

Dynamic Communication with Biased Senders*

Chiara Margaria[†] and Alex Smolin[‡]

October 12, 2016

Abstract

We study dynamic games in which senders with state-independent payoffs communicate to a single receiver. Senders' private information evolves according to an aperiodic and irreducible Markov chain. We prove an analog of a folk theorem—that any feasible and individually rational payoff can be approximated in a perfect Bayesian equilibrium if players are sufficiently patient. In particular, there are equilibria in which the receiver makes perfectly informed decisions in almost every period, even if no informative communication can be sustained in the stage game. We conclude that repeated interaction can overcome strategic limits of communication.

Keywords: Bayesian games, repeated games, communication, folk theorem.

JEL Codes: C72, C73, D82, D83.

1 Introduction

Crawford and Sobel (1982) introduced cheap-talk games as a basis for the analysis of strategic transmission of unverifiable information. They considered a one-shot game between a sender who has private information and a receiver who takes an action. In equilibrium, informative communication can be sustained, but misalignment of players' interests limits the amount of information that can be transmitted. In particular, if players' preferences are misaligned, truthful communication of private information cannot be sustained; otherwise, the sender would bias reports to induce her preferred outcomes. The strategic considerations restrain communication and translate into inefficiency as both players could benefit from a better-informed action.

The strategic limits of communication highlighted in a one-shot game provide insights into many economic situations. A buyer relies more on the recommendation of a friend than

*We are grateful to Larry Samuelson for support and encouragement throughout this project. We also thank Dirk Bergemann, Johannes Hörner and seminar participants at Yale and the 2015 World Congress of the Econometric Society for helpful discussions and suggestions. Finally, we wish to thank the anonymous referees for many valuable comments.

[†]Department of Economics, Yale University, chiara.margaria@yale.edu

[‡]Institute for Microeconomics, University of Bonn, alexey.v.smolin@gmail.com

that of a sales representative; an antitrust legislator takes arguments of firms opposing a newly proposed regulation with a grain of salt; a voter feels skeptical about the promises of a politician running for office. Nonetheless, in many settings in which interests are seemingly misaligned, informative communication is sustained: companies raise funds from investors, government agencies successfully split the state budget, conglomerates allocate resources among different divisions, and so on. In many of these cases, the informed parties have strong biases so that their payoffs are determined solely by the resulting decisions.¹ Importantly, all of these interactions happen repeatedly over time, with the parties trading off immediate opportunistic gains for the prospect of an ongoing relationship.

To address these settings, we analyze a dynamic version of an information transmission game in which private information (states) evolve stochastically over time. We allow for many senders but focus on the case in which the senders' payoffs are state independent. In every period, each sender sends a message to the receiver, who then takes a publicly observable action. The “cheap-talk” nature of messages is preserved: no hard evidence can be presented, the sender cannot commit to a communication strategy, and the receiver never observes extraneous information to test the validity of the past messages. No contracts can be written between the players, so at any point in time it must be in players' interests to follow the equilibrium play.

We obtain an analog of a folk theorem—that any feasible and individually rational payoffs can be approximated in equilibrium as the players become patient.² This payoff set, and hence the set of equilibrium payoffs, admits a simple characterization and includes all Pareto efficient payoffs that satisfy the receiver's individual rationality. Specifically, it includes the receiver's largest feasible “complete information” payoff. In equilibrium, the fraction of periods in which the receiver makes perfectly informed decisions can be arbitrarily close to one, even if no informative communication can be sustained in the stage game.

These results contrast with the conventional wisdom that state independence makes it harder to maintain informative communication. Indeed, in this case, each sender has an unambiguous ranking over actions and is willing to report truthfully only if indifferent among the messages she sends. However, what comes as a curse in a one-shot game turns into a blessing when the game is dynamic. If senders' payoffs depend only on actions and not on states, then the receiver fully observes and controls the payoffs. Our equilibrium construction actively uses this feature—it *targets* the senders' total payoffs and ensures they *do not depend* on the senders' messages. To achieve it, the equilibrium play alternates between communication and adjustment phases. In communication phases, the receiver makes informed decisions relying on the senders' messages. In adjustment phases, the receiver ignores all messages and plays according to a strategy that pulls the senders' payoffs towards the target. Specification of the phases and the transitions between them is tailored to guarantee that the senders' payoffs do not depend on their messages, and thus sustains truth-telling. By the law of large numbers, as players become more patient, the play occurs in the communication phase most of the time and any individually rational payoffs can be achieved by changing the strategy the receiver plays in there.

¹Chakraborty and Harbaugh (2010) provide many more examples of such strong biases.

²There is no standard formulation of a “folk theorem” in repeated games with incomplete information.

Related literature The general idea that an ongoing relationship can overcome strategic limits of interaction is the cornerstone of the literature on repeated games as discussed in depth by Mailath and Samuelson (2006). Further, the idea of using players' payoffs as a determinant of equilibrium construction is reminiscent of Abreu et al. (1990)'s recursive technique of equilibrium payoff decomposition. In fact, if individual states are independently and identically distributed, our dynamic game can be viewed as an infinitely repeated game and the Fudenberg et al. (1994)'s standard method can be used to provide an alternative proof of our folk theorem result. However, the standard method cannot be applied to general stochastic sender-receiver games.

