Theoretical Economics 8 (2013), 95–123 1555-7561/20130095

Implementation of communication equilibria by correlated cheap talk: The two-player case

Péter Vida
Department of Economics, University of Mannheim

Françoise Forges
Université Paris-Dauphine and Institut Universitaire de France

We show that essentially every communication equilibrium of any finite Bayesian game with two players can be implemented as a strategic form correlated equilibrium of an extended game, in which before choosing actions as in the Bayesian game, the players engage in a possibly infinitely long (but in equilibrium almost surely finite), direct, cheap talk.

Keywords. Bayesian game, cheap talk, communication equilibrium, correlated equilibrium, preplay communication.

JEL classification. C72, D70.

1. Introduction

Consider a standard Bayesian game in which the players simultaneously choose actions as a function of their type. Assume that, before making their decision, the players can exchange as many costless messages as they wish, possibly with the help of a mediator. A generalized revelation principle holds: the set of all Nash equilibrium outcomes of all games that extend the Bayesian game by allowing such communication is characterized as the set of all canonical communication equilibria (see Forges 1986 and Myerson 1986, 1991, Chapter 6). Canonical communication equilibria are very tractable but rely on a specific kind of mediator, who invites every player to fully reveal his type and then privately recommends an action to every player. Plain conversation between the players is much more natural and does not require a benevolent third party who becomes aware of the players' private information. Hence the following question:

Péter Vida: vidapet@gmail.com
Françoise Forges: Francoise.Forges@gmail.com
Péter Vida proved the main result of this paper as part of his doctoral thesis at the Universitat Autònoma de Barcelona. A preliminary version of the paper, by Péter Vida, circulated under the title "From Communication Equilibria to Correlated Equilibria." The two authors had discussions on this research in 2005–2006 and completed the present version in 2010–2011. Péter Vida's work on the paper was completed while he was employed by the University of Vienna. We wish to thank Elchanan Ben-Porath, Antoni Calvó-Armengol, Olivier Gossner, Manfred Nermuth, Eilon Solan, Adam Szeidl, and the participants of the workshop "Decentralized Mechanism Design, Distributed Computing and Cryptography" held at Princeton in June 2010. Finally, three referees and, especially, a co-editor of *Theoretical Economics* raised interesting questions that led us to establish additional results.

Copyright © 2013 Péter Vida and Françoise Forges. Licensed under the Creative Commons Attribution-NonCommercial License 3.0. Available at http://econtheory.org.
DOI: 10.3982/TE944

Can all canonical communication equilibrium outcomes be implemented by means of cheap talk, i.e., as Nash equilibrium outcomes of an appropriately designed extended game in which the players can just talk to each other before making their decision?

Partially or even fully positive answers have been given in finite Bayesian games with at least three players.[1] However, for two players, the answer is, in general, negative. For instance, in finite sender–receiver games (namely, two-person Bayesian game with a single informed player, whose actions are not payoff relevant), Forges (1985, 1990a) shows that there may exist communication equilibrium outcomes that cannot be implemented as Nash equilibrium outcomes of any long cheap talk game.[2]

In this paper, we show that the communication equilibrium outcomes of any two-person game can be implemented as *correlated* equilibrium outcomes (in the sense of Aumann 1974, 1987) of a long cheap talk game. In other words, we assume that, before they start to talk, the players can privately observe some signal, a sunspot, that is totally extraneous to the game (i.e., independent of the players' types and without any direct effect on the payoffs).[3] The players are not subject to any deadline and cannot use any common device (like urns, envelopes, or recording machines) while they talk (but each player is of course free to use any personal device to make his own choices). In a tractable representation, the signal of each player before the talking phase is a recommendation on how to talk and how to make a decision at the end of the cheap talk phase.

Our main result can be stated as follows. Fix any finite two-person Bayesian game Γ and any (strictly individually rational) communication equilibrium outcome of Γ. We design a long cheap talk extension $\operatorname{ext}\Gamma$ of Γ, with finitely many messages at every stage, together with a correlation device for the cheap talk game $\operatorname{ext}\Gamma$, with the following properties: (i) no player can gain by unilaterally deviating from the recommendation of the correlation device in $\operatorname{ext}\Gamma$ and (ii) the outcome, namely the conditional probability distributions generated by the correlation device and strategies in $\operatorname{ext}\Gamma$ over actions given types, are exactly the same as in the communication equilibrium. In this construction, the size of the finite set of messages depends on the parameters of the Bayesian game and on the underlying communication equilibrium. By considering a countable set of messages, we can get at once all (strictly individually rational) communication equilibrium outcomes of any Bayesian game as correlated equilibrium outcomes of a *universal* cheap talk game, as in Forges (1990b) for games with at least four players.[4] Our cheap

[1] Game theoretical references involve, e.g., Bárány (1992), Forges (1990b), Ben-Porath (1998, 2003, 2006) and Gerardi (2004), Forges (2009) provides a survey. See, e.g., Halpern (2008) for references in computer science.

[2] As another particular case, if every player has a single type (complete information), communication equilibria coincide with Aumann's (1974, 1987) correlated equilibria, but in a plain conversation, both players know all the messages that they exchange, so that they cannot simulate the private recommendations of a mediator.

[3] As in Forges (1988), we do not reserve the term "sunspot" to a common, public, extraneous signal. The interpretation is that every player observes the sunspots in his own way.

[4] Forges (1990b) also proposes a cheap talk game with a *continuum* of messages that is universal for all three-person games.

talk game ext Γ is possibly infinitely long in the sense that its length is not fixed in advance, in a deterministic way, but depends endogenously on the messages exchanged by the players.

Our result extends Forges (1985), which focuses on the case of a single informed player and a single decision maker. One stage of cheap talk suffices then to implement all communication equilibrium outcomes. Recently, Blume (2010) established a similar result in the context of Crawford and Sobel's (1982) sender–receiver game.[5] Forges's (1985) construction goes through if payoff relevant actions are added for the single informed player. However, the general case, where *both* players are privately informed *and* make decisions, remains open until Vida (2007b) proposes a first solution.[6]

Some papers propose a different approach to implement communication equilibrium outcomes in two-person games without the help of a mediator. Dodis et al. (2000) and Urbano and Vila (2002) assume that the computational ability of the players is limited. Under this assumption, they show that the correlated equilibrium outcomes of any two-person game with complete information can indeed be implemented as (ϵ-) Nash equilibrium outcomes of a cheap talk extension of the game. Ben-Porath (1998) does not rely on any cryptographic tool, but obtains a similar result by allowing the players to make use of urns or envelopes while they talk. Generalizations to games with incomplete information are proposed by Krishna (2007) and Izmalkov et al. (2011) for the latter approach, and by Urbano and Vila (2004) for the cryptographic approach. The common feature of these solutions (as opposed to ours) is that *at every stage*, cheap talk is relaxed in some way: limited computational ability or physical hard devices are used to exchange messages at every stage.

When trying to implement a given communication equilibrium by correlated cheap talk in a two-person game in which both players have private information and must take actions, the main problem is to guarantee that no player learns useful information before the other. Full detection of possible deviations during the cheap talk phase can be of no help if it happens too late. Indeed, there may be no way to "punish" a deviator once he possesses the desired information. To solve the problem, the basic idea is that the correlation device selects a relevant stage t^* of the cheap talk phase, without telling it directly to the players. How do the players figure out when they reach it? At the end of every stage t of cheap talk, they simultaneously discover *from their exchanged messages* whether stage t was relevant (i.e., $t = t^*$). Useful information is exchanged only at the relevant stage t^*, but the players realize this at the end of the stage. In addition, at every stage, each player can check whether the other's message was legitimate. If the stage is not relevant, the players' information is not updated so that illegitimate messages can give rise to punishments.

As indicated in the previous paragraphs, our construction makes use of possibly infinitely long cheap talk. To what extent does our implementation result rely on an infinite horizon? We prove that, at the correlated equilibrium that implements a given

[5] Crawford and Sobel's (1982) model involves types and actions in a real interval and thus does not pertain to the finite setup of this paper.

[6] The main result in this paper can already be found in Vida's (2007b) unpublished doctoral dissertation (see also Vida 2007a). The proof proposed in this paper is a simplification of the original.

communication equilibrium, cheap talk lasts for finitely many stages almost surely. We also propose an example (in Section 5) in which an efficient communication equilibrium outcome cannot be achieved as a correlated equilibrium outcome, in any cheap talk game with a bounded number of stages. Nonetheless, if the underlying game Γ has a strictly individually rational Bayesian-Nash equilibrium (a condition that holds in the previous example), every (strictly individually rational) communication equilibrium payoff can be *approximated* by a correlated equilibrium payoff of a *sufficiently long* cheap talk game (Proposition 1 in Section 4).

The brief sketch above also suggests that our construction makes use of punishments. This should not be surprising in view of the literature on the implementation of a mediator by cheap talk (see Bárány 1992 for an early example, Heller et al. 2012 for a recent one, Ben-Porath 2003, 2006 and Gerardi 2004 for discussions and solutions to the problem, and Forges 2009 for a survey). Are these punishments credible? To formulate this question more precisely, recall that, according to our result, every communication equilibrium of a Bayesian game Γ can be implemented as a *Nash* equilibrium of an extended game $(\text{ext}\,\Gamma)^\mu$, in which a correlation device μ sends private signals to the players before they talk. Can we refine this Nash equilibrium so as to capture the players' sequential rationality, e.g., into a perfect Bayesian equilibrium? Proposition 2 (Section 4) gives sufficient conditions for a positive answer.

We implement communication equilibria of a given Bayesian game as correlated equilibria of the game preceded by cheap talk. Hence we replace the communication device by a correlation device, that is to say, replace a mediator by another! What do we really gain from our construction? As argued by Forges (1985, 1988, 1990b) and recently by Blume (2010), the mediators implicitly involved in the two solution concepts are very different from each other. In a (canonical) communication equilibrium of the original game, the mediator gets to know the whole information of every player. However, in a correlated equilibrium of the cheap talk game, the mediator does not receive *any* information from the players. He makes recommendations on how to exchange messages, but remains fully ignorant of the players' types. With such a mediator, players can preserve their privacy.

Let us turn to the organization of the paper. In the next section, we recall the concepts of Bayesian game and communication equilibrium. Then, in Section 3, we describe the extension of the game in which the players can talk and we define correlated equilibrium in that game. In Section 4, we state the main result as Theorem 1 and the two propositions mentioned above; the reader familiar with our basic concepts can go to the statements right away. Section 5 is devoted to an example that illustrates our results. Section 6 contains the proofs. Finally, Section 7 discusses some variants of the model.

2. Basic game: Communication equilibrium

Let us fix a two-player finite Bayesian game $\Gamma \equiv \langle \{L^i, A^i, g^i\}_{i=1,2}, p\rangle$: for every player $i = 1, 2$, L^i is a finite set of possible types, A^i is a finite set of actions, and $g^i : L \times A \to \mathbb{R}$ is a von Neumann–Morgenstern utility function, where $L = L^1 \times L^2$ and $A = A^1 \times A^2$; $p \in \Delta L$ is the players' common prior over L.

Γ starts with a move of nature, which selects $l = (l^1, l^2) \in L$ according to p. Player i is informed only of his own type l^i, $i = 1, 2$. Then the players simultaneously choose actions $a^1 \in A^1$ and $a^2 \in A^2$, respectively. Let $a = (a^1, a^2)$. The respective payoffs are $g^1(l, a)$ and $g^2(l, a)$.

A (canonical) communication device[7] q for Γ is a transition probability from L to A, $q : L \to \Delta A$, namely a system of probability distributions $q(\cdot|l)$ over A for every $l \in L$. By adding a communication device q to the Bayesian game, one generates an extended game Γ^q, which is played as follows.

