}, \gamma^1 \geq \gamma^2$; see Fig. 4.

On the other hand, each of the extreme points of this triangle is an equilibrium payoff. Indeed, $(\frac{1}{3}, \frac{1}{3})$ is the babbling equilibrium payoff, while $(\frac{2}{3}, \frac{2}{3}) = U(\mu_0, y_1)$ and $(1, \frac{1}{3}) = U(\mu_0, y_2)$, where $y_1 : A \rightarrow \Delta(B)$ plays L when told s_0 , and randomizes between M and R otherwise, while y_2 plays L, M and R in states s_0, s_1 and s_2 respectively. As a result, the set of equilibrium payoffs is equal to the shaded triangle in Fig. 4.

5. Proofs

5.1. Proof of Theorem 1

We here provide most details of the proof of Theorem 1. Some technical details appear in Appendix B. Since $\widehat{E}(\mathcal{M}) \neq \emptyset$, there is $y_0 : A \rightarrow \Delta(B)$ such that $U^2(\mu_0, y_0) > v^2$ and $U^1(\mu_0, y_0) > U^1(\mu, y_0)$ for every $\mu \in \mathcal{M}$, $\mu \neq \mu_0$.

Let $\varepsilon > 0$ and $U(\mu_0, \bar{y}) \in E(\mathcal{M})$ be arbitrary, and define $y := \varepsilon y_0 + (1 - \varepsilon)\bar{y}$. It is sufficient to prove that $U(\mu_0, y)$ is arbitrarily close to some sequential equilibrium payoff of the δ -discounted game, provided δ is high enough.

5.1.1. The strategies

Let some integer $N \in \mathbf{N}$ and a discount factor δ be given. We here define a strategy profile (σ_*, τ_*) . According to (σ_*, τ_*) , the play is divided into consecutive blocks of N stages. At the beginning of each block, players discard past information and re-start playing an N -stage profile (σ_0, τ_0) , where τ_0 is a pure strategy. In case the receiver deviates from the pure strategy τ_0 , the players switch to babbling play forever.

We now construct σ_0 and τ_0 , starting with τ_0 . Consider any block of N stages. According to τ_0 , the receiver “listens” to the reported state a_n in stage n and plays $y(\cdot | a_n)$, as long as no state has been reported too often. As soon as this fails to be the case, the receiver substitutes the actual report of the sender by some fictitious report θ_n , and plays according to $y(\cdot | \theta_n)$. To be formal, we pick a distribution $m_N \in \Delta(S)$ that best approximates the invariant measure m among all distributions $\bar{m} \in \Delta(S)$ such that $N\bar{m}(s)$ is an integer for all s .

For $s \in S$ and $n \in \mathbf{N}$ we denote by $\mathbf{N}_n(s) = |\{k \leq n : a_k = s\}|$ the number of stages where the sender reported state s , and we set

$$q := \min\{1 \leq n \leq N : \mathbf{N}_n(a_n) > Nm_N(a_n)\}$$

($\min \emptyset = +\infty$). Intuitively, each state $s \in S$ is allotted a quota of announcements equal to $Nm_N(s)$. The stage q is the first stage in which at least one of the quotas is no longer met. From stage q until the end of the block the receiver substitutes actual reports by fictitious ones.

Formally, we let (θ_n) ($n = 1, \dots, N$) be a sequence such that

- F1. $\theta_n = a_n$ for $n < q$;
- F2. For each $s \in S$, the equality $|\{n \leq N : \theta_n = s\}| = Nm_N(s)$ always hold;
- F3. Conditional on (a_1, \dots, a_q) , the variables $(\theta_q, \dots, \theta_N)$ are deterministic.

We will refer to θ_n as the *announcement* at stage n . Condition F1 means that the announcements coincide with the sender’s actual announcements prior to stage q ; condition F2 ensures that the entire sequence of announcements always satisfies the quotas; condition F3 ensures in particular that the fictitious announcements are commonly known between the two players.

Thanks to the public randomizing device, the strategy τ_0 may be rewritten as a pure strategy.⁹

The strategy σ_0 is defined to be any pure best-reply strategy of the sender to τ_0 in the N -stage δ -discounted game starting in stage 1. Note that the strategy σ_0 is also a best reply to τ_0 on any of the consecutive blocks of N stages, conditional on past play.¹⁰ In particular, σ_* is a best reply to τ_* .

5.1.2. Equilibrium properties

We here argue that for appropriate choices of N and of δ , the strategy τ_* is a best reply to σ_* , and that (σ_*, τ_*) induces a payoff arbitrarily close¹¹ to $U(\mu_0, y)$.

Proposition 1. For every $\eta > 0$ there exists $N_0 \in \mathbf{N}$ such that the following holds: for every $N \geq N_0$, there is $\delta_0 < 1$ such that for every $\delta \geq \delta_0$ the profile (σ_*, τ_*) is a sequential equilibrium that induces a payoff within η of $U(\mu_0, y)$.

The complete proof of Proposition 1 is in Appendix B.1. The crucial step consists in showing that the fact that σ_0 is a best reply to τ_0 implies that with high probability σ_0 reports the true state in most stages. This is the content of Lemma 2 below. In the statement of the lemma, μ_{σ_0, τ_0} is the (expected, undiscounted) joint distribution of states and reports in a block of N stages. That is, for each $(s, a) \in S \times S$,

$$\mu_{\sigma_0, \tau_0}(s, a) := \mathbf{E}_{\sigma_0, \tau_0} \left[\frac{1}{N} \sum_{n=1}^N \mathbf{1}_{\{s_n=s, \theta_n=a\}} \right]$$

is the expected frequency of the pair (s, a) over N stages (recall that the distribution of the initial state, s_1 , is the invariant measure).

Lemma 2. For every $\eta > 0$ there is $N_0 \in \mathbf{N}$ such that the following holds: for every $N \geq N_0$ there is $\delta_0 < 1$ such that for every $\delta \geq \delta_0$ one has

$$\|\mu_{\sigma_0, \tau_0} - \mu_0\| < \eta.$$

We will provide insights into the proof of Lemma 2 below. For the time being, we show how to deduce Proposition 1 from Lemma 2. Observe first that by the definition of μ_{σ_0, τ_0} the expected average¹² payoff induced by (σ_0, τ_0) over a single block is equal to $U(\mu_{\sigma_0, \tau_0}, y)$. For fixed N , since the profile (σ_*, τ_*) consists of periodic repetitions of (σ_0, τ_0) , the discounted payoff induced by (σ_0, τ_0) converges to $U(\mu_{\sigma_0, \tau_0}, y)$ as $\delta \rightarrow 1$. In particular, it is arbitrarily close to the target payoff $U(\mu_0, y)$.

We now argue that τ_* is a best reply to σ_* . Following any history, the continuation payoff of the receiver is equal to the sum of his payoffs until the end of the current block and of the continuation payoff from the next block on. The latter is equal to the discounted payoff induced

⁹ Indeed, denote by $X_n \sim \mathcal{U}([0, 1])$ the output of the public device in stage n and label the receiver’s actions from 1 to $|B|$. We let the strategy τ_0 instruct the receiver to choose the action $b \in B$ whenever $\sum_{i=1}^{b-1} y(i | \theta_n) \leq X_n < \sum_{i=1}^b y(i | \theta_n)$. Thus, the device performs publicly the desired randomization.

¹⁰ This observation relies on the fact that, in the first block, σ_0 is a best reply to τ_0 , no matter what the distribution of the initial state s_1 is.

¹¹ Off-equilibrium-path beliefs and sequential rationality issues are discussed in Appendix B.

¹² That is, when payoffs in the different stages are not discounted.

by (σ_*, τ_*) , computed using the belief held by the receiver at the beginning of this block. For a fixed $N \in \mathbf{N}$, this continuation payoff thus converges¹³ to $U^2(\mu_{\sigma_0, \tau_0}, y)$ as $\delta \rightarrow 1$.

On the other hand, any deviation from τ_0 , say in stage n , triggers a babbling play, and the receiver’s continuation payoff therefore does not exceed $(1 - \delta) \sum_{k=n}^{\infty} \delta^{k-n} \max_{b \in B} u^2(p_k, b)$, where p_k is the belief that the receiver will hold at stage $k \geq n$ on the state s_k at that stage. Since the sequence of states forms an irreducible and aperiodic chain, p_k converges to m . For a fixed $N \in \mathbf{N}$, this continuation payoff converges to v^2 as $\delta \rightarrow 1$ (again, uniformly over all histories).

Since $U^2(\mu_{\sigma_0, \tau_0}, y) > v^2$, this proves the best-reply property of τ_0 , provided first N , and then δ , are chosen large enough.

We now turn to Lemma 2. We denote by σ_{truth} the strategy of the sender that announces truthfully the current state, no matter what. The proof of Lemma 2 combines several ideas.

First, by a law of large numbers for Markov chains, if N is large enough, then there is a high probability that the realized state frequencies will be consistent with the quotas in most stages. Thus, under (σ_{truth}, τ_0) , there is a high probability that the receiver follows the announcement of the sender in most stages. That is, the distribution $\mu_{\sigma_{truth}, \tau_0}$ is arbitrarily close to μ_0 .

Next, for a fixed $N \in \mathbf{N}$, for every (periodic) strategy σ and viewing τ_0 as a periodic strategy, the discounted payoff $\gamma_\delta(\sigma, \tau_0)$ converges to $U(\mu_{\sigma, \tau_0}, y)$ as δ converges to one.

Finally, the best-reply property of σ_0 implies that $\gamma_\delta^1(\sigma_0, \tau_0) \geq \gamma_\delta^1(\sigma_{truth}, \tau_0)$.