At the same time, the idea of linking decisions motivated a strand of mechanism design literature. Jackson and Sonnenschein (2007) showed that Pareto efficient outcomes can be achieved by linking many identical copies of a collective choice problem with private values into a single mechanism. Frankel (2016) extended these ideas into a dynamic setting where the sender has persistent private information. He introduced discounted quota contracts similar to our equilibrium construction and showed their optimality in many environments. The mechanism design setting, however, differs from ours in that it endows the receiver with commitment power and monetary transfers.

Finally, the closest paper to ours is by Renault, Solan and Vieille (2013), who analyzed a dynamic information transmission game between a single sender and a receiver. They considered a more general payoff structure that, however, did not allow the sender's payoff to be state independent. In their setting, they showed that our analog of a folk theorem does not generally hold. In particular, the players' equilibrium payoffs do not necessarily approach the Pareto efficiency frontier as the players become patient.

Neither Renault et al. (2013)'s nor our proof can be directly extended to cover both the cases of state-dependent and of state-independent payoffs. On one hand, their construction is based on the idea of statistical tests that require the sender to match the message distribution with the state distribution. However, in our setting, because payoffs are state independent, statistical tests cannot be used effectively; faced with these tests, the sender, would induce favorable actions earlier in time independently of the realized states, rendering her messages uninformative. On the other hand, our construction is based on the ability of the receiver to adjust the senders' continuation payoffs exactly to the target without knowing the state. When the payoffs are state dependent, such adjustment must depend on the state and hence rely on the sender's messages running into the same incentive compatibility problem.³

The remainder of the paper proceeds as follows: Section 2 introduces the model, Section 3 characterizes the equilibrium payoffs and discusses, and Section 4 concludes. Appendix contains a detailed proof of the main theorem.

³Escobar and Toikka (2013) also studied limit payoffs with dynamic communication but their setting and techniques are significantly different from ours. They focus on the case of independent private values, so the players know their own payoffs. Their proof relies on statistical tests similar to those used by Renault et al. (2013) which are of limited use in our setting.

2 Model

A receiver (he) repeatedly communicates with n senders (she) indexed by $i \in N = \{1, \dots, n\}$, $n \geq 1$. The stage game is described by the set of states $\Omega \triangleq \times_{i \in N} \Omega_i$, their prior distribution $p \in \Delta(\Omega)$, the set of senders' messages $M \triangleq \times_{i \in N} M_i$, the set of actions of the receiver A , and the stage payoffs of the senders, $v : A \rightarrow \mathbb{R}^n$, and of the receiver $u : A \times \Omega \rightarrow \mathbb{R}$.⁴ Denote the stage game by $\Gamma \triangleq (N, \Omega, p, M, A, v, u)$. We assume that the sets Ω, M , and A are all finite. In the stage game, all senders first privately observe their individual states $\omega_i \in \Omega_i$ and then simultaneously send public messages $m_i \in M_i$ to the receiver who takes a publicly observable action $a \in A$.⁵

The stage game is infinitely repeated at times $t = 1, 2, 3, \dots$ with a common discount factor δ . The state profiles $\omega_t \triangleq \{\omega_{t1}, \dots, \omega_{tn}\}$ evolve according to an irreducible and aperiodic Markov chain with a transition kernel $k(\omega_{t+1} | \omega_t)$. Hence, the individual states can be arbitrarily correlated across senders but their inter-temporal correlation vanishes in time.⁶ Consequently, there exists a unique stationary distribution $p \in \Delta(\Omega)$ that has full support. We assume that the initial state is drawn according to this stationary distribution.

We assume that the validity of senders' messages can never be verified. First, no hard evidence is allowed—the sets of possible messages M_i do not depend on the states. Second, the senders cannot commit to communication strategies, in contrast to the Bayesian persuasion literature. Third, the receiver does not observe any additional information besides the senders' messages; for example, he does not observe any additional signals about past states.⁷ We allow the players to perfectly transmit their private information, $M_i = \Omega_i$, thus concentrating on strategic rather than technological limits of communication.⁸ Denote the message profile by $m_t \in M \triangleq \times_{i \in N} M_i$. Lastly, we assume that there is a public randomization device that produces a uniformly distributed output $y_t \in Y \triangleq [0, 1]$ at the beginning of each period independently of states, and all players publicly observe it. We do not require any public randomization at the interim stage, after the senders send their messages but before the receiver takes an action.

Overall timing within each period of the dynamic game is as follows: the public randomization device produces an output y_t ; the state profile ω_t realizes and each sender i privately observes her individual state ω_{it} ; the senders simultaneously send messages m_{ti} to the receiver; the messages are publicly observed and the receiver takes an action a_t ; the action is

⁴The single action of the receiver affects all senders' payoffs at once. Hence, the problem with many senders is not simply a collection of separate single-sender problems.

⁵We can dispense with the assumption that the messages are public if we allow the receiver to announce the messages at the end of each period.

⁶Notably, our analysis can be applied, with minor modifications, even if the state evolution is action-dependent, $k(\omega_{t+1} | \omega_t, a_t)$. This setting is of particular economic interest because it allows incorporating investment, growth, and precautionary measures. However, it is generally missing from the economic literature because of associated complications as highlighted by Hörner et al. (2015). Details are available upon request.

