

Games with incomplete information: from repetition to cheap talk and persuasion*

Françoise Forges[†]

December 2019

Abstract

This essay aims at showing that repeated games with incomplete information, which were conceived by Aumann and Maschler in the years 1960s, provide basic tools to study information transmission in static interactive decision problems, both when the agents monitoring information can lie (“cheap talk”) and when they cannot (“persuasion”).

Keywords: Incomplete information, repeated game, sender-receiver game, Bayesian persuasion.

JEL classification: C72, C73, D83, B21.

*Material initially prepared for the Econometric Society Summer School, Hokkaido University, Sapporo, August 6-10, 2019 and the Laffont lecture at the 72nd European Meeting of the Econometric Society, University of Manchester, August 26-30, 2019. I am grateful to In-Koo Cho and Ariel Rubinstein, and to Zvika Neeman for inviting me to these events. I wish to thank Frédéric Koessler, Jérôme Renault, Ludovic Renou and most of all, Andrés Salamanca, for helpful discussions during the preparation of the survey. I benefited from comments of Jean-Jacques Herings, Stephen Morris, Debraj Ray, Bruno Strulovici, Joel Sobel and Peter Sorensen.

[†]Université Paris-Dauphine, PSL University, LEDa. Place du Maréchal de Lattre de Tassigny, 75116 Paris. Email: francoise.forges@gmail.com

1 Introduction

The essay that follows is prompted by recent achievements on models of information transmission, which, in the game-theoretic jargon, distinguish between “cheap talk” and “persuasion”. “Cheap talk” refers to situations in which informed agents can send every signal they wish at no cost, so that they can possibly lie. By contrast, the term “persuasion” is relevant when some agents control the signals that are sent to others – as a function of the available information – but cannot lie. In particular, “Bayesian persuasion” denotes the particular case where some agents can design, in an ex ante utility maximizing way, the information structure to be used later by Bayesian decision-makers.

Starting with by now iconic papers like Crawford and Sobel (1982) on cheap talk and Kamenica and Gentzkow (2011) on Bayesian persuasion, the literature on these topics has become huge. The goal of the essay that follows is just to show that many basic insights have spread from the seminal study of repeated games with incomplete information – initiated by Aumann and Maschler in the late 1960’s – to the systematic analysis of static games with communication.

My overview will be awfully selective, with a bias in favor of my own work. The surveys of Bergemann and Morris (2016b, 2019), Kamenica (2019), Kreps and Sobel (1994), Özdoğan (2016) and Sobel (2013) complement the present one.¹ The results that I will cover fit in the game-theoretic tradition initiated by von Neumann, Nash and Harsanyi, namely, assume that agents – players – have finitely many types and actions but arbitrary utility – “pay-off” – functions.² These players can make use of mixed strategies and are expected utility maximizers.

Here is a short description of the paper. Section 2 explains to which extent a basic technique that was developed to solve zero-sum repeated games with a single informed player – “concavification” – turns out to be also useful to study Bayesian persuasion. Yet incentive compatibility conditions are much weaker in the latter model than in the former. Section 3 describes a char-

¹Kreps and Sobel (1994) insist on the differences between “cheap talk” and “signaling.” Kamenica (2019) concentrates on Bayesian persuasion and its applications. Bergemann and Morris (2019) adopt a general, abstract view of information design. Sobel (2013) and Özdoğan (2016) rather deal with parametrized models, allowing for comparative statics.

²Section 8 is an exception: it is devoted to Crawford and Sobel (1982)’s uniform quadratic example, in which types and actions lie in a real interval.

acterization of Nash equilibria in two-person non-zero-sum repeated games with a single informed player. The characterization necessitates “dimartingales” as a further tool. These are dynamic processes which account for a basic belief feasibility condition (to be fulfilled in any model of information transmission, in particular, Bayesian persuasion) but also for strong, truth telling, incentive compatibility conditions faced by the – strategic – informed player (such conditions do not appear in Bayesian persuasion).

Section 4 considers two-person non-zero-sum repeated games of “pure information transmission,” in which the informed player’s actions are not payoff-relevant. Surprisingly, the characterization of Nash equilibrium outcomes described in Section 3 does not become simpler in this particular class of games: dimartingales cannot be dispensed with. However, as explained in Section 5, under pure information transmission, repeated games with incomplete information are closely related to static games preceded by possibly long cheap talk. Hence dimartingales are also an essential tool to characterize the Nash equilibrium outcomes in the latter class of games and players can achieve more Nash equilibrium outcomes by talking longer.