Combining these observations, the following formal statement holds. For every $\varepsilon > 0$ there exist N_0 and δ_0 such that for every $\delta \geq \delta_0$ the following sequence of inequalities holds:

$$\begin{aligned}
 U^1(\mu_{\sigma_0, \tau_0}, y) &\geq \gamma_\delta^1(\sigma_0, \tau_0) - \varepsilon \geq \gamma_\delta^1(\sigma_{truth}, \tau_0) - \varepsilon \geq U^1(\mu_{\sigma_{truth}, \tau_0}, y) - 2\varepsilon \\
 &\geq U^1(\mu_0, y) - 3\varepsilon.
 \end{aligned}
 \tag{2}$$

To conclude, we will rely on Lemma 3 below, which critically depends on the assumption that $U(\mu_0, m_0) \in \widehat{E}(\mathcal{M})$. Denote by \mathcal{M}_e the (finite) set of extreme points of \mathcal{M} . Recall that $\mu_0 \in \mathcal{M}_e$. Set

$$c_1 := \min_{\{\mu_e \in \mathcal{M}_e, \mu_e \neq \mu_0\}} (U^1(\mu_0, y_0) - U^1(\mu_e, y_0)) \quad \text{and} \quad c_2 := \max_{\{\mu_e \in \mathcal{M}_e, \mu_e \neq \mu_0\}} \|\mu_e - \mu_0\|_1,$$

and note that both c_1 and c_2 are positive.

Lemma 3. For every $\mu \in \mathcal{M}$ one has

$$U^1(\mu_0, y) - U^1(\mu, y) \geq \frac{\varepsilon c_1}{c_2} \|\mu - \mu_0\|_1.$$

Lemma 3 may be paraphrased as saying that any strategy that does approximately as well as the truth-telling strategy must be telling the truth in most stages with high probability.

The conclusion of Lemma 2 follows from Eq. (2) combined with Lemma 3.

5.2. Proof of Theorem 2

We here provide insights into the proof of Theorem 2. We let $\delta < 1$ and we fix a Nash equilibrium (σ, τ) of the δ -discounted game (with or without randomizing device). For clarity, we sometimes use boldfaced letters to denote random variables.

¹³ Uniformly over all histories.

For $s \in S$, we define $y(\cdot | s) \in \Delta(B)$ as the expected *discounted* distribution of moves of the receiver in state s . Formally, for $s \in S$, and $b \in B$, we set

$$y(b | s) = \frac{1}{m(s)} \mathbf{E}_{\sigma, \tau} \left[\sum_{n=1}^{\infty} (1 - \delta) \delta^{n-1} 1_{\{s_n=s, \mathbf{b}_n=b\}} \right] \\ = \frac{\mathbf{E}_{\sigma, \tau} [\sum_{n=1}^{\infty} (1 - \delta) \delta^{n-1} 1_{\{s_n=s, \mathbf{b}_n=b\}}]}{\mathbf{E}_{\sigma, \tau} [\sum_{n=1}^{\infty} (1 - \delta) \delta^{n-1} 1_{\{s_n=s\}}]}$$

By construction, one has $\gamma_{\delta}(\sigma, \tau) = U(\mu_0, y)$. Indeed,

$$\gamma_{\delta}(\sigma, \tau) = (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} \mathbf{E}_{\sigma, \tau} [u(\mathbf{s}_n, \mathbf{b}_n)] \\ = (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} \sum_{s \in S, b \in B} \mathbf{E}_{\sigma, \tau} [1_{\{s_n=s, \mathbf{b}_n=b\}}] u(s, b) \\ = \sum_{s \in S, b \in B} u(s, b) m(s) y(b | s) = U(\mu_0, y),$$

as desired.

Since the distribution of \mathbf{s}_n is equal to m for each stage $n \in \mathbf{N}$, one has $\gamma_{\delta}^2(\sigma, \tau) \geq v^2$, and thus $U^2(\mu_0, y) \geq v^2$, so that C2 holds.

We thus need to prove that $U^1(\mu_0, y) \geq U^1(\mu, y)$ for each $\mu \in \mathcal{M}$. The idea of the proof is rather straightforward, but the formal proof is fraught with many technical complications. Let $\mu \in \mathcal{M}$ be given. We will construct a strategy σ' of the sender such that $\gamma_{\delta}(\sigma', \tau) = U(\mu, y)$, so that the desired inequality will follow from the equilibrium property of (σ, τ) .

The strategy σ' is designed as follows. Along the play, the sender will generate a sequence (\mathbf{t}_n) of *fictitious* states that is statistically indistinguishable from the sequence (\mathbf{s}_n) , and such that the average distribution of the pair $(\mathbf{s}_n, \mathbf{t}_n)$ is given by μ . Given such a sequence, in any stage n the sender will play according to σ after substitute the realized state \mathbf{s}_n by the fictitious state \mathbf{t}_n . Formally, following any history $(s_1, t_1, a_1, b_1, \dots, s_n, t_n)$ consisting of realized and fictitious states, messages, and actions up to stage n , the strategy σ' plays the mixed move $\sigma(t_1, a_1, b_1, \dots, t_n)$ that would have been played by σ had the realized states been t_1, \dots, t_n .

We now give some more details. Since the strategy τ may feature complex statistical tests on the successive announcements, the notion of being statistically indistinguishable has to be interpreted in a restrictive sense.

We prove in [Appendix C](#) the following lemma.

Lemma 4. *Assume [Assumption A](#) and let $\mu \in \mathcal{M}$ be given. There exists an S -valued process¹⁴ $(\mathbf{t}_n)_n$ such that the following hold:*

- P1. *Conditional on \mathbf{s}_n , the vector $(\mathbf{t}_1, \dots, \mathbf{t}_n)$ is independent of the future states $(\mathbf{s}_{n+1}, \mathbf{s}_{n+2}, \dots)$.*
- P2. *The law of the sequence $(\mathbf{t}_n)_n$ is the same as the law of the sequence $(\mathbf{s}_n)_n$.*
- P3. *The law of the pair $(\mathbf{s}_n, \mathbf{t}_n)$ is μ for each stage $n \in \mathbf{N}$.*
- P4. *The conditional law of \mathbf{s}_n given $\mathbf{t}_1, \dots, \mathbf{t}_n$ is $\mu(\cdot | \mathbf{t}_n)$.*

¹⁴ The process $(\mathbf{t}_n)_n$ is possibly defined on a probability space which is an enlargement of the one on which $(\mathbf{s}_n)_n$ is defined.

According to P1, the state \mathbf{t}_n can be computed/simulated using only the information available at stage n : past and current states, and past fictitious states. This is a feasibility requirement that ensures that σ' is well defined. Condition P2 ensures that no statistical test can discriminate between the sequences (\mathbf{s}_n) and (\mathbf{t}_n) . Condition P3 provides the desired coupling between \mathbf{s}_n and \mathbf{t}_n .

We now proceed to show that the expected payoff induced by (σ', τ) is equal to $U(\mu, y)$ as claimed. Below we will denote by s_k, t_k, b_k generic values of the random variables $\mathbf{s}_k, \mathbf{t}_k$ and \mathbf{b}_k , respectively. For any given stage $n \in \mathbf{N}$, the following sequence of equalities holds:

$$\begin{aligned} & \mathbf{E}_{\sigma', \tau} [u(\mathbf{s}_n, \mathbf{b}_n)] \\ &= \sum_{s_n, b_n} \mathbf{P}_{\sigma', \tau}(\mathbf{s}_n = s_n, \mathbf{b}_n = b_n) u(s_n, b_n) \\ &= \sum_{s_n} \sum_{t_1, \dots, t_n} \sum_{b_1, \dots, b_n} \mathbf{P}_{\sigma', \tau}(s_n, t_1, \dots, t_n, b_1, \dots, b_n) u(s_n, b_n) \tag{3} \\ &= \sum_{s_n} \sum_{t_1, \dots, t_n} \sum_{b_1, \dots, b_n} \mathbf{P}_{\sigma', \tau}(s_n | t_1, \dots, t_n, b_1, \dots, b_n) \\ &\quad \times \mathbf{P}_{\sigma', \tau}(t_1, \dots, t_n, b_1, \dots, b_n) u(s_n, b_n) \tag{4} \\ &= \sum_{s_n} \sum_{t_1, \dots, t_n} \sum_{b_1, \dots, b_n} \mathbf{P}(s_n | t_1, \dots, t_n) \mathbf{P}_{\sigma', \tau}(t_1, \dots, t_n, b_1, \dots, b_n) u(s_n, b_n) \tag{5} \\ &= \sum_{s_n} \sum_{t_1, \dots, t_n} \sum_{b_1, \dots, b_n} \mu(s_n | t_n) \mathbf{P}_{\sigma', \tau}(t_1, \dots, t_n, b_1, \dots, b_n) u(s_n, b_n) \tag{6} \\ &= \sum_{s_n, t_n, b_n} \mu(s_n | t_n) \mathbf{P}_{\sigma', \tau}(t_n, b_n) u(s_n, b_n), \tag{7} \end{aligned}$$

where (5) holds because the variables $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ are conditionally independent of s_n given $(\mathbf{t}_1, \dots, \mathbf{t}_n)$, and (6) holds by P4.

Using P2 and from the definition of σ' , the δ -discounted sum of $\mathbf{P}_{\sigma', \tau}(\mathbf{t}_n = t_n, \mathbf{b}_n = b_n)$ is equal to $\mathbf{P}_{\sigma, \tau}(\mathbf{s}_n = t_n, \mathbf{b}_n = b_n)$, which is equal to $\mu(t_n) \times y(b_n | t_n)$. By (7) we now obtain

$$\gamma_\delta(\sigma', \tau) = \sum_{s, t, b} \mu(s | t) \mu(t) y(b | t) u(s, b) = U(\mu, y).$$

6. Further results and comments

6.1. On the condition $\widehat{E}(\mathcal{M}) \neq \emptyset$

In light of existing results for repeated games, it is not surprising that some non-empty interi- ority assumption is needed; see [20] for a survey on the topic.

As the next example illustrates, the conclusion of Theorem 1 fails to hold if $\widehat{E}(\mathcal{M}) = \emptyset$.

Example 2. There are two states and the receiver has two actions. The payoffs in the two states are given by the tables in Fig. 5. We assume that the successive states are independent and that the two states are equally likely.