⁷This implies that the receiver does not observe his payoffs. This assumption is standard in repeated games with incomplete information and can be justified in at least two ways. First, it can approximate the situation when payoffs are observed far into the future. Second, the receiver may want to make informed decisions even if she cannot ever confirm their accuracy just like a judge wants to acquit innocents.

⁸This assumption is still *a priori* restrictive because the players cannot commit to their strategies and the revelation principle does not apply. However, it suffices to obtain a folk theorem.

publicly observed and the game proceeds to the next period. Denote the resulting dynamic game by $\Gamma^\infty(\delta)$.

Strategies The timing and monitoring structure of the game $\Gamma^\infty(\delta)$ outlined above result in the following definitions of histories and strategies. A public history h^t at time t consists of past actions, messages and the realizations of the randomization device.⁹ Denote the set of all public histories at time t by H^t . A behavioral strategy of the receiver, α , maps past public history, current output of the randomization device and current messages into an action,

$$h^t \triangleq \{y_s, m_s, a_s\}_{s=1}^{t-1}, \quad t > 1, \quad h^1 \triangleq \emptyset,$$

$$\alpha : \left(\bigcup_{t \geq 1} H^t \right) \times Y \times M \rightarrow A.$$

Senders' private histories h_i^t contain public histories as well as the individual states they observe. Denote the set of private histories of a sender i by H_i^t . A behavioral strategy of a sender i , μ_i , maps her private histories and current output of the randomization device into a message

$$h_i^t \triangleq h^t \cup \{\omega_{si}\}_{s=1}^t,$$

$$\mu_i : \left(\bigcup_{t \geq 1} H_i^t \right) \times Y \rightarrow M_i.$$

Payoffs Players' interests are misaligned. The receiver's stage payoff is state dependent, $u : A \times \Omega \rightarrow \mathbb{R}$, so that he generally prefers to take different actions in different states. The senders' stage payoffs, in contrast, depend only on the receiver's actions and not on the states, $v_i : A \rightarrow \mathbb{R}$. That is each sender had an optimal action that she prefers to be taken in every state. All players discount the future at a common rate $\delta < 1$. Given strategies $\mu \triangleq \{\mu_i\}_{i \in N}$ and α , senders' and receiver's expected normalized payoffs (or simply payoffs) can be written as

$$V_i(\mu, \alpha) = (1 - \delta) \mathbb{E}_{\mu, \alpha} \left[\sum_{t=1}^{\infty} \delta^{t-1} v_i(a_t) \right], \quad i \in N,$$

$$U(\mu, \alpha) = (1 - \delta) \mathbb{E}_{\mu, \alpha} \left[\sum_{t=1}^{\infty} \delta^{t-1} u(a_t, \omega_t) \right],$$

where the conditional expectations take into account the probability law induced by the strategies of the players, evolution of states, and the randomization device.

The set of feasible payoffs in the dynamic game $\Gamma^\infty(\delta)$, $\mathcal{F} \subseteq \mathbb{R}^{n+1}$, with a typical element (V, U) consists of all players' payoffs that can be achieved by some strategies μ, α . Because the first state is drawn from the stationary distribution, the states in all periods are ex ante

⁹We adopt a convention that for any stochastic process x its time- t realization is denoted by subscript x_t and the history up to time t is denoted by superscript $x^t \triangleq \{x_s\}_{s=1}^t$.

identically distributed, and the set of feasible payoffs in the repeated game coincides with the set of feasible payoffs in the stage game. This set is a polytope with finitely many vertices and admits a simple characterization: it is equal to the convex hull of the payoffs in the stage game resulting from all pure mappings from states into actions. Moreover, any payoff $(U, V) \in \mathcal{F}$ can be supported by a stationary strategy $\mathbf{a} : \Omega \times Y \rightarrow A$ played in every period.

$$\mathcal{F} = \text{co} \{ (\mathbb{E}_p v(\mathbf{a}(\omega)), \mathbb{E}_p u(\mathbf{a}(\omega), \omega)) \mid \mathbf{a} : \Omega \rightarrow A \}. \quad (1)$$

The set of individually rational payoffs, $\mathcal{F}^* \subseteq \mathcal{F}$, in a dynamic game $\Gamma^\infty(\delta)$ consists of all feasible payoffs such that all players get at least their minmax payoffs. The senders' minmax payoffs are defined as $\min_{\alpha, \mu_{-i}} \max_{\mu_i} V_i(\alpha, \mu)$ and the receivers' minmax payoff is defined as $\min_{\mu} \max_{\alpha} U(\alpha, \mu)$. Our payoff structure allows for a particularly tractable characterization of \mathcal{F}^* . As the receiver can always ignore the senders' messages his minmax payoff is $\underline{U} \triangleq \max_{a \in A} \mathbb{E}_p [u(a, \omega)]$. At the same time, because the receiver fully controls the senders' payoffs, their individual rationality is innocuous. As a result,

$$\mathcal{F}^* = \{(V, U) \in \mathcal{F} \mid U \geq \underline{U}\}. \quad (2)$$

Our equilibrium concept is a perfect Bayesian equilibrium as described by Fudenberg and Tirole (1991). A strategy profile (μ, α) , together with a system of beliefs, is a perfect Bayesian equilibrium if μ is sequentially rational and Bayes' rule is used to update beliefs whenever possible.¹⁰ Denote by $\mathcal{E}(\delta) \subseteq \mathbb{R}^{n+1}$ the set of equilibrium payoffs.