Section 6 still focuses on games of pure information transmission but is devoted to correlated equilibrium, a solution concept that extends Nash equilibrium by allowing players to observe correlated extraneous signals before the beginning of the game. A simple characterization is proposed for the correlated equilibria of the infinitely repeated game: they are outcome-equivalent to solutions of the static game, the “mediated” equilibria. As a corollary, once the correlated equilibrium is adopted as solution concept, in static games with information transmission, a single phase of cheap talk is enough to obtain the effects of long cheap talk or even of a mediator.

Section 7 summarizes the main single stage information transmission schemes considered in the paper, namely, Bayesian persuasion (Section 2), mediated equilibrium (Section 6) and cheap talk (Section 5). Very recent contributions make it possible to compare these various schemes from the informed player’s viewpoint. This section can be read independently of the developments on repeated games, i.e., directly after Section 2.

Section 8 illustrates the consequences of possibly long cheap talk and mediation in the popular uniform quadratic example of Crawford and Sobel (1982). As Section 7, this section can be followed without entering the details of the other ones. Section 9 goes back to repeated games with incomplete information, to deal with existence issues that were left open in the previous sections. Section 10 concludes with more general models.

2 Concavification as a first illustration

In this section, I first briefly recall the basic features of “Bayesian persuasion.” Most papers on this topic, starting with Kamenica and Gentzkow (2011) cite Aumann and Maschler (1995) for a technique that can be referred to as “concavification.” I compare the way in which this technique is used in Bayesian persuasion and in repeated games with incomplete information.

2.1 Concavification in Bayesian persuasion

In the basic model of “Bayesian persuasion” (Kamenica and Gentzkow (2011)), an individual (who, for reasons to be clear, will be called player 2) has to make a decision whose outcome depends on a state of nature k (which, to fix ideas, lies in a finite set K). This state of nature is unknown to the decision-maker but another individual, the information-designer (or player 1), has the power to design a fully reliable experiment which randomly selects a signal for the decision-maker, as a function of the true state. The decision-maker and the information-designer share a common prior probability distribution p over the possible states (i.e. over the set K), their utility is a function of the state and the decision, the decision-maker maximizes his expected utility at his current belief over the states. This belief is updated as a function of the signal selected by the experiment, using Bayes formula.

More precisely, let us assume that player 2 has to choose an action a_2 in a finite set A_2 . Let us denote his payoff as $V^k(a_2)$ when his action is a_2 and the state of nature is k ; let player 1’s payoff, under the same circumstances, be denoted as $U^k(a_2)$. This defines a basic decision problem, which will be referred to as $DP_0(p)$ all along the paper.

If the decision-maker does not receive any further information on the state of nature k , he just chooses his action a_2 to maximize his expected payoff $\sum_k p^k V^k(a_2)$. Let $u_{NR}(p)$ be the best utility the information-designer can expect by choosing a nonrevealing experiment, which selects signals independently of the state. This is achieved by assuming that, when player 2 is indifferent between two actions, he chooses the one with the highest possible expected utility for player 1.

Kamenica and Gentzkow (2011) show that, by designing his experiment optimally, i.e., to maximize his own expected payoff, the information-designer can achieve the expected utility $cavu_{NR}(p)$, where $cavu$ is the smallest concave function above u . The result is illustrated on a simple example below.

Example 1

Let us assume that there are two states of nature, namely, $K = \{1, 2\}$, and that player 2's set of actions is $A_2 = \{j_1, j_2, j_3, j_4\}$. Let the payoff functions be

$$\begin{aligned} (U^1, V^1)(\cdot) &= \begin{matrix} & j_1 & j_2 & j_3 & j_4 \\ 2, 0 & -1, 3 & 1, 4 & 0, 5 \end{matrix} \\ (U^2, V^2)(\cdot) &= \begin{matrix} & j_1 & j_2 & j_3 & j_4 \\ 0, 5 & 1, 4 & -1, 3 & 2, 0 \end{matrix} \end{aligned}$$