1. Every player i learns his type l^i as in Γ, $i = 1, 2$.
2. Every player i sends a private message $\hat{l}^i \in L^i$ to the communication device q. Let $\hat{l} = (\hat{l}^1, \hat{l}^2)$.
3. The device q selects an action profile $a = (a^1, a^2)$ with probability $q(a|\hat{l})$.
4. The device q sends a^i privately to player i, $i = 1, 2$.
5. The players choose actions and receive payoffs as in Γ.

Some strategies are of special interest in Γ^q: player i is *sincere* in Γ^q if he reveals his type to the communication device at stage 2, namely $\hat{l}^i = l^i$ for every $l^i \in L^i$; player i is *obedient* if at stage 5, he follows the recommendation a^i made by the communication device at stage 4, whatever his type. When both players are sincere and obedient, the expected payoff of player i of type l^i is[8]

$$G^i[q|l^i] = \sum_{l^{-i}} p(l^{-i}|l^i) \sum_a q(a|l^i, l^{-i}) g^i((l^i, l^{-i}), a), \quad l^i \in L^i, i = 1, 2. \tag{1}$$

Let $G[q] = (G^i[q|l^i])_{l^i \in L^i, i=1,2}$ be the pair of vector payoffs associated with q.

Definition 1. Let q be a (canonical) communication device for Γ. q is a (canonical) communication equilibrium of Γ if and only if the sincere and obedient strategies form a Nash equilibrium of Γ^q, namely, if and only if

$$G^i[q|l^i] \geq \sum_{l^{-i}} p(l^{-i}|l^i) \sum_{a^i, a^{-i}} q(a^i, a^{-i}|\hat{l}^i, l^{-i}) g^i((l^i, l^{-i}), r^i(a^i), a^{-i})$$

for $i = 1, 2$, $l^i, \hat{l}^i \in L^i$ and for all $r^i : A^i \to A^i$. ME(Γ) denotes the set of communication equilibrium[9] payoffs of Γ, namely

$$\text{ME}(\Gamma) = \{G[q] \mid q \text{ is a communication equilibrium in } \Gamma\} \subset \mathbb{R}^{|L^1|+|L^2|}.$$

[7]See Forges (1986, 1990b) and Myerson (1986, 1991).

[8]When the index i refers to one of the two players, $-i$ refers to the other one.

[9]We use the notation ME as a reminder of "mediated equilibrium"; we keep CE for "correlated equilibrium."

Thanks to the general *revelation principle* recalled in the Introduction (see, e.g., Forges 1990b), ME(Γ) is the set of all payoffs that can be achieved at a Nash equilibrium of an arbitrary extension of Γ allowing the players to communicate (possibly with infinitely many stages and relying on a mediator at every stage).

Definition 2. A payoff vector $(x^i(l^i))_{l^i \in L^i} \in \mathbb{R}^{|L^i|}$ is (strictly) interim individually rational for player $i = 1, 2$ (or interim supportable with (strict) punishment) in Γ if there is a strategy of the other player in Γ, namely a transition probability $y^{-i}: L^{-i} \to \Delta A^{-i}$, such that for all $l^i \in L^i$,

$$x^i(l^i) \geq (>) \max_{a^i \in A^i} \sum_{l^{-i}} p(l^{-i}|l^i) \sum_{a^{-i}} y^{-i}(a^{-i}|l^{-i}) g^i((l^i, l^{-i}), a^i, a^{-i}).$$

(S)INTIR(Γ) denotes the set of vectors in $\mathbb{R}^{|L^1|+|L^2|}$ that are (strictly) interim individually rational for both players.

Observe that, in general, (S)INTIR(Γ) depends on the prior probability distribution p in Γ. In games with complete information (i.e., when $|L^1| = |L^2| = 1$), the definition reduces to the standard one, namely x^i is (strictly) individually rational for player i if and only if

$$x^i \geq (>) \min_{y^{-i} \in \Delta A^{-i}} \max_{a^i \in A^i} \sum_{a^{-i}} y^{-i}(a^{-i}) g^i(a^i, a^{-i}).$$

The following lemma, which is used later, states that interim individual rationality always holds at a communication equilibrium.

Lemma 1. ME(Γ) ⊆ INTIR(Γ).

Proof. Let q be a communication equilibrium and let $l^i \in L^i$ be a type of player i. For any $b^i \in A^i$ and $\hat{l}^i \in L^i$,

$$\begin{aligned} G^i[q|l^i] &= \sum_{l^{-i}} p(l^{-i}|l^i) \sum_{a} q(a|l^i, l^{-i}) g^i((l^i, l^{-i}), a) \\ &\geq \sum_{l^{-i}} p(l^{-i}|l^i) \sum_{a^i, a^{-i}} q(a^i, a^{-i}|\hat{l}^i, l^{-i}) g^i((l^i, l^{-i}), (b^i, a^{-i})) \\ &= \sum_{l^{-i}} p(l^{-i}|l^i) \sum_{a^{-i}} q(a^{-i}|\hat{l}^i, l^{-i}) g^i((l^i, l^{-i}), (b^i, a^{-i})). \end{aligned}$$

Hence, set $y^{-i}(a^{-i}|l^{-i}) = q(a^{-i}|\hat{l}^i, l^{-i})$ for some $\hat{l}^i \in L^i$ as punishment. □

Observe that in the previous proof, "punishment" is mostly a convenient terminology. More precisely, consider the following strategy of player i in Γ^q: at stage 2, he reports type $\hat{l}^i$ whatever his type; at stage 5, he plays an arbitrary action b^i, independently of the recommendation of the communication device. This strategy of player i can be interpreted as "nonparticipation." If player $j = -i$ plays the strategy y^{-i} in the previous proof, player i's payoff is the same as when he does not participate.

3. Cheap talk game, correlated equilibrium

In this section, we first extend the basic game $\Gamma \equiv \langle \{L^i, A^i, g^i\}_{i=1,2}, p\rangle$ by means of a long cheap talk phase. Then we define correlated equilibria in this extended game.

Let M be a finite set of messages. Let c (continue) and s (stop) be two additional messages available to the players. We define the multistage game $\text{ext}_M \Gamma$ as follows:

Stage 0. Every player i learns his type l^i as in Γ, $i = 1, 2$.

Stage 1. The players simultaneously send the message c or s to each other. If they both selected c, they simultaneously send a message in M to each other and they proceed to stage 2. Otherwise, every player i chooses an action in A^i, payoffs are given as in Γ, and the game stops.

Stage t ($t = 2, 3, \ldots$). If the game has not stopped at an earlier stage, the players simultaneously send the message c or s to each other. If they both selected c, they simultaneously send a message in M to each other and they proceed to stage $t+1$. Otherwise, every player i chooses an action in A^i, payoffs are given as in Γ, and the game stops.

The previous scenario fully describes the players' possible moves in the game $\text{ext}_M \Gamma$ and the payoffs if the moves make the game stop at some stage t. The scenario also allows the game to go on forever, which is unavoidable if the length of communication is not fixed in advance (see, e.g., Forges 1990a, Gossner and Vieille 2001, Aumann and Hart 2003). We thus have to define the payoffs in the case of infinitely long cheap talk, even if this event typically is off the equilibrium path. Since there is no particular outcome to be identified in our general Bayesian game, we assume, as Gossner and Vieille (2001) and Aumann and Hart (2003), that if communication goes on forever, the players make their decisions "at infinity."

Let $H_t = (M \times M)^{t-1}$, $t = 1, 2, \ldots$, be the set of all pairs of messages in M possibly sent before stage t and let $H_\infty = (M \times M)^{\mathbb{N}}$. We provide these sets with a measurable structure, in the standard way. Let $\mathcal{H}_t$ be the algebra over H_∞ generated by cylinder sets of the form $h_{t-1} \times H_\infty$, where h_{t-1} is a sequence in H_t. Let $\mathcal{H}_\infty$ be the σ-algebra over H_∞ generated by the algebras $\mathcal{H}_t$, $t = 1, 2, \ldots$. Finally, $N = \{\{1\}, \{2\}, \{1, 2\}\}$ describes the sets of players possibly choosing s at some stage.

A pure strategy σ^i for player i ($i = 1, 2$) in $\text{ext}_M \Gamma$ is a sequence of measurable mappings $\sigma^i = [(\delta^i_t, m^i_t, d^i_t)_{t\geq 1}, d^i_\infty]$, where

$$\delta^i_t : L^i \times H_t \to \{c, s\}, \qquad m^i_t : L^i \times H_t \to M, \quad t = 1, 2, \ldots$$

$$d^i_t : L^i \times H_t \times N \to A^i, \quad t = 1, 2, \ldots, \qquad d^i_\infty : L^i \times H_\infty \to A^i.$$

These mappings are interpreted as follows: δ^i_t describes player i's decision to continue or stop at stage t if the game is still going on at that stage; m^i_t describes which message in M he sends if both players decide to continue at stage t; d^i_t describes the action he chooses according to which player(s) decide(s) to stop at stage t; d^i_∞ describes the action he chooses if communication goes on forever.

Let $\sigma = (\sigma^1, \sigma^2)$ be a pair of pure strategies in $\text{ext}_M \Gamma$ and let $l = (l^1, l^2)$ be a pair of types chosen at stage 0. If, for these types l, σ induces the game to stop at stage t, namely if σ leads one of the player to choose s at stage t, as a function of the past history, then the payoffs associated with l and σ are computed using the mappings d_t^i and the utility functions g^i. If for these types l, σ induces cheap talk to last forever, the payoffs associated with l and σ are computed in a similar way, using the mappings d_∞^i. Payoffs in $\text{ext}_M \Gamma$ are thus well defined and the definition of the game is complete.

As explained in the Introduction, the players cannot hope to implement all communication equilibrium outcomes of Γ by cheap talk, namely as equilibrium outcomes of $\text{ext}_M \Gamma$ for some set of messages M, without randomizing their strategies in a correlated way.

A correlation device consists of a probability space $(\Omega, \mathcal{B}, \mu)$, together with sub-$\sigma$-algebras $\mathcal{B}^1$ and $\mathcal{B}^2$ of $\mathcal{B}$. The probability space $(\Omega, \mathcal{B}, \mu)$ represents extraneous events ("sunspots"), which happen independently of Γ (and $\text{ext}_M \Gamma$), in particular independently of the types in L; $\mathcal{B}^i$, $i = 1, 2$, represents player i's private information on the extraneous events. To achieve our implementation goal, we make use only of simple and well behaved correlation devices, typically describing discrete random variables.

By adding a correlation device $[(\Omega, \mathcal{B}, \mu), \mathcal{B}^1, \mathcal{B}^2]$ to $\text{ext}_M \Gamma$, we get a new extended game, $(\text{ext}_M \Gamma)^\mu$, in which, before stage 1 of $\text{ext}_M \Gamma$, every player i gets private information in $\mathcal{B}^i$ on an extraneous event, selected in $(\Omega, \mathcal{B})$ according to μ. This lottery can take place before or after stage 0, but is independent of the players' prior p. In $(\text{ext}_M \Gamma)^\mu$, every player i makes his strategic choices as a function of his extraneous information, described by $\mathcal{B}^i$ $(i = 1, 2)$. Proceeding as in Aumann and Hart (2003), a pure strategy $\boldsymbol{\sigma}^i$ for player i $(i = 1, 2)$ in $(\text{ext}_M \Gamma)^\mu$ is a sequence $\boldsymbol{\sigma}^i = [(\boldsymbol{\delta}_t^i, \mathbf{m}_t^i, \mathbf{d}_t^i)_{t \geq 1}, \mathbf{d}_\infty^i]$ of $L^i \times \mathcal{H}_t \times \mathcal{B}^i$-measurable mappings describing player i's move at stage t (including ∞), where

$$\boldsymbol{\delta}_t^i : L^i \times H_t \times \Omega \to \{c, s\}, \qquad \mathbf{m}_t^i : L^i \times H_t \times \Omega \to M, \quad t = 1, 2, \ldots$$

$$\mathbf{d}_t^i : L^i \times H_t \times N \times \Omega \to A^i, \quad t = 1, 2, \ldots, \qquad \mathbf{d}_\infty^i : L^i \times H_\infty \times \Omega \to A^i.$$

Definition 3. A correlated equilibrium of $\text{ext}_M \Gamma$ is a Nash equilibrium of $(\text{ext}_M \Gamma)^\mu$ for some correlation device $[(\Omega, \mathcal{B}, \mu), \mathcal{B}^1, \mathcal{B}^2]$. The set of all correlated equilibrium payoffs of $\text{ext}_M \Gamma$ is denoted as $\text{CE}(\text{ext}_M \Gamma)$ which is a subset of $\mathbb{R}^{|L^1| + |L^2|}$.