The strategy which plays r irrespective of the announcement is weakly dominant in the one- shot game, and thus, $v^2 = 1$. Consider now the stationary strategy y defined by $y(l | L) =$

l	r
0.5, 1	1, 1
<i>State L</i>	

l	r
0, 0	1, 1
<i>State R</i>	

Fig. 5. The game in Example 2.

l	r
1, 1	0, 0
<i>State L</i>	

l	r
1, 1	0, 0
<i>State R</i>	

Fig. 6. The game in Example 3.

$y(r | R) = 1$. The payoff vector $U(\mu_0, y) = (\frac{3}{4}, 1)$ is in $E(\mathcal{M})$. However, we claim that $(1, 1)$ is the unique equilibrium payoff, irrespective of δ . Indeed, consider any equilibrium (σ, τ) . Plainly, the equilibrium payoff of the receiver is equal to 1. In particular, with probability 1 the receiver plays r whenever the current state is R . This implies that in every stage and for almost every past history, there is one (possibly history-dependent) message that is assigned positive probability by σ following which the receiver plays r . But then the sender gets a payoff 1 by assigning probability 1 to this specific message. ♦

As we stressed, the statement of Theorem 1 is unsatisfactory in one important respect: while non-empty interior requirements in existing Folk Theorems are generically satisfied, the condition $\widehat{E}(\mathcal{M}) \neq \emptyset$ does not hold generically, as the next example shows.

Example 3. Consider the game depicted in Fig. 6, where there are two states, and the receiver has two actions. Successive states are independent and the two states are equally likely.

Here, $v^2 = 1$, and the stationary strategy y_* which plays l irrespective of the announcement is the only stationary strategy that satisfies C2. Hence, $E(\mathcal{M})$ contains a single payoff vector, $(1, 1)$, and $\widehat{E}(\mathcal{M})$ is empty. When payoffs are slightly perturbed, the strategy y_* remains the only strategy satisfying C2, therefore $\widehat{E}(\mathcal{M}) = \emptyset$ for any such perturbation. ♦

Example 3 suggests that if all strategies y for which $U(\mu_0, y)$ is in $E(\mathcal{M})$ are constant strategies, then the set $\widehat{E}(\mathcal{M})$ is empty, even when payoffs in the game are slightly perturbed. We build on this intuition, and introduce a new condition.

Condition B. There is a non-constant function $y : S \rightarrow \Delta(B)$ such that $U(\mu_0, y) \in E(\mathcal{M})$.

If Condition B is not met, then all equilibrium payoffs are babbling. In Theorem 3 below we fix the transition function of the Markov chain p and identify a game to a point in the space $\mathbb{R}^{2 \times S \times B}$ of payoff functions.

Theorem 3. Let a game G be given.

If Condition B holds for G , then any neighborhood of G contains a game G' with $\widehat{E}_{G'}(\mathcal{M}) \neq \emptyset$.

If Condition B does not hold for G , there is a neighborhood \mathcal{N} of G such that for every game in \mathcal{N} Condition B does not hold.

l	m	r
3, 0	0, 4	2, 1
State L		

l	m	r
1, -5	4, -4	2, 1
State R		

Fig. 7. The payoffs of the players in the example in Section 6.2.

Theorem 3 allows us to complete the picture provided by **Theorems 1 and 2**, provided the underlying Markov chain satisfies **Assumption A**. Indeed, let G be a game. If **Condition B** holds for the game G , **Theorems 1 and 2** provide a characterization of the limit set of equilibrium payoffs for games arbitrarily close to G . If **Condition B** does not hold for the game G , then all games close enough to G have only babbling equilibrium payoffs.

6.2. On the role of the randomizing device

The randomizing device is not needed in the proof of **Theorem 2** to implement payoffs $U(\mu_0, y)$, whenever $y(\cdot | s)$ is a pure strategy: it assigns probability 1 to some action $b(s)$, for each $s \in S$. However, as soon as $y(\cdot | s)$ is a truly mixed distribution for some state s , it may be impossible to dispense with the randomizing device, as we now argue by means of an example.

Let there be two states, L and R . The successive states are drawn independently in every period, and each of the two states is equally likely. The receiver has three actions, denoted $B = \{l, m, r\}$. The payoffs are given in **Fig. 7**.

One can verify that $v^2 = 1$. Define y_* to be the stationary strategy such that $y_*(\cdot | R)$ assigns probability 1 to r , and $y_*(\cdot | L)$ assigns probabilities $\frac{2}{3}$ and $\frac{1}{3}$ to l and m , respectively. Then $U(\mu_0, y) = (2, \frac{7}{6})$, and one can verify that $U(\mu_0, y_*) \in E(\mathcal{M})$ while $\widehat{E}(\mathcal{M}) \neq \emptyset$. Thus, using **Theorem 1**, the vector $(2, \frac{7}{6})$ can be approximated by sequential equilibrium payoffs, when players are sufficiently patient, provided a randomizing device is available.

We now assume that such a device is not available. Since successive states are independent, the dynamic game can be viewed as a infinite repetition of the one-shot information transmission game. With this interpretation, an *action* of the sender in the one-shot game is a function $x : S \rightarrow A$, while an action of the receiver is a function $y : A \rightarrow B$. Given an action profile (x, y) , payoffs are random and take the value $u(s, y(x(s)))$ with probability $m(s)$, for $s \in S$. Players then receive the public signal $(x(s), y(x(s)))$.

We will rely on [12]’s characterization of the limit set of perfect public equilibrium (PPE) payoffs in repeated games with public signals. Some care is needed, as there are two dimensions according to which our repeated game does not fit into their setup. First, they assume that a player’s payoff depends deterministically on his own action and on the public signal, while payoffs here depend randomly on the entire action profile (x, y) . Second, their result is a characterization of public equilibrium payoffs, while we focus on sequential equilibrium payoffs.

We briefly argue that nevertheless their result applies to our setting. On the one hand, their result is still valid for games where payoffs depend on the entire action profile.¹⁵ Next, it can be verified that the auxiliary game in which stage payoffs are defined to be the *expected* stage payoffs in our game (given the action profile) has the same set of PPE payoffs. Thus, their result provides a characterization of the limit set of PPE payoffs for our game. On the other hand, let (σ, τ) be a sequential equilibrium of our game, and define a public strategy profile $(\bar{\sigma}, \bar{\tau})$

¹⁵ This can be seen from their proof or, alternatively, deduced from [16].

as follows. Let any public history \bar{h} be given. At \bar{h} we let $\bar{\sigma}$ play the expectation of the mixed move played by σ , where the expectation is computed w.r.t. the belief held by the receiver at the information set which contains \bar{h} . We define $\bar{\tau}(h)$ by exchanging the roles of the two players. It can be verified that $(\bar{\sigma}, \bar{\tau})$ is a public perfect equilibrium of the repeated game.

Fudenberg and Levine in [12] showed¹⁶ that $\gamma \in \mathbf{R}^2$ is a limit PPE payoff if and only if for all $\lambda \in \mathbf{R}^2$ we have $\lambda \cdot \gamma \leq k(\lambda)$, where $k(\lambda)$ is the solution to a certain optimization problem $\mathcal{P}(\lambda)$.

We set $\gamma = (2, \frac{7}{6})$ and we will show that it is not a PPE Payoff using the condition of [12] with $\lambda_* = (0, 1)$. We now recall [12]’s definition of $k(\lambda_*)$, and we will show that $\lambda_* \cdot \gamma > k(\lambda_*)$, implying that γ is not a limit PPE payoff.

We denote by $Z = A \times B$ the set of public signals in our game. The quantity $k(\lambda_*)$ is defined as the value of the optimization problem \mathcal{P}

$$\sup V^2,$$

where the supremum is taken over all $(V^1, V^2) \in \mathbf{R}^2$, and all $\phi : Z \rightarrow \mathbf{R}^2$, such that

- $\phi^2(z) \leq 0$ for every $z \in Z$;
- (V^1, V^2) is a Nash equilibrium payoff of the one-shot game, with payoff function defined by

$$\sum_{s \in S} m(s)(u(s, y(x(s))) + \phi(x(s), y(x(s)))) \tag{8}$$

for each action pair (x, y) .

Let $\phi : Z \rightarrow \mathbf{R}^2$ be any function such that $\phi^2(z) \leq 0$ for each $z \in Z$, and let (α, β) be any (possibly mixed) equilibrium of the one-shot game (8) with payoff (V^1, V^2) . We will prove that $V^2 < \frac{7}{6}$. We argue by contradiction, and assume that $V^2 \geq \frac{7}{6}$. We distinguish between two cases.

Assume first that $\alpha : S \rightarrow \Delta(A)$ is pooling: the distribution of messages is the same in both states. Then, since $\phi^2(z) \leq 0$, the expected payoff of the receiver is not higher than

$$\max_{b \in B} \sum_{s \in S} m(s)u^2(s, b) = 1.$$

Thus, $V^2 \leq 1 < \frac{7}{6}$, which is the desired contradiction.

Assume next that α is not pooling. Up to a relabeling of the messages we may then assume w.l.o.g. that the sender always tells the truth with positive probability. That is, $\alpha(s | s) > 0$ for each $s \in S$. We denote by $\tilde{\beta}(\cdot | s)$ the conditional distribution of the receiver’s move under (α, β) , conditional on the state being $s \in S$. Denoting by $s \neq t$ the two states, the equilibrium property for the sender in the game (8) then implies that

$$\sum_{b \in B} \beta(b | s)(u^1(s, b) + \phi^1(s, b)) \geq \sum_{b \in B} \beta(b | t)(u^1(s, b) + \phi^1(t, b)),$$

with equality if $\alpha(\cdot | s)$ assigns positive probability to both messages, and

$$\sum_{b \in B} \beta(b | t)(u^1(t, b) + \phi^1(t, b)) \geq \sum_{b \in B} \beta(b | s)(u^1(t, b) + \phi^1(s, b)).$$

¹⁶ Their result requires that a certain set have a non-empty interior, a condition that can be checked to be met here.

Using the two inequalities, one can verify that

$$u^1(s, \tilde{\beta}(\cdot | s)) + u^1(t, \tilde{\beta}(\cdot | t)) \geq u^1(s, \tilde{\beta}(\cdot | t)) + u^1(t, \tilde{\beta}(\cdot | s)).$$

By Lemma 1, condition C2 therefore holds for the stationary strategy $\tilde{\beta}$.