3 Folk Theorem

Our main result is a characterization of the limit set of equilibrium payoffs as players become arbitrarily patient, $\lim_{\delta \rightarrow 1} \mathcal{E}(\delta)$. We obtain an analog of a folk theorem—any feasible, individually rational payoff can be approximated in an equilibrium if players are sufficiently patient.

A few observations are immediate. First, no payoff outside of \mathcal{F}^* can be supported in equilibrium—the receiver can always guarantee at least his minmax payoff. Second, there are fully uninformative equilibria in the dynamic game Γ^∞ in which all senders babble; that is, they send messages independently of their states, and the receiver plays a myopic best-response. The set of babbling payoffs belongs to the set of equilibrium payoffs for all δ .

Assumption 1. (*Valuable Communication*) *There exists a vector $(V, U) \in \mathcal{F}^*$ such that $U > \underline{U}$.*

Assumption 1 states that communication is valuable—the receiver can strictly benefit from knowing the individual states. The assumption holds in most relevant economic environments. It is weaker than non-empty interior requirements in existing folk theorems, as it applies to the receiver only. This assumption allows to provide incentives for the receiver to follow the expected equilibrium play by threatening to cease the communication and permanently switch to babbling.

¹⁰We imposed the standard restriction (B) in Fudenberg and Tirole (1991); because types are correlated, condition (B) needs to be adjusted (see Fudenberg and Tirole, 1991, pp. 349-350).

Theorem 1. (*Folk Theorem*) For any payoffs $(V, U) \in \text{relint}(\mathcal{F}^*)$ there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ there is an equilibrium with payoffs (V, U) .¹¹ Consequently,

$$\lim_{\delta \rightarrow 1} \mathcal{E}(\delta) = \mathcal{F}^*.$$

The detailed proof is relegated to the Appendix. Here we briefly outline its main ideas. For any target payoff in \mathcal{F}^* we pick the distribution $\mathbf{a} : \Omega \times Y \rightarrow A$ that supports it in the stage game and construct an equilibrium in which senders report truthfully and the receiver plays according to \mathbf{a} most of the time.

In particular, the equilibrium play switches between communication and adjustment phases. Transition between phases is determined by the senders' accumulated payoffs, which depend only on past actions and thus are publicly observed. In the communication phase, each sender reports her individual state truthfully, and the receiver plays according to the stationary distribution \mathbf{a} . In the adjustment phase, the receiver plays according to a strategy that brings all senders' discounted payoffs back to the target. The length and strategy of the adjustment phase depend on the senders' payoffs at the end of preceding communication phase and can always be chosen to hit the target as long as the players are sufficiently patient. Any receiver's randomization is done via the public randomization device so his deviations are immediately observable and trigger permanent babbling play.

The exact specification of the phases and their transition is tailored to ensure the players' obedience in following the equilibrium play. The senders are willing to report truthfully because they are guaranteed to obtain the target payoffs irrespectively of their messages. The receiver is effectively deferred from deviating by the assumption of valuable communication if sufficiently patient.

As players become more patient, the length of the communication phase can be increased and we can appeal to the law of large numbers for Markov chains. Hence, on average, there is less adjustment to be made, so the play occurs in the communication phase most of the time. As a result, the equilibrium payoffs approach the target payoffs.

Example We illustrate the setting and the results in a standard resource allocation example. The receiver is a social planner who decides every period how to allocate an indivisible resource between two ex-ante symmetric regions. The senders are local representatives who privately observe the social values of allocating the resource to their region. They report the values to the planner at a regular meeting. The planner wants to put the resource to the best use and the representatives simply prefer having resource in their region. The social values are independent across regions and positively but imperfectly correlated across periods.

The example fits into our model by setting $n = 2$, $A = \{1, 2\}$, $\Omega_i = \{\omega^0, \omega^1\}$, $v_i(a) = \mathbf{1}(a = i)$, $u(a, \omega) = \omega_a$, $k(\omega_1, \omega_2' \mid \omega_1, \omega_2) = \rho$. Positive but imperfect correlation implies $\rho \in (1/2, 1)$ and corresponds to a uniform stationary distribution. We further normalize $\omega_i^0 = 0$ and $\omega_i^1 = 1$. The set of feasible and individually rational payoffs can be calculated by (1) and (2), and is presented in Figure 1. The set does not depend on the degree of correlation ρ and has an empty interior since the feasibility of allocation requires $V_1 + V_2 = 1$.

¹¹The relative interior of a set \mathcal{F}^* , $\text{relint}(\mathcal{F}^*)$, is its interior under the topology induced on the affine hull of \mathcal{F}^* . Hence the statement is meaningful even if \mathcal{F}^* has a dimension lower than the number of players.