Let $p = (p^1, p^2)$ be the common prior probability over types (hence $p^2 = 1 - p^1$). The optimal decisions of player 2, as a function of p^1 , are: j_1 if $0 \leq p^1 \leq \frac{1}{4}$, j_2 if $\frac{1}{4} \leq p^1 \leq \frac{1}{2}$, j_3 if $\frac{1}{2} \leq p^1 \leq \frac{3}{4}$, j_4 if $\frac{3}{4} \leq p^1 \leq 1$. This generates the following mapping³ u_{NR} :

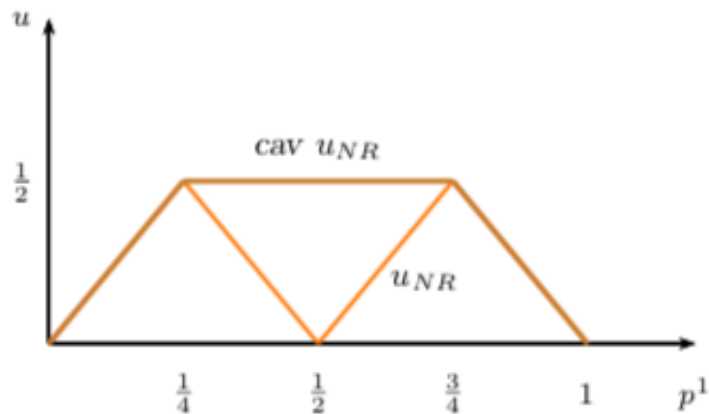
$$u_{NR}(p) = \begin{matrix} 2p^1 & \text{if } 0 \leq p^1 < \frac{1}{4} \\ 1 - 2p^1 & \text{if } \frac{1}{4} \leq p^1 < \frac{1}{2} \\ 2p^1 - 1 & \text{if } \frac{1}{2} \leq p^1 < \frac{3}{4} \\ 2(1 - p^1) & \text{if } \frac{3}{4} \leq p^1 \leq 1 \end{matrix}$$

while the mapping $cav u_{NR}$ is:

$$cav u_{NR}(p) = \begin{matrix} 2p^1 & \text{if } 0 \leq p^1 \leq \frac{1}{4} \\ \frac{1}{2} & \text{if } \frac{1}{4} \leq p^1 \leq \frac{3}{4} \\ 2(1 - p^1) & \text{if } \frac{3}{4} \leq p^1 \leq 1 \end{matrix}$$

The mappings are represented below.

³In this example, the mapping u_{NR} is continuous, which is a bit peculiar (see the remark at the end of this section; see Section 7 for an example in which u_{NR} is just upper-semi-continuous).



A nonrevealing experiment enables player 1 to achieve $cav u_{NR}(p)$ when $p^1 \notin (\frac{1}{4}, \frac{3}{4})$. If $p^1 \in (\frac{1}{4}, \frac{3}{4})$, player 1 designs a state-dependent experiment, so that player 2 computes the posterior probabilities $\frac{1}{4}$ and $\frac{3}{4}$. For instance, when $p^1 = \frac{1}{2}$, such an experiment selects one of two signals, H or L , H is selected with probability $\frac{1}{4}$ when the state is 1 and with probability $\frac{3}{4}$ when the state is 2.

Sobel (2013, footnote 27) notes: “Aumann and Maschler (1995) contained the basic mathematical result in their analysis of repeated games with incomplete information.”

Kamenica and Gentzkow (2011) give a slightly more detailed account in their introduction: “the fact that the informed player’s initial actions have no impact on his long-run average payoffs (and can thus be treated as nothing but a signal) combined with a focus on Nash equilibria (which implicitly allow for commitment [sic]) makes Aumann and Maschler’s problem mathematically analogous to ours.”

We will show below that the informed player’s actions indeed have a signaling role in Aumann and Maschler (1966)’s model and that Kamenica and Gentzkow (2011)’s problem is analogous to the *very first part* of Aumann and Maschler’s one. However the fact that Nash equilibria allow for commitment is disputable.

2.2 Concavification in repeated games with incomplete information

The quotation “Aumann and Maschler (1995)” refers to a book, which gathers reports that were written much earlier, from 1966 until 1968, partly in collaboration with R. Stearns.⁴ These reports were prepared at the initiative of the US Disarmament Agency in the middle of the cold war.⁵ Aumann and Maschler were thus motivated by very practical issues, which were definitely more conceptual than mathematical. A main theme was that by making use of their information, the US forces would also reveal it, which could be disastrous in the long run. It is true that Aumann and Maschler identified many different facets of the strategic use of information, which called for very diverse mathematical tools. This left their followers⁶ with a number of challenges, many of which were solved with the help of rather sophisticated mathematical techniques (see Mertens, Sorin and Zamir (2015)’s book).