On the other hand, since $\phi^2(z) \leq 0$ for each z , the expected payoff V^2 to the receiver does not exceed $U^2(\mu_0, \tilde{\beta})$. Hence, $U^2(\mu_0, \tilde{\beta}) \geq \frac{7}{6}$. This readily implies that $U(\mu_0, \tilde{\beta}) \in E(\mathcal{M})$.

Next, one can verify that the highest payoff $U^2(\mu_0, \tilde{\beta})$ to the receiver, over the whole set $U(\mu_0, \tilde{\beta}) \in E(\mathcal{M})$, is equal to $\frac{7}{6}$. In addition, the unique strategy $\tilde{\beta}$ that achieves such a payoff is the strategy y_* . Since the supports of $y_*(\cdot | L)$ and $y_*(\cdot | R)$ are distinct, it must therefore be that α is truth-telling: $\alpha(s | s) = 1$ for each s . Therefore, β is equal to y_* .

Since $V^2 \geq \frac{7}{6}$ and $V^2 \leq U^2(\mu_0, y_*)$, one also has $V^2 = U^2(\mu_0, y_*)$. In particular, the expectation of $\phi^2(z)$ under the equilibrium profile (α, β) must be equal to zero. Since $\phi^2(z) \leq 0$ for each z , this implies that $\phi^2(z) = 0$, for each public signal z that receives positive probability under (α, β) .

Using this, we finally claim that the equilibrium condition for the receiver in the game (8) is violated. Indeed, when told L , the strategy $\beta = y_*$ assigns positive probability to both l and m . Hence, $\phi^2(L, l) = \phi^2(L, m) = 0$ by the previous paragraph. On the other hand, $u^2(L, m) > u^2(L, l)$, hence the receiver is not indifferent between both actions. This is the desired contradiction.

6.3. Imperfect monitoring

Let us assume here that successive states are independent. Results continue to hold if the receiver only observes a noisy, *public* version of the sender's message, provided the definition of $U(\mu, y)$ is modified in an appropriate way. They still hold if the receiver observes a noisy, *public* signal of the current state, provided the individual rationality level v^2 is modified in the proper way. They also hold, without changes, if the sender only observes a noisy, *public* signal of the receiver's action. What happens in any of these variants when signals are private is beyond the scope of the paper.

We briefly conclude this section by discussing the case where the sender fails to receive any information relative to the receiver's choices. In spite of this feature, the game does not reduce to a sequence of successive, independent, one-shot games, because of the ability of the receiver to monitor the sender. In particular, it is easy to construct examples with equilibrium payoffs that lie outside of the convex hull of the set of equilibrium payoffs in the one-shot game.

We refer to the game where the sender does not observe the actions of the receiver as the *blind* game. Denote by NE_δ^b the set of all Nash equilibrium payoffs of the blind game. We prove that the value of monitoring is positive, in the sense that allowing the sender to monitor the receiver has a non-ambiguous effect on the equilibrium set.

Proposition 2. *The set NE_δ^b is a subset of NE_δ .*

Proof. Let (σ, τ) be a Nash equilibrium of the blind game. Define τ' to be the following strategy that depends only on the sender's announcements, and not on the receiver's past actions: after a sequence (a_1, \dots, a_n) of announcements, τ' plays any action $b \in B$ with the probability that the

l	m	r
2, 2	0, 0	0, 0
<i>State L</i>		

l	m	r
0, 0	2, 2	0, 3
<i>State R</i>		

Fig. 8. The game in Example 4.

n -th action of the receiver according to τ is b , conditional on the sender’s announcements being (a_1, \dots, a_n) :

$$\tau'(a_1, \dots, a_n)[b] = \mathbf{E}[\tau(a_1, b_1, a_2, b_2, \dots, b_{n-1}, a_n)[b] \mid a_1, a_2, \dots, a_n].$$

In words, τ' gets rid of the possible correlation between successive actions of the receiver, that may exist in the strategy τ .

We claim that the strategy profile (σ, τ') is a Nash equilibrium of the blind game. Indeed, τ' is a best reply to σ because it induces the same payoff as τ . σ is a best reply to τ' because any strategy of the sender in the *blind* game induces the same expected payoff against τ or τ' .

We next claim that the strategy profile (σ, τ') is a Nash equilibrium of the non-blind game. Indeed, because under σ , the sender does not condition his play on past actions of the receiver, and because τ' is a best response to σ in the blind game, it follows that τ' is a best response to σ in the non-blind game as well. Because the receiver’s actions are conditionally independent, given the sender’s announcements, any profitable deviation against τ' in the non-blind game is also profitable in the blind game. \square

The inclusion is strict in general, as Example 4 below shows.

Example 4. There are two states $S = \{L, R\}$, and three actions for the receiver, $B = \{l, m, r\}$. Successive states are independent and the two states are equally likely. The payoffs in the two states are given in Fig. 8.

We claim that $(2, 2)$ is an equilibrium payoff when the sender observes the actions of the receiver, but it is no longer an equilibrium payoff when the sender does not observe the receiver’s actions.

Note first that $v^2 = \frac{3}{2}$, and that $\widehat{E}(\mathcal{M}) \neq \emptyset$. By Theorem 1, $(2, 2) \in E(\mathcal{M})$, so that $(2, 2) \in \liminf_{\delta \rightarrow 1} SE_\delta \subseteq \liminf_{\delta \rightarrow 1} NE_\delta$.

We now argue that $(2, 2)$ is bounded away from the set NE_δ^b . Indeed, assume to the contrary that there is some equilibrium profile (σ, τ) of the blind game with a payoff close to $(2, 2)$. In particular, with a probability close to one, there is a positive fraction of the stages in which the current state is R and the receiver plays m . Consider the strategy τ' which plays as τ , except that τ' plays r whenever τ would play m . Because the sender does not observe the receiver’s actions, he cannot tell whether the receiver uses τ or τ' , and therefore τ' is a profitable deviation of the receiver: it yields the receiver payoff close to $2\frac{1}{2}$. \blacklozenge

6.4. Relation to the one-shot game

The characterization implies that every equilibrium payoff of the one-shot game remains an equilibrium payoff in the dynamic game, provided players are patient enough. This property is not obvious *a priori*, since the game is not a repeated game. In particular, it would typically fail to hold if the state were constant throughout the play.

Let (σ, τ) be an equilibrium of the one-shot game. Let $y : A \rightarrow \Delta(B)$ be the stationary strategy defined as

$$y(b | a) = \sum_{s \in S} \sigma(a | s) \tau(a)[b].$$

Note that the expected payoff under (σ, τ) is $U(\mu_0, y)$. We claim that $U(\mu_0, y) \in E(\mathcal{M})$, so that by **Theorem 1** it is a sequential equilibrium payoff in the repeated game. Indeed, because the receiver can guarantee v^2 in the one-shot game, condition C2 holds. Because σ is a best reply to τ in the one-shot game, the inequality in C1 holds for every μ , and in particular for every $\mu \in \mathcal{M}$.

This result has the implication that the lowest equilibrium payoff of the sender in the repeated game cannot be higher than his lowest equilibrium payoff in the one shot game. As the example in Section 3 shows, it can in fact be strictly lower.

On the other hand, the lowest equilibrium payoff of the receiver in both the one-shot game and the repeated game is equal to his babbling equilibrium payoff v^2 .

Appendix A. Proof of Lemma 1

To prove **Lemma 1** we need the following description of \mathcal{M} , which is of independent interest.

A permutation matrix is a (square) matrix with entries in $\{0, 1\}$, such that each row and each column contains exactly one entry equal to 1. We denote by Φ the set of $S \times S$ permutation matrices, and by I the matrix that corresponds to the identity permutation.

Lemma 5. *The set $\mathcal{M}(m)$ is equal to*

$$\mathcal{M}(m) = (\mu_0 - I + \text{co } \Phi) \cap \mathbf{R}_+^{S \times S}.$$

Proof. The inclusion \supseteq is clear. We prove the reverse inclusion. Take μ in $\mathcal{M}(m)$, and define the matrix $J := \mu + I - \mu_0$ in $\mathbf{R}^{S \times S}$. J is a bistochastic matrix, hence it is a convex combination of permutation matrices. Since $\mu = J - I + \mu_0$, the result follows. \square

Proof of Lemma 1. We only prove that C1 is equivalent to C'1. For every permutation ϕ over S denote by $\mu^\phi \in S \times S$ the matrix where the entry (s, t) is equal to 1 if $t = \phi(s)$, and is 0 otherwise. Note that $U^1(I, y) = \sum_{s \in S} u^1(s, y(\cdot | s))$, and $U^1(\mu^\phi, y) = \sum_{s \in S} u^1(s, y(\cdot | \phi(s)))$.

Assume first that C'1 holds, and let $\mu \in \mathcal{M}(m)$. By **Lemma 5**, μ can be written $\mu = \mu_0 - I + \sum_{\phi} \alpha_{\phi} P^{\phi}$, where the α_{ϕ} are non-negative real numbers that sum to one. Because U^1 is linear in μ ,

$$U^1(\mu, y) = U^1(\mu_0, y) - U^1(I, y) + \sum_{\phi} \alpha_{\phi} U^1(\mu^{\phi}, y).$$

By C'1, $U^1(I, y) \geq U^1(\mu^{\phi}, y)$ for every permutation ϕ , and therefore

$$U^1(I, y) \geq \sum_{\phi} \alpha_{\phi} U^1(\mu^{\phi}, y).$$

It follows that $U^1(\mu_0, y) \geq U^1(\mu, y)$. Because this inequality holds for every $\mu \in \mathcal{M}(\mu)$, C1 holds.