4. Implementing communication equilibria by cheap talk

In this section, we first state the main theorem, in terms of the standard correlated equilibrium solution concept. After we deduce two immediate corollaries, we give a sketch of the proof, which clarifies the use of unboundedly long cheap talk and indicates the role of punishments. We then turn to "approximate implementation" with finitely long cheap talk (Proposition 1). Finally, we give sufficient conditions for implementation in sequentially rational strategies (Proposition 2).

The prior probability distribution p over L, the probability distribution μ of a correlation device, and strategies $(\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2)$ in $(\text{ext}_M \Gamma)^\mu$ induce a probability distribution over $\Omega \times L \times H_\infty \times A$, and thus also conditional probability distributions over A, given every $l \in L$.

Theorem 1. *Let* $\Gamma \equiv \langle \{L^i, A^i, g^i\}_{i=1,2}, p\rangle$ *be a finite Bayesian game with two players and let* q *be a communication equilibrium of* Γ *such that* $G[q] \in \text{SINTIR}(\Gamma)$. *There exist a finite set of messages* M *and a correlated equilibrium of* $\text{ext}_M \Gamma$, *the cheap talk extension of* Γ *with messages in* M, *that induce the conditional probability distribution* $q(\cdot|l^1, l^2)$ *over actions (i.e., over* $A^1 \times A^2$*) for every pair of types* $(l^1, l^2) \in L^1 \times L^2$. *In particular, the payoff of the correlated equilibrium is* $G[q]$. *Moreover, the correlated equilibrium of* $\text{ext}_M \Gamma$ *is such that cheap talk lasts for finitely many stages almost surely.*

In this statement, the set of messages depends on the parameters of Γ and of q. If we allow for countably many messages, i.e., if we consider the extended cheap talk game $\text{ext}\,\Gamma$ in which $M = \mathbb{N}$, we can get all strictly individually rational communication equilibrium payoffs at once: $\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma) \subseteq \text{CE}(\text{ext}\,\Gamma)$. Recall that, by Lemma 1, $\text{ME}(\Gamma) \subseteq \text{INTIR}(\Gamma)$; the restriction imposed on communication equilibrium outcomes is thus relatively mild. Conversely, by proceeding as in general versions of the revelation principle, one can show that $\text{CE}(\text{ext}\,\Gamma) \subseteq \text{ME}(\Gamma)$. Hence we get the following corollary.[10]

Corollary 1. $\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma) = \text{CE}(\text{ext}\,\Gamma) \cap \text{SINTIR}(\Gamma)$.

Note that, once $\mathbb{N}$ is the set of messages, cheap talk in $\text{ext}\,\Gamma$ is described in a universal way, i.e., independently of the underlying Bayesian game Γ, as in Forges (1990b).

Corollary 1 can be interpreted as a characterization of the correlated equilibrium payoffs of the long cheap talk game $\text{ext}\,\Gamma$, since it states that $\text{CE}(\text{ext}\,\Gamma)$ and $\text{ME}(\Gamma)$ essentially coincide.[11] To make the relationship between the two sets more precise, let us denote the closure of $\text{CE}(\text{ext}\,\Gamma)$ as $\overline{\text{CE}(\text{ext}\,\Gamma)}$.

Corollary 2. *If* $\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma) \neq \varnothing$, $\overline{\text{CE}(\text{ext}\,\Gamma)} = \text{ME}(\Gamma)$.

Proof. Given that $\text{ME}(\Gamma)$ is closed, $\overline{\text{CE}(\text{ext}\,\Gamma)} \subseteq \text{ME}(\Gamma)$. To see the converse, we have to show that if $\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma) \neq \varnothing$, then $\text{ME}(\Gamma) \subseteq \overline{\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma)}$. Let $x \in \text{ME}(\Gamma)$. By Lemma 1, $x \in \text{INTIR}(\Gamma)$. Let $x^* \in \text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma)$, let α_n be a sequence in $(0, 1)$ such that $\alpha_n \to 1$, and let $x_n = \alpha_n x + (1 - \alpha_n)x^*$. Since $\text{ME}(\Gamma)$ is convex, $x_n \in \text{ME}(\Gamma)$ and from the inequalities in Definition 2, it is readily checked that $x_n \in \text{SINTIR}(\Gamma)$. Finally, $x_n \to x$. □

Sketch of the proof of Theorem 1. Let q be a communication equilibrium of Γ. We gradually construct a set of messages M, a correlation device μ for $\text{ext}_M \Gamma$, and equilibrium strategies in $(\text{ext}_M \Gamma)^\mu$ that induce the transition probability q. The correlation device μ first selects, independently for every $l \in L$, a pair of actions $a_l = (a_l^1, a_l^2) \in A$

[10] In this statement and the next ones, we do not recall that Γ is a finite two-person Bayesian game.

[11] Aumann and Hart (2003) show that even if only one of the players has private information in Γ (if, e.g., $|L^2| = 1$), the characterization of the *Nash* equilibrium payoffs of the game $\text{ext}\,\Gamma$ is fairly complex, as it relies on the martingales generated by the long cheap talk. On the contrary, most correlated equilibrium payoffs of $\text{ext}\,\Gamma$ are characterized in a tractable way, as communication equilibrium payoffs of the original Bayesian game Γ.

according to $q(\cdot|l)$. If the players could reveal their types to each other, the correlation device could send them $(a_l)_{l\in L}$ before the beginning of Γ.

To keep the correlated equilibrium conditions as close as possible to the communication equilibrium conditions, the correlation device μ selects permutations η^i of L^i ($i=1,2$) to encrypt player i's type l^i and permutations $\phi^i_{\eta(l)}$ of A^i ($i=1,2$, $l\in L$) to encrypt player i's recommended action $a^i_{\eta(l)}$. Before the beginning of Γ, the device tells player i how to encrypt his type (namely, η^i) and how to decrypt his recommended action (namely $\phi^i_{\eta(l)}$, $l\in L$). Every player's encrypted, recommended action is transmitted by the other player. More precisely, the correlation device tells the encrypted actions $b^j_{\eta(l)}=\phi^j_{\eta(l)}(a^j_{\eta(l)})$, $l\in L$, to player i, $j\neq i$.

At the first stage of cheap talk, the players can simultaneously send their encrypted type to each other; let $\eta(l)$ be the pair of messages. Let us imagine that, at a second stage of cheap talk, the players send simultaneously the corresponding encrypted actions $b^j_{\eta(l)}$ to each other. The communication equilibrium conditions guarantee that a player cannot gain in lying on his type at the first stage or on deviating (at the decision stage) from the action $(\phi^i_{\eta(l)})^{-1}(b^i_{\eta(l)})$ that he decrypts at the second stage. In other words, at this point, the correlation device mimics q.

However, the communication equilibrium conditions do not ensure that player i correctly transmits the encrypted action $b^j_{\eta(l)}$ of player j at the second stage. To fill this gap, the correlation device μ chooses a "code" $k^i(\eta(l),a^i)$ in some large set, independently and uniformly, for every pair of encrypted types $\eta(l)$ and every possible action a^i, $i=1,2$. The correlation device tells to player i the whole mapping k^i, namely, the code $k^i(\eta(l),a^i)$ associated with *every* encrypted, recommended action a^i that player i might receive from player j, but only $k^j(\eta(l),b^j_{\eta(l)})$, $l\in L$, for $j\neq i$, namely only the codes of encrypted, recommended actions $b^j_{\eta(l)}$ that player i himself must transmit to player j for some $l\in L$. If, given a pair of messages $\eta(l)$, player i transmits $a^j\neq b^j_{\eta(l)}$ to player j, player i will, with high probability, not guess correctly the corresponding code $k^j(\eta(l),a^j)$. In this case, player i is detected and punished by player j, who knows the whole mapping k^j, namely the codes of all the actions that might be transmitted to him.[12]

The equilibrium strategies suggested in the previous paragraph raise two problems. The first one, which is not typical of our construction (see, e.g., Bárány 1992 and Heller et al. 2012), is that punishments may appear as incredible threats. We come back to this below. The second and most important issue is that if player i unilaterally deviates at the second stage of cheap talk and does not transmit the correct encrypted, recommended action $b^j_{\eta(l)}$ to player j, player i typically receives a correct recommendation at the same stage. Player i may thus have updated his beliefs in such a way that player j cannot maintain player i's payoff below the communication equilibrium payoff.[13]

[12]Similar codes were used in Forges (1990b) to allow a player to check with high probability whether another player correctly transmits information generated by a correlation device.

[13]Ben-Porath (2003) identifies this issue in games with three players or more, but does not provide a thorough solution (see Ben-Porath 2006). We illustrate the difficulty on a two-person game in Section 5.

To solve the second problem, we do not fix the number of stages of cheap talk in $\text{ext}_M \Gamma$. Encrypted types are still exchanged only once at the first stage, but the correlation device μ chooses a relevant stage $t^\star \geq 2$ according to a geometric distribution. The previous encrypted, recommended actions, whose distribution is determined by q, are selected for stage $t^\star$ only. For all stages $t \neq t^\star$, the correlation device μ selects encrypted, recommended actions uniformly. Codes are selected at every stage t, as described above.

The key is that the players only discover whether stage t (≥ 2) is relevant, i.e., whether $t = t^\star$, at the end of stage t, after exchanging messages. If player j detects an incorrect code in player i's message at the end of stage t, then, with high probability, $t \neq t^\star$, so that player i's belief over L is still the prior p and player j can punish player i (strictly) below his communication equilibrium payoff (which belongs to $\text{SINTIR}(\Gamma)$). Strict punishment takes care of the small probability that deviation luckily happens at $t^\star$.

There remains to explain how the players discover whether $t = t^\star$ at the end of every stage $t \geq 2$. The correlation device μ selects "labels" λ_t^i such that $\lambda_t^1 = \lambda_t^2$ if and only if $t = t^\star$. By exchanging their labels at the same time as the encrypted, recommended actions and their codes, the players can recognize $t^\star$. To prevent cheating on the labels, codes are associated to the labels as well. For every $t \geq 2$, the correlation device μ tells to player i the code $\kappa^i(t, \lambda_t^i)$ of his own label λ_t^i at stage t, together with the whole mapping $\kappa^j(t, \cdot)$. This completes the description of the correlation device μ. Regarding strategies, if player j detects an incorrect label code in player i's message, player j punishes player i. □

Theorem 1 is proved in full details in Section 6. In particular, we show how to compute precisely the size of the set of messages M and the parameter of the geometric distribution choosing $t^\star$.