Figure 1: The set of feasible payoffs, \mathcal{F} , with a typical element (V_1, V_2, U) . The shaded area indicates individually rational payoffs \mathcal{F}^* .

The assumption of valuable communication is satisfied. In the absence of additional information the planner is indifferent between allocating the resource to one region or another and obtains a payoff $U = 1/2$. However, the resource allocation matters to representatives. The “babbling” payoffs are generically Pareto inefficient and constitute a set

$$\mathcal{B} = \{(V, U) \mid V_1 + V_2 = 1, U = 1/2\}.$$

If the stage game were played only once, then strategic considerations would restrain communication and the babbling payoffs would be a unique equilibrium payoffs. The argument is standard. If the planner were to rely on a representative’s report then the representative would communicate the report that maximizes a probability of getting the resource irrespectively of its social value. This would in turn make the report uninformative and preclude payoffs outside of \mathcal{B} .

However, in a dynamic game any payoffs in \mathcal{F}^* can be approximated in an equilibrium if the players are sufficiently patient as shown in Theorem 1. According to our equilibrium construction the planner should simply ensure the representatives of the total discounted allocation to their regions irrespectively of their reports. It can be achieved by infrequently shutting down the meetings and bringing the total discounted allocation back to the targets. In this way, the resource is allocated efficiently in a fraction of periods arbitrarily close to one.

To recapitulate, in the absence of any contract enforcement or message verification, almost fully informed decision-making can be sustained in a dynamic game even if no informative communication can be sustained in the stage game. Moreover, the result does not require the payoff set to have a non-empty interior in contrast to most of existing folk theorems.

Discussion Before concluding, we discuss two important assumptions of our model that are necessary for our result: ergodicity of state transitions and state-independence of senders’ payoffs.

It is crucial for the folk theorem result that the Markov chain according to which individual states evolve is aperiodic and irreducible. This assumption ensures that any information that the senders have is transient; even though the states can be correlated across periods, this correlation vanishes as the distance between them grows. It gives the game its recurrent structure and allows to split the equilibrium play in “almost independent” blocks. As players become patient, it allows to achieve in equilibrium two seemingly contradictory goals: to preclude the senders from affecting their payoffs and to make the receiver’s action responsive to their messages.

A clear case in which this assumption is violated is when the state is perfectly persistent. It is easy to see that in this case our folk theorem result does not need to hold even if the receiver could commit to his strategy. Since the state is fixed throughout the game for any discount factor the problem is equivalent to a static problem and the senders’ incentive constraints may preclude efficiency.

Another important assumption of our model is that the senders’ payoffs do not depend on the state. One might think that our equilibrium construction might still work if the payoffs are “almost” state-independent. Unfortunately, this is not the case. In our equilibrium, the senders are indifferent between any history of messages and report truthfully in equilibrium. If the payoffs are even slightly state-dependent then the senders’ best-response can be far from truthtelling.

In fact, for a generic payoff perturbation one can apply characterization of Renault et al. (2013) to calculate the limit set of equilibrium payoffs with a single sender. They show that the equilibrium payoffs must satisfy senders’ “incentive compatibility”—she shouldn’t be able to benefit from permuting her reports. This constraint is ordinal and even small payoff perturbations can drastically restrict the equilibrium payoff set.

4 Conclusion

We analyzed dynamic information transmission when senders’ payoffs are state independent. We show that the strategic limits of communication prevalent in a stage-game disappear when the interaction is repeated. Any individually rational payoffs can be approximated in equilibrium if the players are sufficiently patient. This result complements the existing results for state-dependent payoffs and provides a rationale for informative communication to be sustained in a variety of dynamic economic settings between players with seemingly misaligned interests. In fact, our equilibrium construction delivers a clear message how to do so—to induce truthtelling the receiver should track the senders’ payoff and adjust it whenever they are doing too good or too bad. This ensures the senders of their payoffs irrespective of their reports and eliminates incentives to lie.

There are many opportunities for further research. First, the limit analysis for state-dependent payoffs not captured by the previous literature should be completed. Second, one can analyze the joint limit of state independency and patience. Finally, an important open question is the general analysis of equilibrium payoffs when players are impatient. We suspect that in this case finer details of Markov transition and not just its stationary distribution will come into play.

Appendix

Proof of Theorem 1

Without loss of generality, we normalize players' payoffs so that $\min_A v_i(a) = 0, \forall i \in N$, and $\min_{(a,\omega) \in A \times \Omega} u(a, \omega) = 0$. Let $\bar{v}_i \triangleq \max_{a \in A} v_i(a)$, and $\bar{u} \triangleq \max_{(a,\omega) \in A \times \Omega} u(a, \omega)$. Recall that

$$\mathcal{F}^* = \{(V, U) \in \mathcal{F} \mid U \geq \underline{U}\}.$$

where $\underline{U} = \max_{a \in A} \mathbb{E}_p[u(a, \omega)]$. Given any feasible and individually rational payoffs, we construct an equilibrium that achieves payoffs arbitrarily close to it provided that δ is sufficiently high. In particular, define

$$\mathcal{F}^{**} \triangleq \{(V, U) \in \text{relint}(\mathcal{F}^*)\}.$$

and let $|\cdot|$ denote the Euclidean norm. For any $(V, U) \in \mathcal{F}^{**}$ there exists $\eta > 0$ such that $\forall V' : |V' - V| \leq \eta, V' \in \mathcal{F}^{**}$. We will show that for any $\varepsilon > 0$ and any $(V, U) \in \mathcal{F}^{**}$, there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ we can construct an equilibrium (μ^*, α^*) of $\Gamma^\infty(\delta)$ with payoffs (V^*, U^*) , where $V^* = V$ and $|U^* - U| < \varepsilon$. Since we allow for a public randomization device, the result of the theorem will follow.