Assume now that C1 holds. Fix a permutation ϕ and define $\mu_\varepsilon = \mu_0 - \varepsilon I + \varepsilon \mu^\phi$, where $\varepsilon > 0$. Because m has full support, one has $\mu_\varepsilon \in \mathcal{M}(m)$ provided ε is sufficiently small. Now, by C1, it follows that for each such ε ,

$$U^1(\mu_0, y) \geq U^1(\mu_\varepsilon, y) = U^1(\mu_0, y) - \varepsilon U^1(I, y) + \varepsilon U^1(\mu^\phi, y).$$

We conclude that $U^1(\mu^\phi, y) \leq U^1(I, y)$. As this inequality holds for every permutation ϕ , C' 1 holds. \square

Appendix B. Complements to the proof of Theorem 1

The proof of Theorem 1 given in the text is almost complete. For completeness we provide below the proofs of Proposition 1 and Lemma 3, which are missing.

We start by addressing the issue of designing a system of beliefs for the receiver that is consistent with σ_* and that satisfies an additional property. Since the game involves randomizing devices with uncountably many outcomes, the standard definition of consistency does not apply. We denote by $\lambda \in \Delta(S)$ a distribution with full support and, for $\eta < 0$, we denote by σ_η the strategy that, following any history h_n , plays $\eta\lambda + (1 - \eta)\sigma_*(h_n)$.

One can check that for $\eta > 0$ the beliefs of the receiver are uniquely defined by Bayes rule, and have a limit¹⁷ when $\eta = 0$. Note that following any history that is inconsistent with τ_0 , the belief of the receiver in stage n is independent¹⁸ of t_n .

We denote by τ_* a strategy that coincides with τ_0 as long as the sender does not deviate, and that plays in each later stage n an action that (i) maximizes the current expected payoff of the receiver given the belief held by the receiver in stage n , and (ii) does not depend on the announcements made by the sender since the deviation took place.

By construction, the strategy σ_* is sequentially rational at each information set of the sender, while the strategy τ_* is sequentially rational at each information set of the receiver that is inconsistent with τ_0 .

B.1. Proof of Proposition 1

Assume w.l.o.g. that all payoffs belong to the interval $[0, 1]$. Define by σ_{truth} the strategy of the sender that announces truthfully the current state s_n in each stage $n \leq N$, and by τ_{truth} the strategy of the receiver that plays $y(t_n)$ in each stage $n \leq N$. Thus, τ_{truth} coincides with τ_0 until stage q .

Let η be given and set $\xi = \frac{\eta}{|S|+2}$. For every state $s \in S$ and every $n \in \mathbb{N}$ denote by $F_n(s)$ the empirical frequency of visits to s up to (and including) stage n . Since the Markov chain is aperiodic, by the ergodic theorem there is N_0 such that with probability at least $1 - \xi$, we have $F_{(1-\xi)N}(s) \leq m_N(s)$ for every state $s \in S$ for every $N \geq N_0$. It follows that τ_0 coincides with τ_{truth} in the first $(1 - \xi)N$ stages, so that with probability at least $1 - \xi$,

$$\|\mu_{\sigma_{truth}, \tau_0} - \mu_0\|_1 \leq |S|\xi.$$

This implies that

$$U^1(\mu_{\sigma_{truth}, \tau_0}, y) > U^1(\mu_0, y) - (|S| + 1)\xi.$$

¹⁷ And the convergence is uniform w.r.t. the receiver's information set.

¹⁸ That is, should the sender fail to play the babbling announcement \bar{a} , the receiver still interprets the sender's announcements as babbling.

For a fixed $N \in \mathbf{N}$, as δ converges to 1, the discounted payoff in each block converges to the average payoff in that block, and therefore for δ sufficiently large

$$\gamma_\delta^1(\sigma_{truth}, \tau_0) > U^1(\mu_0, y) - (|S| + 2)\xi.$$

Because σ_0 is a best reply to τ_0 , we deduce that

$$\gamma_\delta^1(\sigma_0, \tau_0) \geq \gamma_\delta^1(\sigma_{truth}, \tau_0) > U^1(\mu_0, y) - (|S| + 2)\xi = U^1(\mu_0, y) - \eta.$$

We again use the fact that for a fixed $N \in \mathbf{N}$, as δ goes to 1, the payoff $\gamma_\delta(\sigma_0, \tau_0)$ converges to $U(\mu_{\sigma_0, \tau_0}, y)$ to deduce that

$$U^1(\mu_{\sigma_0, \tau_0}, y) > U^1(\mu_0, y) - \eta. \tag{9}$$

For a fixed $N \in \mathbf{N}$, and for every δ , the marginal distributions of $\mu_{\sigma_0, \tau_0} \in \Delta(S \times A)$ on S and A are respectively equal to m and to m_N .

Since the approximation m_N converges to m as $N \rightarrow +\infty$, the distribution μ_{σ_0, τ_0} converges to the set \mathcal{M} of copulas. Using Lemma 3, Proposition 1 therefore follows from (9).

B.2. Proof of Lemma 3

Let a copula $\mu \in \mathcal{M}$ be given. Present μ as a convex combination of the extreme points $(\mu_e)_e$ of \mathcal{M} : $\mu = \sum_{\mu_e \in \mathcal{M}_e} \alpha_e \mu_e$, with $\alpha_e \geq 0$ and $\sum_{\mu_e \in \mathcal{M}_e} \alpha_e = 1$. Recall that μ_0 is one of the extreme points of \mathcal{M} .

On the one hand, since U^1 is bi-linear, one has

$$\begin{aligned} U^1(\mu_0, y) - U^1(\mu, y) &= \varepsilon U^1(\mu_0, y_0) + (1 - \varepsilon)U^1(\mu_0, y_1) - \varepsilon U^1(\mu, y_0) \\ &\quad - (1 - \varepsilon)U^1(\mu, y_1) \\ &\geq \varepsilon(U^1(\mu_0, y_0) - U^1(\mu, y_0)) \end{aligned} \tag{10}$$

$$= \varepsilon \left((1 - \alpha_0)U^1(\mu_0, y_0) - \sum_{\mu_e \neq \mu_0} \alpha_e U^1(\mu_e, y_0) \right) \tag{11}$$

$$\geq \varepsilon(1 - \alpha_0)c_1, \tag{12}$$

where the inequality (10) holds because $y_1 \in Y(\mathcal{M})$ and by C1.

On the other hand one has $\mu - \mu_0 = \sum_{\mu_e \in \mathcal{M}_e} \alpha_e (\mu_e - \mu_0)$, hence

$$\|\mu - \mu_0\|_1 \leq c_2 \sum_{\mu_e \in \mathcal{M}_e, \mu_e \neq \mu_0} \alpha_e = c_2(1 - \alpha_0). \tag{13}$$

The result follows from (12) and (13).

Appendix C. Complements to the proof of Theorem 2

We here prove Lemma 4. For clarity, we introduce yet another copy T of the set S . Intuitively, fictitious states are T -valued, while realized ones are S -valued.

Define $\mathcal{M}' \subseteq \Delta(S \times T)$ to be the set of distributions $\mu \in \mathcal{M}$ such that the following Property P holds:

Property P. For every $(s, t) \in S \times T$, one has

$$\sum_{s' \in S} \mu(s' | t) p(s | s') = \sum_{t' \in T} \mu(s | t') p(t' | t). \tag{14}$$

We will prove

Lemma 6. Under Assumption A, the set \mathcal{M}' coincides with the set \mathcal{M} .

Lemma 7. Let $\mu \in \mathcal{M}'$ be given. There exists an S -valued process¹⁹ $(\mathbf{t}_n)_n$, such that the following hold:

- P1. The law of the sequence $(\mathbf{t}_n)_n$ is the same as the law of the sequence $(\mathbf{s}_n)_n$.
- P2. The law of the pair $(\mathbf{s}_n, \mathbf{t}_n)$ is μ for each stage $n \in \mathbf{N}$.
- P3. The conditional law of \mathbf{s}_n given $\mathbf{t}_1, \dots, \mathbf{t}_n$ is $\mu(\cdot | \mathbf{t}_n)$.
- P4. Conditional on \mathbf{s}_n , the vector $(\mathbf{t}_1, \dots, \mathbf{t}_n)$ is independent of the future states $(\mathbf{s}_{n+1}, \mathbf{s}_{n+2}, \dots)$.

We emphasize that only Lemma 6 makes use of Assumption A. This has the following consequence. Given $\mu \in \mathcal{M}'$, using Lemma 7 and the construction presented in the paper, one has $U^1(\mu_0, y) \geq U^1(\mu, y)$. Thus, the conclusion $U^1(\mu_0, y) = \max_{\mu \in \mathcal{M}'} U^1(\mu, y)$ holds, irrespective of whether Assumption A is met or not.

C.1. Proof of Lemma 7

Let $\mu \in \mathcal{M}'$ be given, and define $\bar{\mu} \in \Delta(T \times S \times T)$ by

$$\bar{\mu}(t', s, t) = \mu(s, t) p(t | t') \frac{m(t')}{m(t)}, \quad (t', s, t) \in T \times S \times T. \tag{15}$$

For every two indices $i, j \in \{1, 2, 3\}$ with $i < j$ denote by $\bar{\mu}_{i,j}$ the marginal of $\bar{\mu}$ on the i -th and j -th coordinates.

We will use the following properties of $\bar{\mu}$.

Lemma 8. One has

- 1. $\bar{\mu}_{2,3} = \mu$;
- 2. $\bar{\mu}_{1,3}(t', t) = m(t') p(t | t')$ for every $t, t' \in T$;
- 3. $\bar{\mu}_{1,2}(t', s') = \sum_{s \in S} \bar{\mu}_{2,3}(s, t') p(s' | s)$, for each $t' \in T, s' \in S$;
- 4. $\bar{\mu}(s | t', t) = \mu(s | t)$ for each $(t', s, t) \in T \times S \times T$.

Proof. We prove the four claims in turn. Let $s, t \in S \times T$ be given. One has

$$\bar{\mu}_{2,3}(s, t) = \sum_{t' \in T} \bar{\mu}(t', s, t) = \sum_{t' \in T} \mu(s, t) p(t | t') \frac{m(t')}{m(t)}$$

¹⁹ The process $(\mathbf{t}_n)_n$ is possibly defined on a probability space which is an enlargement of the one on which $(\mathbf{s}_n)_n$ is defined.

$$= \frac{\mu(s, t)}{m(t)} \sum_{t' \in T} p(t | t') m(t') = \mu(s, t),$$

which proves the first claim.