Let me go on with some details on Aumann and Maschler (1966)’s results. They start with a standard two-person game. Let A_1 (resp., A_2) be the finite set of actions of player 1 (resp., 2). The payoffs depend on a state of nature k , which can take finitely many values. Let us denote them as $U^k(a_1, a_2)$ and $V^k(a_1, a_2)$ for player 1 and player 2 respectively, when the state is k and the actions $(a_1, a_2) \in A = A_1 \times A_2$. The players share a common prior p over the states. A main difference with the model of the previous subsection is that, here, player 1 knows the state of nature, which can thus be referred to as player 1’s *type*, in the sense of Harsanyi (1967). It is convenient to represent the players’ asymmetric information by making the game start with a move of nature selecting the state k , player 1 being the only one to be informed of k . Then, at every stage, the players simultaneously choose actions, which are observed by both of them, stage after stage. Payoffs are undiscounted, i.e., evaluated as limits of averages. This defines the game $\Gamma_\infty(p)$.

Let $f_{NR}(p)$ be the best expected payoff player 1 can guarantee himself, whatever player 2’s reaction, by playing independently of the state. $f_{NR}(p)$ can be computed as the minmax level of player 1 in the expected one-shot

⁴Each chapter of Aumann and Maschler (1995) is followed by “postscripts” summarizing the main results obtained between 1966 and 1995.

⁵This is acknowledged in Kamenica (2019) in “a few words on the intellectual history of information design and the concavification approach.”

⁶Just to mention a few from the “first generation:” S. Hart, R. Simon, S. Sorin, J.-F. Mertens, S. Zamir.

game.⁷ Aumann and Maschler (1966) show that, in $\Gamma_\infty(p)$, by making use of his information, player 1 can guarantee himself $cavf_{NR}(p)$, and no more, against any strategy of player 2.⁸

One of the keys to the result is that, in the absence of discounting, player 1's actions at the early stages of the game can be interpreted as costless signals. Player 1 can thus modify player 2's probability distribution over the states by choosing his actions s for the first n stages of the game (in $S = (A_1)^n$, for some $n = 1, 2, \dots$) by means of a type-dependent mixed strategy. Which posteriors p_s can be reached in this way? All the p_s 's whose expectation, with respect to some probability distribution over the signals, is the prior p . This characterization, which is just a straightforward consequence of Bayes formula, is known as the “splitting lemma” in the literature on repeated games with incomplete information (see, e.g., Sorin (2002)). As understood in the previous subsection, it is also a corner stone in “Bayesian persuasion.” Using the splitting lemma, Aumann and Maschler (1966) first establish that if player 1 can guarantee himself some $f(p)$ in $\Gamma_\infty(p)$, he can as well guarantee himself $cavf(p)$.

This looks at first sight quite similar to what we have seen for Bayesian persuasion. There are however a number of differences, which reflect that $u_{NR}(p)$ and $f_{NR}(p)$ are computed in quite different games. First of all, in Aumann and Maschler (1966)'s framework, the state of nature corresponds to player 1's “type”. And player 1 is a strategic player, who cannot commit to randomly select his action as a function of the true state of nature when he maximizes his payoff against player 2. In particular, player 1's signaling (minmax) strategy is not only constrained by the straightforward condition on posteriors but also by incentive compatibility conditions, which guarantee that he cannot profit from lying about his type.⁹

By contrast, an implicit condition behind Bayesian persuasion is that the information-designer chooses his experiment without knowing the state

⁷A precise formula is: $f_{NR}(p) = \min_{\tau \in \Delta(A_2)} \max_{\sigma \in \Delta(A_1)} (\sum_{k \in K} p^k U^k(\sigma, \tau))$.

⁸In the previous presentation, there is no restriction on the payoff matrices, so that $\Gamma_\infty(p)$ can be viewed as a non-zero-sum game. $f_{NR}(p)$ (resp., $cavf_{NR}(p)$) is the minmax level of player 1 in the expected one-shot game (resp., in $\Gamma_\infty(p)$), where player 1 is the maximizer. In Aumann and Maschler (1966)'s original model, the game is assumed to be zero-sum and player 1 is the maximizer.