Approximation with finite cheap talk

Theorem 1 is stated in terms of the *infinitely* long cheap talk game $\text{ext}_M \Gamma$. For $T \geq 2$, let us denote as $\text{ext}_M^T \Gamma$ the extension of Γ in which the players cannot talk for more than T stages. In Section 5, we show in an example that there may exist a communication equilibrium payoff that cannot be achieved as a correlated equilibrium payoff of $\text{ext}_M^T \Gamma$, for any T. A natural question is thus whether every communication equilibrium payoff of a Bayesian game Γ can be *approximated* by a correlated equilibrium payoff of a *sufficiently long* cheap talk game extending Γ. A positive answer is given in the following proposition and is illustrated in Section 5. The result is formally established after the proof of Theorem 1, in Section 6.

Proposition 1. *Let us assume that* Γ *has a Bayesian-Nash equilibrium payoff that belongs to* $\text{SINTIR}(\Gamma)$. *Let* $x = G[q] \in \text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma)$ *and let* $\delta > 0$. *There exist a finite set of messages* M, *a finite number of stages* T, *and a payoff vector* x_δ *that is* δ*-close to* x *such that* $x_\delta \in \text{CE}(\text{ext}_M^T \Gamma)$.

Implementation in sequentially rational strategies

The proof of Theorem 1 makes use of punishment strategies that may not be "credible," in the sense that they apply to any communication equilibrium payoff in SINTIR(Γ) and are thus akin to minmax strategies. A standard way to guarantee credible punishments is to focus on communication equilibrium payoffs that are not only strictly individually rational, but even dominate a Bayesian-Nash equilibrium. Ben-Porath (2003, 2006) studies the implementation of such particular communication equilibria in Bayesian games with three players or more.

Definition 4. A payoff vector $((x^i(l^i))_{l^i \in L^i})_{i=1,2} \in \mathbb{R}^{|L^1|+|L^2|}$ in Γ is Nash-dominating if there is a Bayesian-Nash equilibrium payoff $\xi = (\xi^i(l^i)_{l^i \in L^i})_{i=1,2}$ in Γ such that

$$x^i(l^i) > \xi^i(l^i) \quad \text{for every } i = 1, 2 \text{ and } l^i \in L^i.$$

For Nash-dominating payoffs,[14] Theorem 1 can be restated in terms of a version of the *perfect Bayesian equilibrium* (PBE) solution concept, which we call *semi-weak* PBE.[15] More precisely, we require sequential rationality at every information set and Bayesian updating on the equilibrium path as in the *weak* PBE (see, e.g., Mas-Colell et al. 1995 and Myerson 1991, who refer to "weak" sequential equilibrium). We further impose a natural restriction on the players' beliefs over histories that are out off equilibrium paths, in the vein of the condition of "action-determined beliefs" of Osborne and Rubinstein (1994) (see also Fudenberg and Tirole 1991), which we define precisely below, after the statement of Proposition 2.

Proposition 2. *Let us assume that the prior p of $\Gamma \equiv \langle \{L^i, A^i, g^i\}_{i=1,2}, p\rangle$ has full support and that $x = G[q] \in \mathrm{ME}(\Gamma)$ is Nash-dominating. There exist a finite set of messages M and a correlation device $[(\Omega, \mathcal{B}, \mu), \mathcal{B}^1, \mathcal{B}^2]$ for $\mathrm{ext}_M \Gamma$ such that x is a semi-weak perfect Bayesian equilibrium payoff of $(\mathrm{ext}_M \Gamma)^\mu$.*

The reason for restricting to a prior p with full support is well explained in Gerardi (2004). As his Example 1 illustrates, without full support of the prior p in Γ, there may exist communication equilibria that can be achieved only by means of a communication device recommending a strictly dominated action to one of the players when a type profile of zero probability under p is reported. Such communication equilibria cannot be implemented with sequentially rational strategies, even if the implementation process does not rely on any punishment.

To make precise the condition on beliefs behind our semi-weak PBE, let $\boldsymbol{\sigma}$ be an equilibrium of $(\mathrm{ext}_M \Gamma)^\mu$ and let h_{t-1} be a sequence of messages before stage t, i.e., $h_{t-1} \in H_t = (M \times M)^{t-1}$; let $m_t = (m_t^1, m_t^2)$ be a pair of messages at stage t. Assume that

[14] Nash domination is by no means a necessary condition, as illustrated, for instance, by the sender–receiver case (see Forges 1985 and Section 7.2).

[15] We limit ourselves to strengthening the rationality of the specific equilibrium strategies constructed in the proof of Theorem 1, without addressing the question of an appropriate definition of refined correlated equilibrium (see, e.g., Dhillon and Mertens 1996 for a discussion of this topic).

the probability of h_{t-1} given l^i and $\mathcal{B}^i$ is positive under the distribution induced by p, μ, and $\boldsymbol{\sigma}$. Then the belief of player i over L^j given l^i, $\mathcal{B}^i$, h_{t-1}, and m_t does not depend on m_t^i. This condition guarantees that, in the previous sketch of the proof of Theorem 1, a "lucky deviator" (who does not transmit the correct encrypted action to the other player at stage t, but correctly guesses its code and then discovers that $t^\star = t$) updates his belief on the other player's type.

In Section 6.3, we establish Proposition 2 by using the same correlation device $[(\Omega, \mathcal{B}, \mu), \mathcal{B}^1, \mathcal{B}^2]$ and the same set M of messages as in the proof of Theorem 1.

5. An example

We consider a variant of the "secret sharing" problem, which is well known in computer science (see, for instance, Abraham et al. 2008). We show that there is a communication equilibrium in which the players reach an efficient outcome by sharing the secret, but that this outcome cannot be achieved as a correlated equilibrium outcome of any cheap talk game involving a bounded number of stages. This illustrates the need for unboundedly long cheap talk in Theorem 1. We also show that Proposition 1 applies to the example so that the efficient outcome can be reached approximately with sufficiently long, bounded cheap talk.

The secret sharing game Γ is derived from an auxiliary game $\hat{\Gamma}$, in which both players have two equally likely possible types in $S^1 = S^2 = \{0, 1\}$, which we refer to as payoff types. The payoff types of the players are chosen independently of each other. Every player has two possible actions: $A^1 = A^2 = \{0, 1\}$. The payoff functions $g^i : S^1 \times S^2 \times A^1 \times A^2 \to \mathbb{R}$, $i = 1, 2$, are summarized in the following table:

g	s^2	0		1	
s^1	A	0	1	0	1
0	0	3, 3	6, −2	0, 0	−2, 6
	1	−2, 6	0, 0	6, −2	3, 3
1	0	0, 0	−2, 6	3, 3	6, −2
	1	6, −2	3, 3	−2, 6	0, 0

The interpretation is as follows. The secret is $s = s^1 + s^2 (\text{mod}\, 2)$. Given the secret $s \in \{0, 1\}$, the "right" (resp., "wrong") action is to play according to the secret, namely $a^i = s$ (resp., $a^i \neq s$). Both players have the same preferences: being the only one to take the right action is preferred to both taking the right action, which is preferred to both taking the wrong action, which is itself preferred to being the only one to take the wrong action.

In the game $\hat{\Gamma}$, the pair of expected payoffs $(3, 3)$ can be achieved only as a completely revealing outcome, in which both players take the right action.[16] But complete revelation cannot be achieved at a communication *equilibrium* of $\hat{\Gamma}$: every player can

[16]To see this, let $q(\cdot|l)$, $l \in L$, be conditional probability distributions over actions given types achieving the pair of expected payoffs $(3, 3)$ in the game $\hat{\Gamma}$. Every $q(\cdot|l)$ is a distribution over the *same* payoffs $\{(0, 0), (-2, 6), (3, 3), (6, -2)\}$, in which $(3, 3)$ is an extreme point.

gain in lying unilaterally about his payoff type so as to be the only one to take the right action.

We now modify $\hat{\Gamma}$ into a more complex game Γ. In Γ, the payoff type of every player is enriched into a "full type," which is highly correlated to the full type of the other player. More precisely, player i's type in Γ is denoted as l^i and consists of a 4-tuple. The first component of l^i is player i's payoff type s^i. To define the other three components of l^i, let E be a finite set and let $|E|$ denote the number of elements in E. The set E is interpreted as a set of "codes." The sets of types in Γ are $L^i = \{0, 1\} \times E \times E \times E$, $i = 1, 2$. In Γ, nature first makes the following choices:

1. Choose a pair of payoff types (s^1, s^2), as in $\hat{\Gamma}$.

2. Choose four codes e_0^1, e_1^1, e_0^2, e_1^2 in E, independently of each other, with probability $1/|E|$ each.

Player i's type is $l^i = (s^i, e^i_{s^i}, e_0^{-i}, e_1^{-i})$, $i = 1, 2$, i.e., player i is informed of his payoff type s^i, of the code $e^i_{s^i} \in E$ of his payoff type s^i, and of the codes e_0^{-i} and e_1^{-i} of the two possible payoff types of the other player. Player i is not informed of the code of the other possible payoff type he might have or on the payoff type of the other player, of course. The action sets and the payoff functions in Γ are the same as in $\hat{\Gamma}$, in the sense that payoffs only depend on payoff types and actions.

If player i can talk to the other player $j = -i$ and wants to reveal his payoff type s^i to him, player i also sends the code $e^i_{s^i}$, so that player j, who knows the code of the two possible payoff types of player i, namely, e^i_0 and e^i_1, can check that player i's reported payoff type is consistent with the codes. If player i wants to lie on his payoff type, he has to guess the corresponding code, with a probability of $1 - 1/|E|$ of being detected by player j.[17]

Even if no communication device is available, every player can detect the other's lie with high probability by checking the codes, but this typically happens *after* that useful information has been transmitted. The situation is very different when there is a communication device. In this case, the device does not release any information when it detects cheating, which protects the honest player. This effect cannot be simulated at a Nash equilibrium of a cheap talk game extending Γ.

Let us show that the vector of conditional expected payoffs $((3, 3), (3, 3))$ is in $\mathrm{ME}(\Gamma)$. For that, we describe a canonical communication device $q : L \to \Delta A$. Every player i, $i = 1, 2$, reports a type $(r^i, e^i, \varepsilon_0^{-i}, \varepsilon_1^{-i})$ to the communication device q, which then recommends actions as follows:

1. If $e^i = \varepsilon^i_{r^i}$ and $e^j = \varepsilon^j_{r^j}$, q computes $r = r^1 + r^2 (\mathrm{mod}\, 2)$ and sets $a^1 = a^2 = r$.

2. Otherwise, q chooses an action profile (a^1, a^2) uniformly.

Let us check that q defines a communication equilibrium. Assume that player j is honest and obedient, and consider player $i = -j$ with type $(s^i, e^i_{s^i}, e_0^{-i}, e_1^{-i})$. Suppose

[17] The technique of codes is also useful in the proof of Theorem 1. However, in the current example, codes are not generated by a correlation device, but as part of the types in the Bayesian game.

first that $r^i \neq s^i$, namely that player i lies on his component of the secret. Player i has no information on the code $e^i_{r^i}$, which has been chosen with probability $1/|E|$ in E; he thus guesses it correctly with probability $1/|E|$. In this case, the device recommends actions $a^1 = a^2 = r^i + s^j$. By playing against the recommendation of the device, player i gets the highest possible payoff, 6. Otherwise, if player i does not guess $e^i_{r^i}$ correctly, the device selects actions uniformly, and player i can as well play against the recommendation of the device. His total expected payoff is $1/|E| \times 6 + (1 - 1/|E|) \times [\frac{1}{4} \times 3 + \frac{1}{4} \times 6 + \frac{1}{4} \times (-2)]$, which is <3 as soon as $|E| \geq 4$. All other possible deviations of player i, e.g., involving cheating in the other player's codes, either give rise to a higher probability of being detected and reduce his expected payoff or have no effect on the payoffs. As we already observed above, while completely revealing in terms of the payoff types (in $S^1 \times S^2$), the communication equilibrium expected payoff $(3, 3)$ cannot be achieved as a *Nash* equilibrium of a cheap talk game like $\text{ext}_M \Gamma$.