To prove the second claim fix $t', t \in T$. One has

$$\bar{\mu}_{1,3}(t', t) = \sum_{s \in S} \bar{\mu}(t', s, t) = \sum_{s \in S} \mu(s, t) p(t | t') \frac{m(t')}{m(t)} = p(t | t') m(t'),$$

where the last equality holds since the marginal distribution of μ on S is m .

We turn to the third claim. Let $t' \in T, s' \in S$ be given. By the first claim and since $\mu \in \mathcal{M}$ one has

$$\sum_{s \in S} \bar{\mu}_{2,3}(s, t') p(s' | s) = \sum_{s \in S} \mu(s, t') p(s' | s) = m(t') \sum_{t \in T} \mu(s' | t) p(t | t'). \tag{16}$$

On the other hand,

$$\bar{\mu}_{1,2}(t', s') = \sum_{t \in T} \bar{\mu}(t', s', t) = \sum_{t \in T} \mu(s', t) p(t | t') \frac{m(t')}{m(t)}. \tag{17}$$

The third claim follows from (16) and (17).

Finally let $(t', s, t) \in T \times S \times T$ be given. By the second claim,

$$\bar{\mu}(s | t', t) = \frac{\bar{\mu}(t', s, t)}{\bar{\mu}_{1,3}(t', t)} = \frac{\mu(s, t) p(t | t')}{p(t | t') m(t')} \times \frac{m(t')}{m(t)} = \mu(s | t),$$

and the fourth claim follows. \square

We construct the sequence $(\mathbf{t}_n)_n$ as follows. The initial values \mathbf{t}_0 and \mathbf{t}_1 are drawn according to the conditional distribution $\bar{\mu}(\cdot | \mathbf{s}_1) \in \Delta(T \times T)$. For $n \neq 2, \mathbf{t}_n$ is drawn according to the conditional distribution $\bar{\mu}(\cdot | \mathbf{t}_{n-1}, \mathbf{s}_n)$. In this construction, \mathbf{t}_0 is used to unify the treatment of \mathbf{s}_1 with that of $(\mathbf{s}_n)_{n \geq 2}$. Property P4 thus holds by construction. Properties P1 and P2 follow from the next lemma.

Lemma 9. *The law of $(\mathbf{t}_{n-1}, \mathbf{s}_n, \mathbf{t}_n)$ is equal to $\bar{\mu}$, for each stage $n \geq 1$.*

Proof. We argue by induction. Observe that the law of \mathbf{s}_1 is equal to m . Therefore,

$$\mathbf{P}((\mathbf{t}_0, \mathbf{s}_1, \mathbf{t}_1) = (t', s, t)) = m(s) \bar{\mu}(t', t | s) = \bar{\mu}(t', s, t).$$

Assume that the claim holds for some $n \in \mathbf{N}$. We will prove that the law of $(\mathbf{t}_n, \mathbf{s}_{n+1})$ is equal to $\bar{\mu}_{1,2}$. This follows from the following sequence of equalities, which holds for every $t' \in T, s \in S$:

$$\begin{aligned} \mathbf{P}((\mathbf{t}_n, \mathbf{s}_{n+1}) = (t', s)) &= \sum_{s' \in S} \mathbf{P}((\mathbf{s}_n, \mathbf{t}_n, \mathbf{s}_{n+1}) = (s', t', s)) \\ &= \sum_{s' \in S} \mathbf{P}((\mathbf{s}_n, \mathbf{t}_n) = (s', t')) \times \mathbf{P}(\mathbf{s}_{n+1} = s | (\mathbf{s}_n, \mathbf{t}_n) = (s', t)) \\ &= \sum_{s' \in S} \bar{\mu}_{2,3}(s', t') p(s | s') = \bar{\mu}_{1,2}(t', s), \end{aligned}$$

where the last equality follows from Lemma 8(3) and P4. Since the conditional law of \mathbf{t}_{n+1} given $(\mathbf{t}_n, \mathbf{s}_{n+1})$ is equal to $\bar{\mu}(\cdot | \mathbf{t}_n, \mathbf{s}_{n+1})$, this yields the claim for $n + 1$. \square

Finally, property P3 follows from the second part of the next lemma. The first part of the lemma is needed to the proof of the second part.

Lemma 10. (1) *The conditional law of \mathbf{t}_n given $(\mathbf{t}_0, \dots, \mathbf{t}_{n-1})$ coincides with the conditional law of \mathbf{t}_n given \mathbf{t}_{n-1} .*

(2) *The conditional law of \mathbf{s}_n given $(\mathbf{t}_0, \dots, \mathbf{t}_{n-1}, \mathbf{t}_n)$ coincides with the conditional law of \mathbf{s}_n given \mathbf{t}_n .*

Proof. The proof is by induction. For $n = 1$, the first statement trivially holds, while the second statement holds by Lemma 8(1). Assume that the claim holds for some $n \in \mathbf{N}$. For brevity, we denote by t_n, s_n, \dots generic values of $\mathbf{t}_n, \mathbf{s}_n \dots$ and we write $\mathbf{P}(t_n, s_n)$ instead of $\mathbf{P}((\mathbf{t}_n, \mathbf{s}_n) = (t_n, s_n))$.

Observe first that by the definition of (\mathbf{t}_n) ,

$$\begin{aligned} \mathbf{P}(t_{n+1} | t_0, \dots, t_n) &= \sum_{s_{n+1} \in S} \mathbf{P}(s_{n+1} | t_0, \dots, t_n) \mathbf{P}(t_{n+1} | t_0, \dots, t_n, s_{n+1}) \\ &= \sum_{s_{n+1} \in S} \mathbf{P}(s_{n+1} | t_0, \dots, t_n) \times \bar{\mu}(t_{n+1} | t_n, s_{n+1}). \end{aligned} \tag{18}$$

Moreover,

$$\begin{aligned} \mathbf{P}(s_{n+1} | t_0, \dots, t_n) &= \sum_{s_n \in S} \mathbf{P}(s_n | t_0, \dots, t_n) \mathbf{P}(s_{n+1} | s_n, t_0, \dots, t_n) \\ &= \sum_{s_n \in S} \mathbf{P}(s_n | t_0, \dots, t_n) \times p(s_{n+1} | s_n) \\ &= \sum_{s_n \in S} \mathbf{P}(s_n | t_n) \times p(s_{n+1} | s_n), \end{aligned} \tag{19}$$

where the last equality holds by the induction hypothesis. Note that the right-hand side of (19) is independent of $(t_0, t_1, \dots, t_{n-1})$, and therefore

$$\mathbf{P}(s_{n+1} | t_0, \dots, t_n) = \mathbf{P}(s_{n+1} | t_n). \tag{20}$$

Plugging (19) in (18), one obtains

$$\mathbf{P}(t_{n+1} | t_0, \dots, t_n) = \sum_{s_{n+1} \in S} \sum_{s_n \in S} \mathbf{P}(s_n | t_n) \times p(s_{n+1} | s_n) \times \bar{\mu}(t_{n+1} | t_n, s_{n+1}).$$

The right-hand side is independent of t_1, \dots, t_{n-1} , hence it is equal to $\mathbf{P}(t_{n+1} | t_n)$, and the first part of the lemma follows.

We turn to the second statement. One has

$$\begin{aligned} \mathbf{P}(s_{n+1} | t_0, \dots, t_{n+1}) &= \frac{\mathbf{P}(s_{n+1}, t_{n+1} | t_0, \dots, t_n)}{\mathbf{P}(t_{n+1} | t_0, \dots, t_n)} \\ &= \frac{\mathbf{P}(s_{n+1} | t_0, \dots, t_n) \times \mathbf{P}(t_{n+1} | s_{n+1}, t_0, \dots, t_n)}{\mathbf{P}(t_{n+1} | t_0, \dots, t_n)} \\ &= \frac{\mathbf{P}(s_{n+1} | t_n) \bar{\mu}(t_{n+1} | s_{n+1}, t_n)}{\mathbf{P}(t_{n+1} | t_n)} \end{aligned}$$

$$\begin{aligned}
 &= \mathbf{P}(s_{n+1} | t_n) \frac{\mu(s_{n+1}, t_{n+1})}{\bar{\mu}(t_n, s_{n+1})} \frac{m(t_n)}{m(t_{n+1})} \\
 &= \frac{\mu(s_{n+1}, t_{n+1})}{m(t_{n+1})} = \mathbf{P}(s_{n+1} | t_{n+1}),
 \end{aligned}$$

where the third equality holds by (20), the construction of $(t_n)_n$, and the first claim, and the fourth equality holds by (15). This concludes the proof of the induction step. \square

The proof of Lemma 7 is now completed.

C.2. Proof of Lemma 6

We here verify that if Assumption A holds then $\mathcal{M} = \mathcal{M}'$. Let p be a transition function such that $p(s' | s) = \alpha_{s'}$ for every two states $s \neq s'$, and $p(s | s) = 1 - \sum_{s' \neq s} \alpha_{s'}$. Set $C = \sum_{s \in S} \alpha_s$. One can verify that the invariant measure of p is given by $m(s) = \frac{\alpha_s}{C}$ for each $s \in S$.