In what follows, fix $\delta < 1, \eta > 0$, the target payoff vector $(V, U) \in \mathcal{F}^{**}$, and the corresponding supporting strategy $\mathbf{a} : \Omega \times Y \rightarrow A$.

Strategies We first describe the equilibrium strategies (μ^*, α^*) . Fix some $T_c \in \mathbb{N}$ as a function of δ such that $\lim_{\delta \rightarrow 1} \delta^{T_c(\delta)} = 1, \lim_{\delta \rightarrow 1} T_c(\delta) = \infty$.¹² The equilibrium play is divided into consecutive blocks, each starting with a communication phase of length T_c followed by an adjustment phase of (random) length T_a .

On the equilibrium path, the behavior within each block is described by strategies (μ_1, α_1) . According to μ_1 , the senders always report truthfully. According to α_1 , the receiver's behavior depends on the current phase within a block. In the communication phase, he plays according to the strategy \mathbf{a} . In the adjustment phase, he ignores the senders' reports and plays according to an adjustment strategy that depends on the profile of senders' normalized discounted payoffs at the end of the communication phase,

$$\bar{v}_c \triangleq \frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=1}^{T_c} \delta^{t-1} v(a_t),$$

and the target payoff vector V . In particular, define

$$\lambda(\bar{v}_c) \triangleq \frac{\bar{v}_c - V}{|\bar{v}_c - V|}, \quad V_a(\bar{v}_c) \triangleq V - \lambda(\bar{v}_c)\eta.$$

Let $\hat{T}_a \in \mathbb{R}_+$ satisfy

$$\delta^{T_c} \left(1 - \delta^{\hat{T}_a}\right) \eta = (1 - \delta^{T_c}) |\bar{v}_c - V|. \quad (3)$$

¹²For example, $T_c(\delta) = \lceil (1 - \delta)^{-1/2} \rceil$. In what follows we will often omit dependence of T_c on δ .

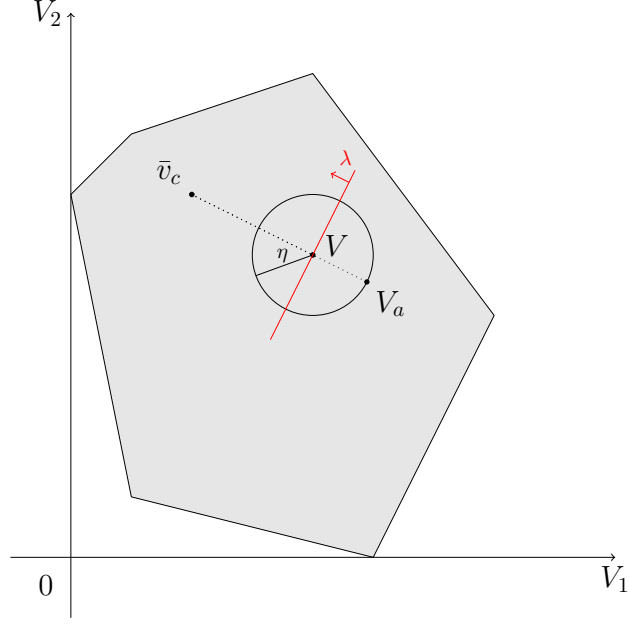


Figure 2: The choice of adjustment payoffs V_a given the realized senders' payoff at the end of the communication phase \bar{v}_c and the target equilibrium payoffs V . Schematic illustration for the case of two senders, $n = 2$.

Notice that \hat{T}_a is well defined as long as $\eta > (1 - \delta^{T_c})\delta^{-T_c} \sum_{n=1}^N \bar{v}_i$, which, in light of our choice of δ and $T_c(\delta)$, is verified for δ high enough. Define $T_a \triangleq \lceil \hat{T}_a \rceil + 1$, and let $0 \leq r(\bar{v}_c) < \eta$ solve

$$\frac{1 - \delta}{1 - \delta^{T_c + T_a}} \left(\sum_{t=1}^{T_c} \delta^{t-1} v(a_t) + \sum_{t=T_c+1}^{T_c+T_a-1} \delta^{t-1} V_a(\bar{v}_c) + \delta^{T_c+T_a-1} (V - \lambda(\bar{v}_c)r(\bar{v}_c)) \right) = V,$$

The adjustment phase lasts T_a periods. In the first $T_a - 1$ periods, the possibly random action $\mathbf{a}_a \in \Delta A$ is played with $V_a(\bar{v}_c) = \mathbb{E}v(\mathbf{a}_a)$. In the last period, the possibly random action $\mathbf{a}_r \in \Delta A$ is played with $V - \lambda(\bar{v}_c)r = \mathbb{E}v(\mathbf{a}_r)$.