⁹Player 1's incentive compatibility conditions somehow come for free in Aumann and Maschler (1966)'s analysis, whose aim is to characterize the informed player's individually rational level. We will be more precise on incentive compatibility conditions in the next section.

of nature and can make sure that the experiment be fully reliable, namely, selects the signal as a function of the true state. Another difference is that Aumann and Maschler (1966) are interested in player 1's minmax level, so that player 2 appears as player 1's opponent.

But the most important difference between the two frameworks is that, in Aumann and Maschler (1966), it is far from obvious that player 1 cannot guarantee *more* than $cav f_{NR}(p)$. Actually, easy examples (like the one below) illustrate that it may happen that player 1 can do much better in the *finitely* repeated game (the same holds in the discounted infinitely repeated game). A result of Blackwell (1956) can be used to show how, in the undiscounted *infinitely* repeated game, the uninformed player can keep the payoff of the informed one below $cav f_{NR}(p)$.

Example 2

Player 1 is informed; he has two types, $K = \{1, 2\}$, and two actions H and L . Player 2 is uninformed and has three actions ℓ , c and r . As in Example 1, let $p = (p^1, p^2)$ be the common prior probability over types. The game is zero-sum, the utility function of player 1 is:


		ℓ	c	r
$k = 1$	H	2	0	1
	L	2	0	-1
		ℓ	c	r
$k = 2$	H	0	2	-1
	L	0	2	1

If the informed player does not make use of his private information, he gets, at best,

$$f_{NR}(p) = \min \{2p^1, 2(1 - p^1), \max \{2p^1 - 1, 1 - 2p^1\}\}$$

It can be checked that $f_{NR}(p)$ coincides with the mapping $u_{NR}(p)$ computed in Example 1 (see the figure below).¹⁰

¹⁰The mapping $f_{NR}(p)$ is continuous, a property that always holds in Aumann and Maschler's framework. In Bayesian persuasion, if player 2 has finitely many actions and breaks ties in favor of player 1, u_{NR} is upper-semi-continuous (but not necessarily continuous, see Section 7).



In the one-shot game, player 1 can make full use of his information without taking any risk, namely, play H (resp., L) when his type is 1 (resp., 2); he gets, at best,

$$f_1(p) = \min \{2p^1, 2(1 - p^1)\}.$$

In the infinitely repeated game, player 1's first move can be used as a signal. According to Aumann and Maschler (1966) he gets, at best, the level $cav f_{NR}(p) = cav u_{NR}(p)$ computed in Example 1.

We observe that

$$\text{for every } p \text{ s.t. } p^1 \in \left(\frac{1}{4}, \frac{3}{4}\right), f_1(p) > cav f_{NR}(p).$$

More generally, one can show that player 1's minmax levels in the n times repeated game define a sequence of concave mappings $f_n(p) \geq cav f_{NR}(p)$, which converge to $cav f_{NR}(p)$ (see, e.g., Sorin (2002)).

3 Joint plans and dimartingales

In this section, I describe what I view as the main achievement in Aumann and Maschler (1995), namely, fundamental steps toward the characterization of the Nash equilibria of the undiscounted infinitely repeated nonzero-sum

game $\Gamma_\infty(p)$ introduced in the previous section. This is the topic of Aumann, Maschler and Stearns (1968), reprinted as the last chapter of Aumann and Maschler (1995). With this view, the results of Aumann and Maschler (1966) and Blackwell (1956) appear as auxiliary tools to compute the players' individually rational levels in the game $\Gamma_\infty(p)$.

The name “Folk theorem” had not yet been coined in the 1960's, but Aumann, Maschler and Stearns knew that, under complete information, the Nash equilibrium payoffs of an undiscounted infinitely repeated game coincide with the feasible individually rational payoffs which, in this case, are defined in the one-shot game.