Let $\mu \in \mathcal{M}$. We will prove that for every $(t, s') \in T \times S$, the equality

$$\sum_{s \in S} \mu(s | t) p(s' | s) = \sum_{t' \in T} \mu(s' | t') p(t' | t) \tag{21}$$

holds. Fix $t \in T$ and $s' \in S$. Observe that

$$\begin{aligned}
 \sum_{s \in S} \mu(s | t) p(s' | s) &= \mu(s' | t) \left(1 - \sum_{s \neq s'} \alpha_s \right) + \sum_{s \neq s'} \alpha_{s'} \mu(s | t) \\
 &= \mu(s' | t) \left(1 - \sum_{s \neq s'} \alpha_s \right) + \alpha_{s'} (1 - \mu(s' | t)) \\
 &= \alpha_{s'} + \mu(s' | t) (1 - C).
 \end{aligned} \tag{22}$$

We also have

$$\sum_{t' \in T} \mu(s' | t') p(t' | t) = \mu(s' | t) \left(1 - \sum_{t' \neq t} \alpha_{t'} \right) + \sum_{t' \neq t} \mu(s' | t') \alpha_{t'} \tag{23}$$

$$= \mu(s' | t) (1 - C + \alpha_t) + \sum_{t' \neq t} \mu(s' | t') \alpha_{t'}. \tag{24}$$

Subtracting (23) from (22) one obtains

$$\begin{aligned}
 \sum_{s \in S} \mu(s | t) p(s' | s) - \sum_{t' \in T} \mu(s' | t') p(t' | t) &= \alpha_{s'} - \mu(s' | t) \alpha_t - \sum_{t' \neq t} \mu(s' | t') \alpha_{t'} \\
 &= \alpha_{s'} - \sum_{t' \in S} \mu(s' | t') \alpha_{t'} \\
 &= \alpha_{s'} - C \sum_{t' \in S} \mu(s' | t') m(t') \\
 &= \alpha_{s'} - C m(s') = 0,
 \end{aligned} \tag{25}$$

where (25) and (25) hold since $\alpha(s) = C m(s)$ for every $s \in S$. This proves (21), as desired.

Appendix D. Proof of Theorem 3

The proof of Theorem 3 consists of two independent parts. We first prove that if Condition B does not hold for some game G , then it does not hold throughout some neighborhood of G .

Proposition 3. *Let G be a game that does not satisfy Condition B. Then there is a neighborhood \mathcal{N} of G such that no game in \mathcal{N} satisfies Condition B.*

Proof. The proof relies on the theory of semi-algebraic sets. We refer to [7] for the results used below. Recall that the set of extreme points of the polytope \mathcal{M} is denoted by \mathcal{M}_e .

We will use the following two properties, that hold for constant functions $y : S \rightarrow \Delta(B)$.

- R1. If $y : S \rightarrow \Delta(B)$ is constant, then $U^1(\mu_0, y) = U^1(\mu, y)$ for every $\mu \in \mathcal{M}$.
- R2. If $y : S \rightarrow \Delta(B)$ is constant, then $U^2(\mu_0, y) \leq v^2$.

Property R1 holds because when y is constant the payoff is independent of the sender’s announcements. Property R2 holds because v^2 is the maximum of $U^2(\mu_0, y)$ over all constant functions y .

Given a payoff function $\tilde{u} : S \times B \rightarrow \mathbf{R}^2$, we denote by $\mathcal{S}(\tilde{u})$ the system of inequalities

$$\tilde{U}^2(\mu_0, y) \geq v_{\tilde{u}}^2 \quad \text{and} \quad \tilde{U}^1(\mu_0, y) \geq \tilde{U}^1(\mu, y), \quad \text{for all } \mu \in \mathcal{M}_e,$$

with unknowns $y : S \rightarrow \Delta(B)$, where $v_{\tilde{u}}^2 = \max_{b \in B} \tilde{U}^2(\mu_0, b)$ is the minmax value of the receiver in the game with payoffs \tilde{u} .

We say that a vector $y \in \mathbf{R}^{S \times B}$ is *constant* if $y(s, b)$ only depends on b .

Let u denote the payoff function of G . By assumption, any solution y to $\mathcal{S}(u)$ is constant. We will show that this implies that all solutions to $\mathcal{S}(\tilde{u})$ are constant, for all \tilde{u} in a neighborhood of u .

Assume to the contrary that for every $\varepsilon > 0$ there is a payoff function $u_\varepsilon \in \mathbf{R}^{2(S \times B)}$ such that

- (i) $\|u - u_\varepsilon\| < \varepsilon$, and (ii) the system $\mathcal{S}(u_\varepsilon)$ has a non-constant solution $y_\varepsilon \in \mathbf{R}^{S \times B}$.

This implies that there is a *semi-algebraic* function $\varepsilon \in (0, 1) \mapsto (u_\varepsilon, y_\varepsilon)$ such that

- (i) $\lim_{\varepsilon \rightarrow 0} u_\varepsilon = u$, and (ii) y_ε is a *non-constant* solution to $\mathcal{S}(u_\varepsilon)$ for every $\varepsilon > 0$ small enough.

In particular, the function $\varepsilon \mapsto y_\varepsilon$ has an expansion to a Puiseux series in a neighborhood of zero: there exist $\varepsilon_0 > 0$, a natural number r , and vectors $y_k \in \mathbf{R}^{S \times B}$ for $k \geq 0$ such that

$$y_\varepsilon = \sum_{k=0}^{\infty} \varepsilon^{\frac{k}{r}} y_k,$$

for every $\varepsilon \in (0, \varepsilon_0)$, and a similar expansion exists for the function $\varepsilon \mapsto u_\varepsilon$.

Note that $y_0 = \lim_{\varepsilon \rightarrow 0} y_\varepsilon$. This implies in particular that $y_0(\cdot, s) \in \Delta(B)$ for every $s \in S$, and that y_0 is a solution to $\mathcal{S}(u)$. In particular, y_0 is constant.

Because $y_\varepsilon(\cdot, s) \in \Delta(B)$, it follows that $\sum_{b \in B} y_\varepsilon(b | s) = 1$ for every $\varepsilon > 0$ and every $s \in S$, so that $\sum_{b \in B} y_k(b | s) = 0$ for every $k \geq 1$ and every $s \in S$.

Let $l \geq 0$ be the maximal integer such that y_0, y_1, \dots, y_l are constant functions. Because y_ε is not constant for every $\varepsilon > 0$, we have $l < \infty$. Define a vector $d \in \mathbf{R}^B$ by

$$d(b) = \min_{s \in S} y_{l+1}(b, s), \quad \forall b \in B.$$

Note that

$$y_\varepsilon = \sum_{k=0}^{\infty} \varepsilon^{\frac{k}{r}} y_k = \left(\sum_{k=0}^l \varepsilon^{\frac{k}{r}} y_k + \varepsilon^{\frac{l+1}{r}} d \right) + \varepsilon^{\frac{l+1}{r}} (y_{l+1} - d) + \sum_{k=l+2}^{\infty} \varepsilon^{\frac{k}{r}} y_k. \tag{26}$$

The first term $(\sum_{k=0}^l \varepsilon^{\frac{k}{r}} y_k + \varepsilon^{\frac{l+1}{r}} d)$ is independent of s , and all its coordinates are non-negative because y_ε is non-negative for every $\varepsilon > 0$. Set

$$z_\varepsilon = \frac{\sum_{k=0}^l \varepsilon^{\frac{k}{r}} y_k + \varepsilon^{\frac{l+1}{r}} d}{1 + \varepsilon^{\frac{l+1}{r}} \sum_{b \in B} d(b)} \in \mathbf{R}^{S \times B}.$$

Then $z_\varepsilon(\cdot, s) \in \Delta(B)$ for every $s \in S$, and z_ε is independent of s . Set

$$w(\cdot, s) = \frac{y_{l+1}(\cdot, s) - d}{-\sum_{b \in B} d(b)} \in \mathbf{R}^{S \times B}, \quad \forall s \in S.$$

Then $w(s) \in \Delta(B)$ and w is non-constant. We will show that w solves $\mathcal{S}(u)$, contradicting the assumption that all solutions of $\mathcal{S}(u)$ are constant.

By R2, for every $\varepsilon > 0$ we have $\tilde{U}^2(\mu_0, z_\varepsilon) \leq v_{u_\varepsilon}^2$. But $\tilde{U}^2(\mu_0, y_\varepsilon) \geq v_{u_\varepsilon}^2$ and y_ε is a convex combination of z_ε , w , and a “tail” which is of a lower order of ε ; by taking the limit $\varepsilon \rightarrow 0$ and using $v_{u_\varepsilon}^2 \rightarrow v^2$ we obtain $U^2(\mu_0, w) \geq v^2$.

Fix $\mu \in \mathcal{M}_e$. By R1 it follows that $\tilde{U}^1(\mu_0, z_\varepsilon) = \tilde{U}^1(\mu, z_\varepsilon)$. Because $\tilde{U}^1(\mu_0, y_\varepsilon) \geq \tilde{U}^1(\mu, y_\varepsilon)$, it follows for the same reasoning as above that $U^1(\mu_0, w) \geq U^1(\mu, w)$. \square

We turn to the second part of the proof.

Proposition 4. *Let G be a game such that Condition B holds. Then any neighborhood of the game G contains a game G' such that $\tilde{E}_{G'}(\mathcal{M}) \neq \emptyset$.*

Proof. The proof combines three independent lemmas. We first show that there are perturbations of u^2 such that the inequality in (i) holds strictly for the perturbed game. Next, we show that the function y may be assumed to be one-to-one. Finally, we construct perturbations of u^1 such that the inequalities in (ii) will be strict.