The vector of conditional expected payoffs $((3, 3), (3, 3))$ is in SINTIR(Γ): by playing both actions with probability $\frac{1}{2}$, independently of his type, player j guarantees that player $i = -j$'s payoff does not exceed $\frac{7}{4}$, whatever his type and his action.[18] Obviously, this punishment depends on the fact that player i does not know player j's share of the secret. By Theorem 1, $((3, 3), (3, 3))$ can thus be achieved as a correlated equilibrium of a long cheap talk game $\text{ext}_M \Gamma$ for some finite set of messages M. We show below that in any extended cheap talk game in which the number of stages is fixed, the players cannot reach $((3, 3), (3, 3))$.

Let us fix an extension $\text{ext}^T_M \Gamma$ of Γ in which the cheap talk phase cannot exceed T stages. Every stage $t = 0, 1, \ldots, T$ of $\text{ext}^T_M \Gamma$ can be described as in $\text{ext}_M \Gamma$ for some set M of messages, but the moves in $\{c, s\}$ are not necessary: the game goes on for $T + 1$ stages, with final decisions at stage $T + 1$, whatever the history.[19] Let us assume that $\text{ext}^T_M \Gamma$ has a correlated equilibrium achieving the expected payoff $(3, 3)$, namely complete revelation of the secret. At the last stage T, both players must know the secret on every possible history on the equilibrium path. Without loss of generality, this does not happen at stage $T - 1$; otherwise the deadline could be $T - 1$.

Thus, at the end of stage $T - 1$, there exists a history $\mathbf{h}_{T-1} = (l, \omega, h_{T-1})$, where h_{T-1} is the sequence of messages up to stage $T - 1$, which has positive probability at equilibrium, for which at least one of the players, say player 1, does not know the secret, namely player 1's posterior probability that player 2's type is 0 *is not* 0 *or* 1. Hence, on $\mathbf{h}_{T-1}$, player 1 relies on player 2's message at stage T to learn the secret. Note that the history $\mathbf{h}_{T-1}$ involves the choice $\omega = (\omega^1, \omega^2)$ of the underlying correlation device, hence is not necessarily fully identified by player 2. But player 2 can select his message uniformly, independently of the past, at stage T. If player 2 deviates in this way (only at stage T), while player 1 does not deviate, player 2 learns the secret at stage T, on every possible history, while player 1 does not learn it at least on $\mathbf{h}_{T-1}$. In the next paragraph, we complete player 2's deviation by describing how he chooses his action and we show that his deviation is profitable.

[18] In fact, the strategies consisting of playing both actions with the same probability, independently of the type, form a Bayesian-Nash equilibrium.

[19] Hence, on some histories, cheap talk may become vacuous from some stage on.

At the end of stage $T-1$, player 2's information consists of his type l^2, the private extraneous signal from the correlation device ω^2 and the messages exchanged at stages $1, \dots, T-1$. Given his information, player 2 determines the message m_T^2 he should send at stage T as if he did not deviate. Since there is no deviation at any stage $1, \dots, T-1$, player 1 sends his message m_T^1 at stage T as in equilibrium. Even if player 2 deviates at stage T, he has the same information at the end of stage T as when he does not deviate. In particular, m_T^1 and m_T^2 are part of player 2's information. We complete his deviation as follows: after having sent his (uniformly selected) message $\widetilde{m}_T^2$ to player 1 and having received player 1's message m_T^1, he chooses his action in A^2 according to his equilibrium strategy as if the messages at stage T were (m_T^1, m_T^2). This guarantees him a payoff strictly higher than 3 if the history $\mathbf{h}_{T-1}$ identified above occurs and no less than 3 otherwise. Hence player 2's deviation is profitable.

The constructive proof of Theorem 1 avoids the obstacles of a bounded cheap talk phase by introducing extra uncertainty for the players about the time $t^\star$ at which they reveal their part of the secret to each other. In such a construction, the number of conversation stages cannot be deterministically bounded. Nevertheless, in equilibrium, the players stop talking with probability 1. The probability that a deviator can affect the conversation in a way that it lasts forever can be made arbitrarily small. The main idea is that, at every stage t, player i, say, does not know whether $t = t^\star$, i.e., whether he will receive useful information from player $j = -i$ at that stage. Hence player i may not have any incentive to send a message that differs from the one prescribed by the correlation device. In particular, in our construction, with large probability, a deviation of player i is detected by player j before that player i learns the secret, so that player j can stop the conversation and punish player i in the initial Bayesian game Γ, with prior p.

To sum up, the proof of Theorem 1 confirms that, in the secret sharing game, the players learn the secret with probability 1 after a random finite number $t^\star$ of stages of correlated cheap talk (namely $x = ((3,3),(3,3)) \in \mathrm{CE}(\mathrm{ext}_M \Gamma)$). We show that this result cannot be true if $t^\star$ is imposed not to exceed a fixed, deterministic bound T, i.e., that $x \notin \bigcup_{T \geq 1} \mathrm{CE}(\mathrm{ext}_M^T \Gamma)$. The requirement that the players learn the secret *with probability* 1 is essential to this observation, as follows from Proposition 1. To see that this proposition applies to the secret sharing game, let us set $\xi = ((2,2),(2,2))$. The payoff ξ is associated with the Bayesian-Nash equilibrium in which one of the players chooses the action 0, independently of his type, and the other player chooses the action 1, independently of his type. Furthermore, $\xi \in \mathrm{SINTIR}(\Gamma)$, since as noticed above, every player can guarantee that the other's interim expected payoff does not exceed $\frac{7}{4}$. Thus, from Proposition 1, for every $\delta > 0$, there is a finite number of stages T such that the game $\mathrm{ext}_M^T \Gamma$ has a correlated equilibrium at which the players learn the secret with probability at least $1-\delta$ and thus get approximately the desired payoff x, i.e., for every $\delta > 0$, there exist a finite number of stages T and a payoff x_δ that is δ-close to x such that $x_\delta \in \mathrm{CE}(\mathrm{ext}_M^T \Gamma)$. Finally, observe also that the payoff x Nash-dominates ξ (or $((\frac{7}{4}, \frac{7}{4}), (\frac{7}{4}, \frac{7}{4}))$) so that, by Proposition 2, it can be achieved with sequentially rational strategies.

6. Proof of the results

6.1 *Proof of Theorem 1*

Let us fix a communication equilibrium q of Γ, such that $G[q] \in \mathrm{SINTIR}(\Gamma)$. We construct a set of messages M and a correlated equilibrium of $\mathrm{ext}_M \Gamma$ that satisfy the requirements of the theorem. The precise size of M is determined when we check the equilibrium conditions. We start by describing a correlation device, namely a probability space $(\Omega, \mathcal{B}, \mu)$, and private signals for every player, namely sub-σ-algebras $\mathcal{B}^1$ and $\mathcal{B}^2$. Then we define the players' strategies.

Items selected by the correlation device: $(\Omega, \mathcal{B}, \mu)$ We make a list of the items selected by the correlation device. Unless specified otherwise, these items are selected uniformly in the finite set to which they belong and they are all selected independently of each other.

The correlation device makes the following choices.

1. For $i = 1, 2$, a permutation η^i of L^i. Let $\eta = (\eta^1, \eta^2)$ and $\eta(l) = (\eta^1(l^1), \eta^2(l^2))$ for every $l = (l^1, l^2) \in L$.
2. A stage $t^\star \in \{2, 3, \ldots\}$, according to a geometric distribution with success parameter $z > 0$ to be specified later.
3. For every $l \in L$, a pair of actions $a_{t^\star, \eta(l)} \in A$, according to $q(\cdot|l)$.
4. For every $l \in L$ and every $t \in \{2, 3, \ldots\}$, $t \neq t^\star$, a pair of actions $a_{t,\eta(l)} \in A$.
5. For $i = 1, 2$, every $l \in L$, and every $t \in \{2, 3, \ldots\}$, a permutation $\phi^i_{t,\eta(l)}$ of A^i. Let us set $b^i_{t,\eta(l)} = \phi^i_{t,\eta(l)}(a^i_{t,\eta(l)})$.
6. For $i = 1, 2$, every $l \in L$, every action $b^i \in A^i$, and every $t \in \{2, 3, \ldots\}$, a code $k^i(t, \eta(l), b^i) \in M$.
7. For $i = 1, 2$ and every $t \in \{2, 3, \ldots\}$, a pair of labels $\lambda^i_t \in M$ such that $\lambda^1_{t^\star} = \lambda^2_{t^\star}$ and $\lambda^1_t \neq \lambda^2_t$ if $t \neq t^\star$.
8. For $i = 1, 2$, every $l \in L$, every $t \in \{2, 3, \ldots\}$, and every label $\lambda \in M$, a code $\kappa^i(t, \lambda) \in M$.

To sum up, only $t^\star$ in choice 2 and $a_{t^\star,\eta(l)}$, $l \in L$, in choice 3 are selected according to a specific, nonuniform probability distribution. The stage $t^\star$ is the only random variable that is not finite. In choice 7, the labels λ^1_t and λ^2_t at stage t are not independent from each other or from t. The parameter z represents the probability that $t^\star$ is the next stage; z and the size of M are computed at the end of the proof (see the expression (4) below).

Private extraneous information: $\mathcal{B}^i$, $i = 1, 2$ The correlation device sends the following private signal[20] to player i, $i = 1, 2$.

— The permutation η^i of L^i selected in choice 1.

[20]It is understood that functions over $L = L^1 \times L^2$ are described as $L^1 \times L^2$ tables for a given order on L^1 and L^2.

— The permutations $\phi^i_{t,\eta(l)}$ of A^i, $l \in L$, $t \in \{2, 3, \ldots\}$, selected in choice 5.

— The (encrypted, recommended) actions (for the other player, $-i$) $b^{-i}_{t,\eta(l)} \in A^{-i}$ for every $l \in L$, $t \in \{2, 3, \ldots\}$ defined in choice 5, together with their associated code $k^{-i}(t, \eta(l), b^{-i}_{t,\eta(l)})$ selected in choice 6.

— The code functions $k^i(t, \eta(l), \cdot) : A^i \to M$ for every $l \in L$, $t \in \{2, 3, \ldots\}$, selected in choice 6.

— The labels λ^i_t, $t \in \{2, 3, \ldots\}$ selected in choice 7, together with their associated code $\kappa^i(t, \lambda^i_t)$ selected in choice 8.

— The code functions (of the other player, $-i$) $\kappa^{-i}(t, \cdot) : M \to M$ for every $t \in \{2, 3, \ldots\}$, selected in choice 8.

We denote player i's private signal as

$$\omega^i = \begin{bmatrix} \eta^i \\ (\phi^i_{t,\eta(l)})_{t\geq 2, l\in L} \\ (b^{-i}_{t,\eta(l)}, k^{-i}(t, \eta(l), b^{-i}_{t,\eta(l)}), k^i(t, \eta(l), \cdot))_{t\geq 2, l\in L} \\ (\lambda^i_t, \kappa^i(t, \lambda^i_t), \kappa^{-i}(t, \cdot))_{t\geq 2} \end{bmatrix}. \tag{2}$$

At this point, the description of the game $(\mathrm{ext}_M \Gamma)^\mu$ is complete.