If the receiver ever deviated, then (μ^*, α^*) prescribes the senders to babble; that is, to send messages irrespectively of their individual states, and the receiver to play an action that is optimal given no information.

As for the deviation of each sender, because individual states are correlated, some deviations during the communication phase are detectable as they lead to an inconsistent sequence of messages. During the communication phase, whenever the current messages m_t are inconsistent with the messages from the previous period m_{t-1} the receiver plays \mathbf{a} after replacing the reports with an artificially generated state $\hat{\omega} \in \Omega$ consistent with the previous messages and the underlying Markov chain. The strategy μ^* prescribes each sender to keep playing according to μ_1 after any detected deviation of another sender, and if she ever privately deviated.

Finally, transitions between and within blocks are as follows. Each block starts with a communication phase that lasts for T_c periods and is followed by the adjustment phase,

which lasts for T_a periods according to (3). At period $T_a + 1$, a new block starts and the timer is reset to $t = 1$.

Payoffs By construction, the equilibrium strategies deliver expected continuation payoffs V to the senders at the beginning of each block and, in particular, at the beginning of a game. We show that for sufficiently high δ , the equilibrium strategies deliver the expected continuation payoff within ε of the target payoff U to the receiver after any public history.

First, we show the expected continuation payoff of the receiver is within ε of the target payoff U at the beginning of each block. Consider the average realized payoffs within the each block at the end of communication and adjustments phases:

$$\begin{aligned}\bar{u}_c &\triangleq \frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=1}^{T_c} \delta^{t-1} u(a_t, \omega_t), \\ \bar{u}_a &\triangleq \frac{1 - \delta}{1 - \delta^{T_c + T_a}} \sum_{t=1}^{T_c + T_a} \delta^{t-1} u(a_t, \omega_t),\end{aligned}$$

so that $\bar{u}_c, \bar{u}_a \in \text{proj}_{n+1} \mathcal{F}^*$ and $U^* = \mathbb{E}_p[\bar{u}_a]$, $U = \mathbb{E}_p[\bar{u}_c]$.

We can bound the difference between \bar{u}_a and \bar{u}_c as follows:

$$\begin{aligned}\bar{u}_a - \bar{u}_c &= \frac{\delta^{T_c} (1 - \delta^{T_a})}{1 - \delta^{T_c + T_a}} \left(\frac{1 - \delta}{1 - \delta^{T_a}} \sum_{t=T_c+1}^{T_c+T_a} \delta^{t-T_c-1} u(a_t, \omega_t) - \bar{u}_c \right), \\ |\bar{u}_a - \bar{u}_c| &\leq \frac{\delta^{T_c} (1 - \delta^{T_a})}{1 - \delta^{T_c + T_a}} \bar{u} \leq \frac{\delta |\bar{v}_c - V| + o(1 - \delta)\eta}{\delta |\bar{v}_c - V| + (1 + o(1 - \delta))\eta}\end{aligned}\quad (4)$$

where the second equality follows by (3) and the fact that $\delta^{T_c} \rightarrow 1$ as $\delta \rightarrow 1$. By construction $\mathbb{E}_p \bar{v}_c = V^* = V$. We now argue that as δ goes to 1, the difference $|V - \bar{v}_c|$ goes to 0 almost surely independently of starting distribution. To show this, note that for any fixed T_c the random variable $\bar{v}_c = \frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=1}^{T_c} \delta^{t-1} v(a(\omega_t, y_t))$ viewed as a statistic test has bounded differences, that is

$$\begin{aligned}\frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=1}^{T_c} \delta^{t-1} v(a(\omega_t, y_t)) - \left(\frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=1}^{\tau-1} \delta^{t-1} v(a(\omega_t, y_t)) + \frac{1 - \delta}{1 - \delta^{T_c}} \sum_{t=\tau+1}^{T_c} \delta^{t-1} v(a(\omega_t, y_t)) + \right. \\ \left. \frac{1 - \delta}{1 - \delta^{T_c}} \delta^{\tau-1} v_i(a(\omega_t, y_t)) \right) \leq \frac{1 - \delta}{1 - \delta^{T_c}} \delta^{\tau-1} \max_{i \in N} \bar{v}_i \quad \forall \tau = 1, \dots, T_c - 1\end{aligned}$$

It follows by Paulin (2015) who generalized the McDiarmid (1989)'s concentration inequality from independent random variables to Markov chains that

$$\Pr(|V - \bar{v}_c| > \varepsilon) \leq 2 \exp \left(- \frac{2\varepsilon^2}{9t^{\text{mix}} (\max_{i \in N} \bar{v}_i)^2} \frac{(1 + \delta) \sum_{t=1}^{T_c} \delta^{t-1}}{(1 + \delta^{T_c})} \right), \quad (5)$$

where t^{mix} denotes the mixing time of the corresponding Markov chain. Since the states evolve according to an aperiodic Markov chain, t^{mix} is finite.¹³ Consequently, because

¹³The mixing time is finite if and only if the chain is uniformly ergodic. Any aperiodic and irreducible Markov chain on finite state space is uniformly ergodic.