Lemma 11. *Let G be a game with payoff function u and let $y : S \rightarrow \Delta(B)$ be a non-constant function such that $U^2(\mu_0, y) \geq v^2$. Then any neighborhood of u^2 contains payoff functions \tilde{u}^2 such that $\tilde{U}^2(\mu_0, y) > \tilde{v}^2$.*

Proof. Define $P \in \Delta(S \times B)$ by $P(s, b) := m(s)y(b | s)$ for every $s \in S$ and $b \in B$, and let $\varepsilon > 0$ be given. We abuse notations and still denote by P the two marginals of P over S and B . Note that $P(s) = m(s) > 0$ for each $s \in S$. Define $\tilde{u}^2 : S \times B \rightarrow \mathbf{R}$ by $\tilde{u}^2(s, b) = u^2(s, b)$ if $P(b) = 0$, and

$$\tilde{u}^2(s, b) = u^2(s, b) + \varepsilon \frac{P(s, b)}{P(s)P(b)} \quad \text{if } P(b) > 0.$$

We claim that $\tilde{U}^2(\mu_0, y) > \tilde{v}^2$. Since ε is arbitrary, the result will follow. Note first that for every $b \in B$ that satisfies $P(b) > 0$ one has

$$\tilde{U}^2(\mu_0, b) = U^2(\mu_0, b) + \varepsilon \sum_{s \in S} m(s) \frac{P(s, b)}{m(s)P(b)} = U^2(b, \mu_0) + \varepsilon.$$

Hence $\tilde{v}^2 = v^2 + \varepsilon$; see Eq. (1). On the other hand, since $y(b | s) = \frac{P(s,b)}{m(s)} = \frac{P(s,b)}{P(s)} = P(b | s)$,

$$\begin{aligned} \tilde{U}^2(\mu_0, y) &= U^2(\mu_0, y) + \varepsilon \sum_{s \in S, b \in B} m(s)y(b | s) \frac{P(s, b)}{P(s)P(b)} \\ &= U^2(\mu_0, y) + \varepsilon \sum_{s \in S} m(s) \sum_{b \in B} \frac{P(b | s)^2}{P(b)}. \end{aligned}$$

Viewed as a function of the probability distribution $q \in \Delta(B)$, the expression $\sum_{b \in B} \frac{(q(b))^2}{P(b)}$ is strictly convex, and admits a unique minimum equal to 1 when $q = P$. Thus, for every fixed state $s \in S$ one has $\sum_{b \in B} \frac{P(b|s)^2}{P(b)} \geq 1$ with a strict inequality whenever the conditional distribution $P(\cdot | s)$ differs from P . Since y is not constant, there exists at least one state s such that $P(\cdot | s) \neq P$. Therefore,

$$\tilde{U}^2(\mu_0, y) > U^2(\mu_0, y) + \varepsilon \geq v^2 + \varepsilon = \tilde{v}^2,$$

as desired. \square

Lemma 12. *Let G be a game with payoff function u and let $y : S \rightarrow \Delta(B)$ be such that $U^1(\mu_0, y) \geq U^1(\mu, y)$ for each $\mu \in \mathcal{M}$. Then any neighborhood of y in $\mathbf{R}^{S \times B}$ contains a one-to-one function $\tilde{y} : S \rightarrow \Delta(B)$ such that $U^1(\mu_0, \tilde{y}) \geq U^1(\mu, \tilde{y})$ for each $\mu \in \mathcal{M}$.*

Proof. It suffices to show the existence of a one-to-one function $\tilde{z} : S \rightarrow \Delta(B)$ such that $U^1(\mu_0, \tilde{z}) \geq U^1(\mu, \tilde{z})$ for each $\mu \in \mathcal{M}$. Indeed, the conclusion of the lemma then follows by setting $\tilde{y} = (1 - \varepsilon)y + \varepsilon\tilde{z}$ for $\varepsilon > 0$ small enough.

Let $(z_s)_{s \in S}$ be arbitrary distinct elements of $\Delta(B)$. Let $\tilde{\phi}$ be a permutation over S that maximizes the sum $\sum_{s \in S} u^1(s, z_{\psi(s)})$ over all permutations ψ and set $\tilde{z}_s = z_{\tilde{\phi}(s)}$ for every $s \in S$. By construction one has

$$\sum_{s \in S} u^1(s, \tilde{z}_s) \geq \sum_{s \in S} u^1(s, \tilde{z}_{\phi(s)}),$$

for every permutation ϕ over S . By Lemma 1, this implies $U^1(\mu_0, \tilde{z}) \geq U^1(\mu, \tilde{z})$ for every $\mu \in \mathcal{M}$, as desired. \square

Lemma 13. *Let G be a game with payoff function u and let $y : S \rightarrow \Delta(B)$ be a one-to-one function such that $U^1(\mu_0, y) \geq U^1(\mu, y)$ for each $\mu \in \mathcal{M}$. Then any neighborhood of u^1 contains payoff functions \tilde{u}^1 such that $\tilde{U}^1(\mu_0, y) > \tilde{U}^1(\mu, y)$ for each $\mu \in \mathcal{M}$, $\mu \neq \mu_0$.*

Note that the existence of a stationary strategy y that satisfies the requirements follows from Lemma 12.

Proof. Let G , u and y be as stated. Given $\varepsilon > 0$ we define $\tilde{u}^1 : S \times B \rightarrow \mathbf{R}$ by

$$\tilde{u}^1(s, b) = u^1(s, b) + \varepsilon y(b | s).$$

We will prove that for every $\varepsilon > 0$ one has $\tilde{U}^1(\mu_0, y) > \tilde{U}^1(\mu, y)$ for each $\mu \in \mathcal{M} \setminus \{\mu_0\}$.

Given a permutation ϕ over S we denote by $Y_\phi \in \mathbf{R}^{S \times B}$ the vector whose (s, b) -component is equal to $y(b | \phi(s))$. Then,

$$\sum_{s \in S} \tilde{u}^1(s, y(\cdot | \phi(s))) = \sum_{s \in S, b \in B} y(b | \phi(s)) \tilde{u}(s, b) \quad (27)$$

$$= \sum_{s \in S, b \in B} y(b | \phi(s)) u(s, b) + \varepsilon \sum_{s \in S, b \in B} y(b | \phi(s)) y(b | s) \quad (28)$$

$$= \sum_{s \in S} u^1(s, y(\cdot | \phi(s))) + \varepsilon \langle Y_{Id}, Y_\phi \rangle, \quad (29)$$

where $\langle Y_{Id}, Y_\phi \rangle = \sum_{s \in S, b \in B} y(b | \phi(s)) y(b | s)$ is the standard scalar product in $\mathbf{R}^{S \times B}$.

Since y is one-to-one, the vectors Y_ϕ and Y_{Id} are not co-linear as soon as $\phi \neq Id$. By Cauchy–Schwarz inequality, it follows that

$$\langle Y_{Id}, Y_\phi \rangle < \|Y_{Id}\|_2 \|Y_\phi\|_2 = \|Y_{Id}\|^2 = \langle Y_{Id}, Y_{Id} \rangle \quad (30)$$

where the first equality holds since the components of Y_ϕ are obtained by permuting the components of Y_{Id} .

By Lemma 1 one has

$$\sum_{s \in S} u^1(s, y(\cdot | \phi(s))) \leq \sum_{s \in S} u^1(s, y(\cdot | s)). \quad (31)$$

Plugging (31) into (27), one obtains

$$\sum_{s \in S} \tilde{u}^1(s, y(\cdot | \phi(s))) < \sum_{s \in S} u^1(s, y(\cdot | s)) + \langle Y_{Id}, Y_{Id} \rangle = \sum_{s \in S} \tilde{u}^1(s, y(\cdot | s)).$$

By Lemma 1 this yields $\tilde{U}^1(\mu_0, y) > \tilde{U}^1(\mu, y)$ for every $\mu \in \mathcal{M}$, $\mu \neq \mu_0$, as desired. \square

The proof of Proposition 4 follows from Lemmas 11, 12 and 13.

References

- [1] S. Athey, K. Bagwell, Collusion with persistent cost shocks, *Econometrica* 76 (2008) 493–540.
- [2] R.J. Aumann, S. Hart, Long cheap talk, *Econometrica* 71 (2003) 1619–1660.
- [3] R.J. Aumann, M.B. Maschler, Repeated Games with Incomplete Information, MIT Press, 1995.
- [4] M. Battaglini, Long-term contracting with Markovian consumers, *Amer. Econ. Rev.* 95 (2005) 637–658.
- [5] D. Baron, D. Besanko, Regulation and information in a continuing relationship, *Info. Econ. Pol.* 1 (1984) 447–470.
- [6] D. Besanko, Multiperiod contracts between principal and agent with adverse selection, *Econ. Letters* 17 (1985) 33–37.
- [7] J. Bochnak, M. Coste, M.F. Roy, *Real Algebraic Geometry*, Springer, 1998.
- [8] V.P. Crawford, J. Sobel, Strategic information transmission, *Econometrica* 50 (1982) 1431–1451.
- [9] M. Ekmekci, O. Gossner, A. Wilson, Impermanent types and permanent reputations, *J. Econ. Theory* 147 (2012) 142–178.
- [10] J.F. Escobar, J. Toikka, Efficiency in games with Markovian private information, mimeo, 2012.
- [11] J. Farrell, M. Rabin, Cheap talk, *J. Econ. Perspect.* 10 (1996) 103–118.
- [12] D. Fudenberg, K. Levine, Efficiency and observability with long-run and short-run players, *J. Econ. Theory* 62 (1994) 103–135.
- [13] M. Golosov, V. Skreta, A. Tsyvinski, A. Wilson, Dynamic strategic information transmission, preprint, 2009.
- [14] J.R. Green, N.L. Stokey, A two-person game of information transmission, *J. Econ. Theory* 135 (2007) 90–104.
- [15] J. Hörner, D. Rosenberg, E. Solan, N. Vieille, On a Markov game with one-sided incomplete information, *Operations Res.* 58 (4) (2010) 1107–1115.
- [16] J. Hörner, T. Sugaya, S. Takahashi, N. Vieille, Recursive methods in discounted stochastic games: an algorithm for $\delta \rightarrow 1$ and a folk theorem, *Econometrica* 79 (2011) 1277–1318.

- [17] M.O. Jackson, H.F. Sonnenschein, Overcoming incentive constraints by linking decisions, *Econometrica* 75 (2007) 241–258.
- [18] V. Krishna, J. Morgan, A model of expertise, *Quart. J. Econ.* 116 (2001) 747–775.
- [19] V. Krishna, J. Morgan, Contracting for information under imperfect commitment, *RAND J. Econ.* 39 (2008) 905–925.
- [20] G.J. Mailath, L. Samuelson, *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, 2006.
- [21] C. Phelan, Public trust and government betrayal, *J. Econ. Theory* 130 (2006) 27–43.
- [22] J. Renault, The value of Markov chain games with lack of information on one side, *Math. Operations Res.* 31 (2006) 490–512.
- [23] J. Sobel, Signaling games, in: *Encyclopedia of Complexity and Systems Science*, vol. 19, Springer, 2009, pp. 8125–8139.
- [24] T. Wiseman, Reputation and impermanent types, *Games Econ. Behav.* 62 (2008) 190–210.