Equilibrium strategies $(\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2)$ *in* $(\mathrm{ext}_M \Gamma)^\mu$ We first give a rough description of the strategies $(\boldsymbol{\sigma}^1, \boldsymbol{\sigma}^2)$ and of the way in which they combine with each other. The basic idea is that the geometric random variable $t^\star$ describes the only relevant stage, in which players determine the actions $a^i_{t^\star,\eta(l)}$, $i = 1, 2$, to be played in the Bayesian game. For every l, the pair of actions $a_{t^\star,\eta(l)}$ selected in choice 3 is distributed according to $q(\cdot|l)$. However, the players cannot fully reveal their types to each other or know more than their own action. Hence permutations are applied both to the types (η^i, selected in choice 1) and to the actions ($\phi^i_{t^\star,\eta(l)}$, selected in choice 5). At stage 1, the players send hidden types, $\eta^i(l^i)$, $i = 1, 2$, to each other. At stage $t^\star$, every player i sends the message $b^{-i}_{t^\star,\eta(l)}$ to the other player. If player i indeed receives the message $b^i_{t^\star,\eta(l)}$ from the other player, he is able to evaluate his action as $a^i_{t^\star,\eta(l)} = (\phi^i_{t^\star,\eta(l)})^{-1}(b^i_{t^\star,\eta(l)})$ by applying the inverse of the permutation $\phi^i_{t^\star,\eta(l)}$, and this action is distributed as in the communication equilibrium. There remains to make every player able to identify $t^\star$, *only after* having transmitted his recommended action $b^{-i}_{t^\star,\eta(l)}$ to the other player. This is the role of the labels selected in choice 7. By construction, as in the communication equilibrium, player i does not gain by pretending another type at stage 1 or deviating from his recommended action $a^i_{t^\star,\eta(l)}$. But player i must transmit a recommended action $b^{-i}_{t^\star,\eta(l)}$ to the other player, which has no counterpart in the communication equilibrium. This is the role of the codes[21] selected in choice 6. To prevent cheating in the labels, further

[21]Restricted to two stages, $t = 1$ and $t^\star$ chosen deterministically equal to 2, the correlation device is a variant of the one used in Forges (1990b) in the case of three players.

codes are needed, selected in choice 8. We detail the equilibrium strategies in the next paragraph.

Given his private extraneous signal ω^i described above, player i's equilibrium strategy in $\text{ext}_M \Gamma$ is as follows.

— At stage 1, player i chooses c. If both players select c, player i announces $\eta^i(l^i)$ if his type is l^i; otherwise, he plays a punishment action against the other player and the game stops (recall that $G[q] \in \text{SINTIR}(\Gamma)$, so that player i can select a punishment[22] according to some $y^i(\cdot|l^i) \in \Delta A^{-i}$). Let $\eta(l)$ be the pair of announcements at the first stage (if (c, c) was chosen).

— At stage 2, player i chooses c.

— At every stage $t \geq 2$, if both players select c, player i sends the message

$$b^{-i}_{t,\eta(l)}, \qquad k^{-i}(t, \eta(l), b^{-i}_{t,\eta(l)}), \qquad \lambda^i_t, \qquad \kappa^i(t, \lambda^i_t).$$

— At stage 2, if (c, c) was not selected, player i punishes the other player, as in stage 1.

— At every stage $t \geq 2$, if (c, c) was selected, then, right after having received the other player's last message, player i checks whether the latter is consistent with the codes, namely that player $-i$'s announcement $(b^i_t, k^i_t, \lambda^{-i}_t, \kappa^{-i}_t)$ satisfies

$$k^i_t = k^i(t, \eta(l), b^i_t), \qquad \kappa^{-i}_t = \kappa^{-i}(t, \lambda^{-i}_t).$$

If these equalities do not hold at stage t, player i stops the cheap talk, that is, he chooses s at the beginning of stage $t+1$ and plays a punishment action against the other player as above.

— At every stage $t \geq 2$, if (c, c) was selected, player i also checks whether his label λ^i_t coincides with the label sent by the other player, namely whether $\lambda^i_t = \lambda^{-i}_t$. If yes, and no deviation was detected, player i concludes that $t = t^\star$; he chooses to stop (namely s) at the beginning of stage $t+1$, the cheap talk ends, and player i determines his action a^i by applying the inverse of the permutation $\phi^i_{t,\eta(l)}$ (which he received from the correlation device) to the message b^i_t (which he received from the other player):

$$(\phi^i_{t,\eta(l)})^{-1}(b^i_t) = a^i.$$

— If, at the beginning of some stage $t \geq 3$, player i chooses c but the other player j $(= -i)$ chooses s, player i punishes player j as above.

— Should cheap talk last forever, $\mathbf{d}^i_\infty$ $(i = 1, 2)$ could be defined in an arbitrary way.

To sum up, if both players follow the prescribed strategies, the conversation lasts for at least two stages. Stage 1 is the only stage where the players send a type dependent

[22] To be consistent with our definition of strategies in $(\text{ext}_M \Gamma)^\mu$, in which all randomizations are made by the correlation device, possible punishment strategies should in fact be selected by the correlation device.

message, but posteriors are not updated until stage $t^\star$ is reached. Stages $t \geq 2$ are used for possible coordination. Coordination happens when $\lambda_t^1 = \lambda_t^2$, namely when $t = t^\star$; in this case, final decisions are made at stage $t^\star + 1$.

To check that the prescribed strategies form an equilibrium in the game $(\text{ext}_M \Gamma)^\mu$, we assume for simplicity that player 2 does not deviate in $(\text{ext}_M \Gamma)^\mu$ and consider possible deviations of player 1. Let l^1 be his type and let ω^1 be his extraneous signal, described as in (2). Since player 1's payoff $G^1[q]$ is individually rational, he cannot benefit from choosing s at the beginning of stage 1 of $(\text{ext}_M \Gamma)^\mu$. Thus let $\eta^1(\hat{l}^1)$ be his further message at stage 1, with $\hat{l}^1$ possibly different from l^1. We distinguish between several deviations of player 1. We start with deviations that are already feasible in the communication equilibrium and we show that they are unprofitable, namely that the correlated equilibrium of $(\text{ext}_M \Gamma)^\mu$ "mimics" the communication equilibrium.

Equilibrium conditions: Undetectable deviations Let us assume that from stage 2 on, player 1 sends all his messages as prescribed by his correlated strategy. More precisely, let $t \geq 2$ be a stage t at which the conversation is still going on. Given our current assumptions, it must be that $\lambda_r^1 \neq \lambda_r^2$ for every stage r such that $2 \leq r < t$. At the beginning of stage t, player 1 has not learnt anything on $t^\star$, player 2's type, or recommended actions, since all items that player 1 can interpret in ω^1 have been selected uniformly (this holds in particular for *every* action $b^2_{t,\eta^1(\hat{l}^1),\eta^2(l^2)}$, including $b^2_{t^\star,\eta^1(\hat{l}^1),\eta^2(l^2)}$, which is obtained by applying a random permutation to $a^2_{t,\eta^1(\hat{l}^1),\eta^2(l^2)}$). Furthermore, at the beginning of stage t, given ω^1 and the sequence of moves in $(\text{ext}_M \Gamma)^\mu$ up to stage t (including his first move $\eta^1(\hat{l}^1)$), player 1 anticipates that the pair of actions to be determined (but not necessarily played) at the further stage $t^\star$ are

$$(\phi^i_{t^\star,\eta^1(\hat{l}^1),\eta^2(l^2)})^{-1}(b^i_{t^\star,\eta^1(\hat{l}^1),\eta^2(l^2)}) = a^i_{t^\star,\eta^1(\hat{l}^1),\eta^2(l^2)}, \quad i = 1, 2. \tag{3}$$

By construction, given player 1's information at the beginning of stage t, this pair of actions is distributed according to $q(\cdot|\hat{l}^1, l^2)$. In other words, if player 2 does not deviate and player 1 of type l^1 sends $\eta^1(\hat{l}^1)$ at the first stage and all his other messages as prescribed, the actions computed by the players at stage $t^\star$, namely (3), are distributed exactly as the actions recommended by the communication device q when player 1's reported type is $\hat{l}^1$ and player 2's type is l^2. Hence player 1 does not deviate at the first stage by lying on his type and/or at $t^\star + 1$ by choosing an action other than the one computed in (3).

The previous paragraph also shows that if both players follow the prescribed strategies at every stage, the conditional probability distribution over actions (i.e., over $A^1 \times A^2$) given types $(l^1, l^2) \in S^1 \times S^2$ is $q(\cdot|l^1, l^2)$; in particular, the expected payoffs are $G[q]$.

We consider further possible deviations of player 1.

Equilibrium conditions: Deviations that are detectable with high probability Let l^1, ω^1, and $\hat{l}^1$ be as above. As already observed for stage 1, if player 2 does not deviate, player 1 cannot gain in sending his messages as prescribed and choosing s at the beginning of a stage at which he should choose c, since his payoff $G^1[q]$ is individually rational.

As above, let us consider a stage t at which the conversation is still going on. Assume that player 1 does not send (at least one of) the prescribed variables $b^2_{t,\eta^1(\hat{l}^1),\eta^2(l^2)}$ and λ^1_t in his message to player 2. Since the codes are chosen uniformly in M, the corresponding codes $k^2(t, \eta^1(\hat{l}^1), \eta^2(l^2), b^2_{t,\eta^1(\hat{l}^1),\eta^2(l^2)})$ and $\kappa^1(t, \lambda^1_t)$, are incorrect with probability (at least) $1 - 1/|M|$, in which case player 2 detects an inconsistency, stops the conversation, and chooses his action according to a punishment strategy $y^2(\cdot|l^2)$. If it turns out that $\lambda^1_t \neq \lambda^2_t$, player 1 does not learn anything; in particular, his probability distribution over L^2 is still $p(\cdot|l^1)$. In this case, player 2 can pick the strategy $y^2(\cdot|l^2)$ in such a way that player 1's payoff does not exceed $G^1[q|l^1] - \epsilon$ for some $\epsilon > 0$, since $G^1[q]$ is *strictly* individually rational in the original game Γ (whose parameters involve the prior p). However, if $\lambda^1_t = \lambda^2_t$, so that $t = t^\star$, player 1 acquires new information; the effect of the punishment strategy becomes unclear, except for the fact that player 1's payoff cannot exceed the largest possible payoff in the game Γ, which we denote by α. Finally, if player 1's deviation is not detected, his payoff can also be bounded by α (in this case, the conversation could be infinite). By recalling that, at every stage t at which the game has not yet stopped, the probability that $t = t^\star$ is z, we compute the following upper bound on player 1's payoff $G^1_{\text{dev}}(l^1)$ when he deviates as described above:

$$G^1_{\text{dev}}(l^1) \leq (1 - 1/|M|)\big(z\alpha + (1 - z)(G^1[q|l^1] - \epsilon)\big) + \alpha/|M|. \tag{4}$$

If the set M of messages[23] is large enough and the probability $z > 0$ is small enough, the previous bound does not exceed $G^1[q|l^1]$, namely

$$G^1_{\text{dev}}(l^1) \leq G^1[q|l^1].$$

We have thus shown that the correlated strategies described above form an equilibrium of the game $(\text{ext}_M \Gamma)^\mu$ that achieves the conditional probability distributions $q(\cdot|l)$ of the communication equilibrium; in particular, the payoff $G[q]$. At equilibrium, given the geometric distribution of $t^\star$, the conversation ends with probability 1. □

6.2 *Proof of Proposition 1*

Let ξ be a Bayesian-Nash equilibrium payoff such that $\xi \in \text{SINTIR}(\Gamma)$ and let $x \in \text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma)$. There exists $\epsilon > 0$ such that player i ($i = 1, 2$) has strategies y^i_x and y^i_ξ to punish player $j = -i$ in Γ at the payoff vectors $(x^j(l^j) - \epsilon)_{l^j \in L^j}$ and $(\xi^j(l^j) - \epsilon)_{l^j \in L^j}$, respectively.