$T_c(\delta) \rightarrow \infty$ as $\delta \rightarrow 1$ by (4) and (5), for any $\varepsilon > 0$ there is $\underline{\delta}_1$ such that for all $\delta > \underline{\delta}_1$,¹⁴

$$|U^* - U| = |\mathbb{E}_p \bar{u}_a - \mathbb{E}_p \bar{u}_c| < \varepsilon.$$

We conclude the section by showing the receiver's continuation payoff after *any* public history is close to U^* and hence to U . Since the receiver's stage payoffs are bounded by \bar{u} and the equilibrium play is reset in every block, it suffices to show that (i) a contribution of a single block to the total payoff is negligible and (ii) the impact of current beliefs on payoffs in future blocks is negligible. For (i), observe that by (3)

$$(1 - \delta^{T_a - 1})\eta \leq \frac{1 - \delta^{T_c}}{\delta^{T_c}} |V - \bar{v}_c| \leq \frac{1 - \delta^{T_c}}{\delta^{T_c}} n \max_{i \in N} \bar{v}_i.$$

Consequently, $\delta^{T_a - 1} \rightarrow 1$ as $\delta \rightarrow 1$ and $\delta^{T_c} \rightarrow 1$, so the contribution of a single block, weighted by $1 - \delta^{T_c + T_a}$, goes to 0 as $\delta \rightarrow 1$. For (ii), observe that we can apply the same argument of variational inequalities, which does not require any assumption on the initial distribution, to show that \bar{u}_a converges to U^* almost surely as $\delta \rightarrow 1$ for any starting belief.

Incentives We now check players' incentives to follow the suggested strategies (μ^*, α^*) ; that is, that they are sequentially rational. Consider senders' incentives. Because babbling is an equilibrium of a stage game, the off-path behavior $(\mu_\emptyset, \alpha_\emptyset)$ is sequentially rational. Also, since the senders' expected continuation payoffs are independent of their messages, they have no incentives to deviate. Thus, μ_1 is sequentially rational. Consider the receiver's incentives. The receiver's continuation payoff following a deviation at time t equals

$$\bar{u}_B \triangleq (1 - \delta) \sum_{\tau=1}^{\infty} \delta^{\tau-1} \mathbb{E}_{\pi_t, \alpha'} [u(a_t, \omega_\tau)]$$

where π_t is the receiver's belief at time t and α' is his best-reply strategy during the punishment. By the same argument as before, the receiver can benefit from the current information only in the short run because for any starting belief \bar{u}_B converges to \bar{u} almost surely as $\delta \rightarrow 1$. As we showed before, the equilibrium continuation payoff after any history is in the neighborhood of U . By assumption of valuable communication, $U > \underline{U}$ as any other payoff in \mathcal{F}^{**} . Hence, the receiver is willing to obey if sufficiently patient facing the threat of switching to babbling. It follows that α_1 is sequentially rational.

To complete the proof, we specify the system of beliefs. Bayes rule uniquely pins down players' belief, unless a sequence of inconsistent messages are reported. Recall that when the current messages m_t are inconsistent with the messages from the previous period m_{t-1} , the receiver takes action $\mathbf{a}(\hat{\omega})$, where $\hat{\omega}$ is an artificially generated state. At these histories, the receiver updates his belief according to Bayes rule as if the reported state was $\hat{\omega}$. The belief of the deviator is computed by Bayes rule. The belief of any other sender i equals the posterior belief based only on her private information ω_i^t .

¹⁴In fact, the proof shows a stronger statement, $\bar{u}_a \rightarrow_{a.s.} \bar{u}_c$ as $\delta \rightarrow 1$.

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990) “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- Chakraborty, Archishman and Rick Harbaugh (2010) “Persuasion by Cheap Talk,” *American Economic Review*, 100, 2361–2382.
- Crawford, Vincent P. and Joel Sobel (1982) “Strategic Information Transmission,” *Econometrica*, 50, 1431–1451.
- Escobar, Juan F. and Juuso Toikka (2013) “Efficiency in games with Markovian private information,” *Econometrica*, 81, 1887–1934.
- Frankel, Alexander (2016) “Discounted quotas,” *Journal of Economic Theory*, 166, 396–444.
- Fudenberg, Drew, David Levine, and Eric Maskin (1994) “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039.
- Fudenberg, Drew and Jean Tirole (1991) *Game Theory*: MIT Press.
- Hörner, Johannes, Satoru Takahashi, and Nicolas Vieille (2015) “Truthful Equilibria in Dynamic Bayesian Games,” *Econometrica*, 83, 1795–1848.
- Jackson, Matthew O. and Hugo F. Sonnenschein (2007) “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75, 241–257.
- Mailath, George J. and Larry Samuelson (2006) *Repeated Games and Reputations: Long-Run Relationships*, New York: Oxford University Press.
- McDiarmid, Colin (1989) “On the method of bounded differences,” *Surveys in combinatorics*, 141, 148–188.
- Paulin, Daniel (2015) “Concentration inequalities for Markov chains by Marton couplings and spectral methods,” *Electronic Journal of Probability*, 20.
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille (2013) “Dynamic sender-receiver games,” *Journal of Economic Theory*, 148, 502–534.