Let δ be given. Let us choose $\rho \in (0, 1)$ such that

$$x_\delta \equiv (1 - \rho)x + \rho\xi$$

is δ-close to x (i.e., $\rho\|\xi - x\| \leq \delta$).

[23]The bound (4) reflects the required size of M as far as codes are concerned. The set M should of course be also large enough to contain the other messages to be transmitted by the players (i.e., $|M| \geq \max\{|L^i|, |A^i|, i = 1, 2\}$).

We temporarily fix T. Let the correlation device make its choices as in the proof of Theorem 1; in particular, let $t^\star$ be chosen as in choice 2, according to a geometric distribution with parameter z (which is also fixed for the moment) and let actions be chosen as in choice 3, using q. Let player i's prescribed strategy be as in the proof of Theorem 1 as long as he does not detect any deviation and if he deduces that $t^\star \leq T$. Let player i play the Bayesian-Nash equilibrium strategy associated with payoff ξ if he detects no deviation until stage T and concludes that $t^\star > T$.

If the players follow the previous strategies, the expected payoff of player j of type l^j, just before stage t, i.e., conditionally on $t^\star \geq t$, is

$$x_t^j(l^j) = x_t^j(z, T)(l^j) \equiv (1-(1-z)^{T-t+1})x^j(l^j) + (1-z)^{T-t+1}\xi^j(l^j). \tag{5}$$

By the above properties of x and ξ, player i can punish player $j = -i$ at the payoff vector $(x_t^j(l^j) - \epsilon)_{l^j \in L^j}$ if player j deviates at stage t, for every $t \leq T$, by playing y_x^i with probability $1-(1-z)^{T-t+1}$ and y_ξ^i with probability $(1-z)^{T-t+1}$. Hence, by proceeding as in (4), if player j deviates at stage t, his expected payoff cannot exceed

$$(1-1/|M|)\big(z\alpha + (1-z)(x_t^j(l^j) - \epsilon)\big) + \alpha/|M|.$$

We can choose $|M|$ and z (as a function of ϵ) in such a way that, for every j, l^j and $t \leq T$, this bound is less than or equal to $x_t^j(l^j)$.[24] This guarantees that no player can gain in deviating at any stage $t \leq T$. The corresponding equilibrium payoff is computed from (5) at $t = 1$. More precisely, we choose $z' < z$ and $T = T(z')$ such that $x_1(z', T) = x_1(z', T(z')) = x_\delta$, i.e., $\rho = (1-z')^T$. □

Remark (The assumption $\xi \in \mathrm{SINTIR}(\Gamma)$ in Proposition 1). In the previous proof, it is an essential assumption that $\xi \in \mathrm{SINTIR}(\Gamma)$. Intuitively, just before the last stage of the game, i.e., when $t = T$, the players' equilibrium expected payoff becomes $x_T(z, T) = zx + (1-z)\xi$ (omitting the indices j, l^j). If ξ itself is in $\mathrm{SINTIR}(\Gamma)$, we can choose the level of punishment to be $x_T - \epsilon$. When choosing z to be sufficiently small, the difference between the equilibrium payoff and the punishment stays constant at ϵ. If $\xi \in \mathrm{INTIR}(\Gamma) \setminus \mathrm{SINTIR}(\Gamma)$, by proceeding as in Corollary 2, $x_T(z, T)$ can still be supported by strict punishment, namely by $v_T = z(x-\epsilon) + (1-z)\xi$. However, in this case, the difference $x_T(z, T) - v_T = z\epsilon$ converges to 0 as z goes to 0. Hence choosing a smaller z decreases the effectiveness of punishment, which in turn necessitates the choice of an even smaller z.

6.3 *Proof of Proposition 2*

Let us fix a communication equilibrium q of Γ such that the associated payoff $G[q]$ is Nash-dominating, namely higher than the expected payoff of some Nash equilibrium $\zeta = (\zeta^1, \zeta^2)$ of Γ for every type of every player. Let us consider the set of messages M and the correlation device μ constructed in the proof of Theorem 1; μ induces a scenario,

[24] More precisely, $\forall \epsilon > 0\ \exists |M| \in \mathbb{N}$, $\exists 0 < z < 1$ such that for *any* feasible payoff γ in Γ, $(1-1/|M|)[z\alpha + (1-z)(\gamma - \epsilon)] + \alpha/|M| \leq \gamma$.

namely a prescribed plan of actions for every player. As in the proof of Theorem 1, player i's strategy $\boldsymbol{\sigma}^i$ in $\text{ext}_M(\Gamma)^\mu$ first consists of following the prescribed plan of actions and of punishing player $j = -i$ if it appears that player j did not follow the plan ($i = 1, 2$).[25] Player i now plays ζ^i if he has to punish the other player, but we have to further complete the description of strategies and beliefs to show that they form a semi-weak PBE. Strategy $\boldsymbol{\sigma}^i$ consists of stopping the game and choosing an action in Γ according to ζ^i at basically *all* information sets such that the prescribed plan of actions was not followed at some earlier stage, possibly by the player who has to move at that information set.[26] Such a specification of $\boldsymbol{\sigma}^i$ makes player i's strategy sequentially rational provided that player i's belief over L^j is his prior and that he expects that player j will play ζ^j. There is only one class of information sets out off the equilibrium path at which the underlying player updates his belief over the other player's type. In the next paragraph, we describe these information sets for player i and how $\boldsymbol{\sigma}^i$ operates at them.

Let $t \geq 2$. Assume that player i followed the prescribed plan of actions at all stages $<t$ but does not follow it at stage t. Assume also that player i observed correct codes in player j's messages at all stages less than or equal to t and that player j's reported label at stage t, λ_t^j, coincides with the label λ_t^i that player i received from the correlation device μ. At such an information set, player j's moves are exactly the same as on the equilibrium path. Since our semi-weak PBE requires that player i's beliefs on player j's type l^j do not depend on player i's own last move, player i must update his belief over L^j. We also assume that player i believes that player j will detect his deviation and thus punish him immediately by stopping the game at the beginning of stage $t+1$ and playing ζ^j. Note that semi-weak PBE does not restrict player i's belief over ω^j, so that player i can indeed believe that his deviation is detected with probability 1, even if M is finite. So that player i is sequentially rational at the described information set, $\boldsymbol{\sigma}^i$ specifies that he stops the game and plays a best response against ζ^j *given his updated belief* over L^j.

If at some stage, player i updates his belief in the way just described and the game goes on, even for infinitely many stages, player i keeps his belief over L^j without modifying it any further. At all other information sets out of the equilibrium path, player i does not update his belief over L^j, stops the game, and plays ζ^i. In particular, consider the following situation: player i first deviates at some stage t, then concludes that $t^\star > t$, but the game nevertheless goes on until stage $t' > t$ and player i again deviates at stage t'; even if the labels at stage t' lead player i to the conclusion that $t^\star = t'$, player i does not update his beliefs over L^j at t', stops the game, and plays ζ^i.[27]

[25]Player i detects that player j has deviated from the plan typically when he receives an incorrect code. If q is such that some actions of player i have zero probability given his type, player i may also detect a deviation of player j from his computed action when he concludes that $t^\star$ has been reached. We focus on the typical case, but the second one can be handled similarly.

[26]To be complete, we must also consider the case of information sets occurring at the second substage of some stage $t \geq 2$, after which both players chose "continue" while the prescribed plan of actions was not followed at an earlier stage. In this case, $\boldsymbol{\sigma}^i$ prescribes to choose a message uniformly and to stop at the next stage.

[27]Player i's beliefs over L^j at stage t' are not restricted by our semi-weak PBE, but are coherent with the belief that player j detects player i's first deviation at stage t, fails to punish player i, and chooses messages uniformly from then on.

Consider now the typical case where player i followed the plan of actions induced by μ at all stages less than or equal to t for some $t \geq 2$, and discovers, through the codes that he receives from player j, that player j did not follow the plan at stage t. This must be player j's first observed defection, since otherwise player i, who followed the scenario, would have stopped the game. As soon as player j deviates from the prescribed scenario, no constraint must be imposed on player i's belief over L^j, which can thus be kept at the prior. In particular, even if only player j's code on player i's encrypted action is incorrect, while player j's reported label λ_t^j coincides with the label λ_t^i that player i received from the correlation device, player i can believe that player j's reported label λ_t^j is, in fact, incorrect and that player j luckily picked his label code. As a consequence, player i can believe that player j did not update his belief over L^i and expects that player j will play ζ^j. Recall that, as detailed above, it is indeed sequentially rational for player j to stop the game and play ζ^j after his deviation in this case. With these beliefs, player i is sequentially rational by stopping the game at the beginning of stage $t+1$ and punishing player j using ζ^i.

The reasoning of the previous paragraphs can be applied to any information set occurring after some finite stage $t \geq 2$. Let us come to the case where cheap talk never stops. If player i always followed the prescribed plan of actions, it means that he did not detect any incorrect code, but could not identify $t^\star$. Player j must thus have deviated from the prescribed scenario, say at stage t', by not reporting the correct $\lambda_{t'}^j$, and must have been lucky in picking the associated code. Player i's belief over L^j is still the prior and, furthermore, player i can expect that player j will play ζ^j because player j himself did not update his belief over L^i. Indeed, player i can believe that player j's first deviation occurred at some stage $t < t'$ such that, just after stage t, player j deduced that $t^\star > t$; player j expected to be punished by player i at stage $t+1$ and having realized that player i did not follow the plan, player j did not update his belief at stage t' (player i's belief is coherent with the fact that player j only updates his belief over L^i at his first deviation, as explained above). Finally, assume again that cheap talk never stops and that player i did not follow the plan at some stage t. If he updated his belief over L^j at the corresponding information set as described above, he kept this belief since stage t and plays a best response against ζ^j given this belief. Otherwise, he plays ζ^i. □

7. Discussion: Variants of the model

We start with a variant of the strategic form correlated equilibria considered up to now. Then we consider two particular cases in which Theorem 1 takes a much simpler form. Finally, we address a question mostly motivated by Ben-Porath (2003, 2006).

7.1 *Extensive form correlated equilibria*

The proof of Theorem 1 makes use of typical correlation devices for the long cheap talk game $\mathrm{ext}_M \Gamma$, which select, before the beginning of the game, an *infinite* sequence of extraneous signals to be used gradually by the players. The corresponding correlated equilibria can be denoted as *strategic form correlated equilibria*. What if the players do not have access to (or cannot generate[28]) infinite sequences of correlated extraneous signals, at once, at the beginning of the game? One could then consider *extensive form, autonomous* correlation devices that send one private signal to every player at every stage of $\text{ext}_M \Gamma$ (see Forges 1986 and Myerson 1986, 1991). Such devices generate sunspots every day. They are independent of the cheap talk game, in the sense that they do not receive any input from the players and do not get any information on the players' messages. They thus preserve the players' privacy. The previous proof shows that Theorem 1 still holds if "correlated equilibrium" is replaced by "extensive form, autonomous correlated equilibrium using *finitely* many signals at every stage." Corollary 1 also holds for the set $\widetilde{\text{CE}}(\text{ext}\,\Gamma)$ of extensive form, autonomous correlated equilibrium payoffs, since $\text{CE}(\text{ext}\,\Gamma) \subseteq \widetilde{\text{CE}}(\text{ext}\,\Gamma) \subseteq \text{ME}(\Gamma)$.

7.2 *Sender–receiver games*

As a particular case, let us assume that only player 1 possesses private information ($|L^2| = 1$) and that only player 2 makes a decision ($|A^1| = 1$). Under these assumptions, the cheap talk game becomes a sender–receiver game, in which the length of the players' conversation is not fixed in advance (as in, e.g., Forges 1990a, Aumann and Hart 2003,[29] Koessler and Forges 2008). We deduce from the proof of Theorem 1 that $t^\star$ can be chosen in a deterministic way, as $t^\star = 1$. Let us set $L = L^1$ and $A = A^2$, and let us consider a correlation device as above, which selects

1. a permutation η of L
2. for every $l \in L$, an action $a_{\eta(l)} \in A$, according to $q(\cdot|l)$
3. for every $l \in L$, a permutation $\phi_{\eta(l)}$ of A; let us set $b_{\eta(l)} = \phi_{\eta(l)}(a_{\eta(l)})$
4. for every $l \in L$ and every action $b \in A$, a code $k(\eta(l), b) \in M$.

The correlation device transmits

— to player 1, η and $(b_{\eta(l)}, k(\eta(l), b_{\eta(l)}))_{l\in L}$

— to player 2, $(\phi_{\eta(l)}, k(\eta(l), \cdot))_{l\in L}$.

Given the signal from the correlation device and his type l, player 1's equilibrium strategy is to send $\eta(l)$, $b_{\eta(l)}$, and $k(\eta(l), b_{\eta(l)})$ to player 2 at a single stage of information transmission. Given his private signal $(\phi_{\eta(l)}, k(\eta(l), \cdot))_{l\in L}$ and player 1's message $(\hat{l}, b, m)$, player 2 checks whether the code is correct, namely that $m = k(\hat{l}, b)$. If it is the case, he chooses the action $(\phi_{\hat{l}})^{-1}(b)$; otherwise, he chooses his action according to $q(\cdot|l)$ for some arbitrary $l \in L$. By proceeding as above, one shows that these correlated strategies form an equilibrium, which is equivalent to the communication equilibrium q. Forges (1985, Lemma 2) establishes a slightly stronger result, namely that *every*

[28]Players can simulate finite correlation devices by themselves by using simple machines (like Turing machines; see Dodis et al. 2000 and Urbano and Vila 2002) or the AND signalling function (see Vida and Āzacis (2012)).

[29]Aumann and Hart (2003) assume one sided private information, namely $|L^2| = 1$, but allow both players to make decisions.

communication equilibrium payoff (even not in SINTIR(Γ)) can be achieved as a correlated equilibrium payoff of the cheap talk game. As already pointed out, Blume (2010) proves an analog in Crawford and Sobel's (1982) model.

7.3 *Uniform punishments*

The proof of Theorem 1 dramatically simplifies if the communication equilibrium payoff of Γ to be implemented as a correlated equilibrium payoff of $\text{ext}_M\,\Gamma$ is higher than a punishment payoff that can be achieved *for every probability distribution* $p \in \Delta L$. This happens, for instance, if Γ has a "bad outcome" that every player can enforce, *whatever their types are*.

More precisely, recalling expression (1), let $G[q] = (G^i[q|l^i]_{l^i \in L^i})_{i=1,2} \in \text{ME}(\Gamma)$ be a communication equilibrium payoff for which there exist $y^i : L^i \to \Delta A^i$, $i = 1, 2$, such that, for every $i = 1, 2$, $l = (l^i, l^{-i}) \in L$, $a^i \in A^i$,

$$G^i[q|l^i] \geq \sum_{a^{-i}} y^{-i}(a^{-i}|l^{-i}) g^i((l^i, l^{-i}), a^i, a^{-i}).$$

Then, in the proof of Theorem 1, to achieve $G[q]$ as a payoff in $\text{CE}(\text{ext}_M\,\Gamma)$, $t^\star$ can be chosen in a deterministic way, as $t^\star = 2$. The correlation device can dispense with selecting the labels and all items associated with $t > 2$. Indeed, if player i's code $k^{-i}(2, \eta(l), b^{-i}_{2,\eta(l)})$ at stage 2 is not correct, player $-i$ can punish him by playing the strategy y^{-i} in Γ, which guarantees that player i's payoff does not exceed $G^i[q|l^i]$, *independently of the information that player i may have acquired at stage 2*, i.e., even if player i learns the type l^{-i} of player $-i$.

However, in many interesting situations, when a player obtains further information on the other's type, it becomes impossible to punish him below his communication equilibrium payoff. This is exactly what happens in the example in Section 5 once a player knows the secret.

7.4 *Cheap talk with delayed messages*

The terminology "cheap talk" is used to cover more or less sophisticated forms of communication between the players. In this paper, we just allow the players to talk for as long as they like by sending simultaneous messages to each other. Bárány (1992) and Ben-Porath (2006) consider more flexible procedures, like the safe recording, at some stage t, of a message that can possibly be released at some further stage t', as a function of the history at stage t'.

If such a relaxed form of cheap talk is allowed in the framework of the current paper, the proof of Theorem 1 can easily be modified so as to achieve every payoff in $\text{ME}(\Gamma) \cap \text{SINTIR}(\Gamma)$ with *only four stages of cheap talk*. To see this, let us slightly modify the correlation device of the proof of Theorem 1 by choosing $t^\star$ uniformly in some finite set T and interpreting it as an index (rather than a stage). At the first stage of cheap talk, the players exchange information $\eta(l)$ on their types as before. Then every player i secretly prepares $|T|$ envelopes, with envelope t containing the encrypted recommended action $b^{-i}_{t,\eta(l)}$ of the other player, its code $k^{-i}(t, \eta(l), b^{-i}_{t,\eta(l)})$, and player i's code function $k^i(t, \eta(l), \cdot)$. At the second stage of cheap talk, the players exchange their extraneous signals on the labels for all $t \in T$ at once (namely $(\lambda^i_t, \kappa^i(t, \lambda^i_t), \kappa^{-i}(t, \cdot))_{t\in T}$). If no deviation is detected at this stage, they identify the index $t^\star$. At the third stage of cheap talk, they reveal to each other the content of all envelopes with index $t \neq t^\star$ and check that the codes are consistent. If again no deviation is detected, they open the two envelopes with index $t^\star$.

The conclusion from this exercise is that allowing delayed messages in cheap talk is by no means innocuous. Indeed, in Section 5, we exhibit a communication equilibrium payoff that cannot be achieved as a correlated equilibrium payoff of any game in which cheap talk lasts for a fixed number of stages and does not involve any delayed message.

References

Abraham, Ittai, Danny Dolev, and Joseph Y. Halpern (2008), "Lower bounds on implementing robust and resilient mediators." In *Proceedings of the Fifth Conference on Theory of Cryptography*, 302–319, Springer, Heidelberg. [107]

Aumann, Robert J. (1974), "Subjectivity and correlation in randomized strategies." *Journal of Mathematical Economics*, 1, 67–96. [96]

Aumann, Robert J. (1987), "Correlated equilibrium as an expression of Bayesian rationality." *Econometrica*, 55, 1–18. [96]

Aumann, Robert J. and Sergiu Hart (2003), "Long cheap talk." *Econometrica*, 71, 1619–1660. [101, 102, 103, 119]

Bárány, Imre (1992), "Fair distribution protocols or how the players replace fortune." *Mathematics of Operation Research*, 17, 327–340. [96, 98, 104, 120]

Ben-Porath, Elchanan (1998), "Correlation without mediation: Expanding the set of equilibrium outcomes by 'cheap' pre-play procedures." *Journal of Economic Theory*, 80, 108–122. [96, 97]

Ben-Porath, Elchanan (2003), "Cheap talk in games with incomplete information." *Journal of Economic Theory*, 108, 45–71. [96, 98, 104, 106, 118]

Ben-Porath, Elchanan (2006), "A correction to 'Cheap talk in games with incomplete information'," *Journal of Economic Theory*, 108, 2003, 45–71. Unpublished paper. [96, 98, 104, 106, 118, 120]

Blume, Andreas (2010), "Correlated equilibria in sender–receiver games." Unpublished paper, University of Pittsburgh. [97, 98, 120]

Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 50, 1431–1451. [97, 120]

Dhillon, Amrita and Jean-François Mertens (1996), "Perfect correlated equilibria." *Journal of Economic Theory*, 68, 279–302. [106]

Dodis, Yevgeniy, Shai Halev, and Tal Rabin (2000), "A cryptographic solution to a game theoretic problem." In *Advances in Cryptology* (Mihir Bellare, ed.), 112–130, Springer, Berlin. [97, 119]

Forges, Françoise (1985), "Correlated equilibria in a class of repeated games with incomplete information." *International Journal of Game Theory*, 14, 129–149. [96, 97, 98, 106, 119]

Forges, Françoise (1986), "An approach to communication equilibria." *Econometrica*, 54, 1375–1385. [95, 99, 119]

Forges, Françoise (1988), "Can sunspots replace a mediator?" *Journal of Mathematical Economics*, 17, 347–367. [96, 98]

Forges, Françoise (1990a), "Equilibria with communication in a job market example." *Quarterly Journal of Economics*, 105, 375–398. [96, 101, 119]

Forges, Françoise (1990b), "Universal mechanisms." *Econometrica*, 58, 1341–1364. [96, 98, 99, 100, 103, 104, 112]

Forges, Françoise (2009), "Correlated equilibrium and communication in games." In *Encyclopedia of Complexity and Systems Science* (Robert A. Meyers, ed.), 1587–1596, Springer, New York. [96, 98]

Fudenberg, Drew and Jean Tirole (1991), *Game Theory*. MIT Press, Cambridge, Massachusetts. [106]

Gerardi, Dino (2004), "Unmediated communication in games with complete and incomplete information." *Journal of Economic Theory*, 114, 104–131. [96, 98, 106]

Gossner, Olivier and Nicolas Vieille (2001), "Repeated communication through the mechanism *and*." *International Journal of Game Theory*, 30, 41–60. [101]

Halpern, Joseph Y. (2008), "Computer science and game theory: A brief survey." In *The New Palgrave Dictionary of Economics*, second edition (Steven N. Durlauf and Lawrence E. Blume, eds.), Palgrave Macmillan, Basingstoke. [96]

Heller, Yuval, Eilon Solan, and Tristan Tomala (2012), "Communication, correlation and cheap-talk in games with public information." *Games and Economic Behavior*, 74, 222–234. [98, 104]

Izmalkov, Sergei, Matt Lepinski, and Silvio Micali (2011), "Perfect implementation." *Games and Economic Behavior*, 71, 121–140. [97]

Koessler, Frédéric and Françoise Forges (2008), "Multistage communication with and without verifiable types." *International Game Theory Review*, 10, 145–164. [119]

Krishna, R. Vijay (2007), "Communication in games of incomplete information: Two players." *Journal of Economic Theory*, 132, 584–592. [97]

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995), *Microeconomic Theory*. Oxford University Press, New York. [106]

Myerson, Roger B. (1986), "Multistage games with communication." *Econometrica*, 54, 323–358. [95, 99, 119]

Myerson, Roger B. (1991), *Game Theory*. Harvard University Press, Cambridge, Massachusetts. [95, 99, 106, 119]

Osborne, Martin J. and Ariel Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts. [106]

Urbano, Amparo and Jose E. Vila (2002), "Computational complexity and communication: Coordination in two-player games." *Econometrica*, 70, 1893–1927. [97, 119]

Urbano, Amparo and Jose E. Vila (2004), "Computationally restricted unmediated talk under incomplete information." *Economic Theory*, 23, 283–320. [97]

Vida, Péter (2007a), "From communication equilibria to correlated equilibria." Unpublished paper, University of Vienna. [97]

Vida, Péter (2007b), *Long Pre-play Communication in Games*. Ph.D. thesis, Universitat Autònoma de Barcelona. [97]

Vida, Péter and Helmuts Āzacis (2012), "A detail-free mediator." Unpublished paper, University of Vienna and Cardiff University. [119]

Submitted 2011-2-16. Final version accepted 2012-1-12. Available online 2012-1-15.