Equipped with a double background (Folk theorem under complete information and individually rational levels under incomplete information), Aumann, Maschler and Stearns (1968) start by characterizing the “nonrevealing” Nash equilibrium payoffs in the game $\Gamma_\infty(p)$. A noticeable feature is that these payoffs are sustained by strategies of the informed player that are nonrevealing on the equilibrium path but may very well be type-dependent off path, i.e., to punish player 2 by means of a trigger strategy keeping him below his individually rational level.¹¹

Aumann, Maschler and Stearns (1968) propose easy examples to show that a nonrevealing equilibrium may not exist for some values of the prior probability distribution p . Furthermore, a nonrevealing Nash equilibrium of the game with incomplete information $\Gamma_\infty(p)$, when it exists, may not reflect a lesson that is familiar under complete information, namely, that the repetition of a game enables the players to cooperate. Aumann, Maschler and Stearns (1968) go on with “joint plan equilibria” in which the informed player first sends a signal s to the other one, who updates his prior p to a posterior p_s . Then, given the signal s , the players play a nonrevealing equilibrium of $\Gamma_\infty(p_s)$. As in the previous section, the absence of discounting enables the informed player to use his early actions as costless signals.

A “signaling strategy” takes the form of a mixed strategy $\sigma(\cdot | k)$ to choose a signal for every type k . It is formally identical to the experiment of an information-designer in Bayesian persuasion. The posteriors p_s are just constrained by the fact that their expectation is the prior p (splitting). But here, player 1 chooses his signal by himself, which implies very demanding

¹¹Let $g_{NR}(p)$ be the best expected payoff player 2 can guarantee himself if player 1 does not make use of his information. Player 2's individually rational level is $ve x g_{NR}(p)$, where $ve x g$ is the largest convex function below g .

incentive compatibility conditions. Indeed, if being of type k , player 1 sends signals s and s' with positive probability (namely, $\sigma(s | k) > 0$, $\sigma(s' | k) > 0$), he must be indifferent between s and s' . Denoting as x_s^k the expected payoff of player 1 of type k after having sent signal s , we must have $x_s^k = x_{s'}^k$.¹²

Summing up loosely, let $(x, \beta) = ((x^k)_{k \in K}, \beta)$ be a pair of interim expected payoffs for player 1 and player 2 respectively¹³; (x, β) is a joint plan equilibrium expected payoff if and only if there exist a probability distribution over signals and for every signal s , a posterior p_s over the states, an interim payoff $x_s = (x_s^k)_{k \in K}$ for player 1 and an expected payoff β_s for player 2 such that (i) for every s , (x_s, β_s) is a nonrevealing equilibrium payoff of $\Gamma_\infty(p_s)$ (ii) $x_s = x_{s'}$ for every s, s' and (iii) (p, x, β) is the expectation of (p_s, x_s, β_s) , in particular, the prior p is the expectation of the posteriors p_s .

Having characterized the joint plan equilibria, Aumann, Maschler and Stearns (1968) propose examples of games with further equilibrium payoffs, which cannot be achieved by a joint plan (and a fortiori, in a nonrevealing way). For this, they first identify a new coordination procedure that is made feasible in the undiscounted repeated game $\Gamma_\infty(p)$: the jointly controlled lottery.

In a jointly controlled lottery, which can take place any time in $\Gamma_\infty(p)$, player 1 does not make use of his information and both players choose their action by means of a mixed strategy. To give the simplest example, suppose that at some stage of $\Gamma_\infty(p)$, both players select independently of each other their first action (“1”) or their second action (“2”), with equal probability $\frac{1}{2}$. Then, as soon as one of the players randomizes with equal probability over “1” and “2”, the probability that the outcome is “11” or “22” is $\frac{1}{2}$ (and similarly for “12” or “21”). The procedure is immune to unilateral deviations and is easily generalized to any probability distribution. It enables the players to generate by themselves the effect of an extraneous public device.¹⁴ Given two equilibrium payoffs (x, β) and (x', β') of $\Gamma_\infty(p)$, the players can thus decide to achieve (x, β) with some probability π and (x', β') with probability $1 - \pi$.

¹²At first sight, we need to add conditions of the form $x_s^k \geq x_{s'}^k$ if $\sigma(s | k) > 0$, $\sigma(s' | k) = 0$. These inequalities can be turned into equalities by introducing fictitious payoffs (see Hart (1985) and e.g., Forges (1994), Peski (2014)).

¹³For the informed player 1, we consider interim payoffs, indexed by his types, namely, $(x^k)_{k \in K}$, with K denoting the set of all states of nature.

¹⁴See, e.g., Matthews and Postlewaite (1989) and Celik and Peters (2016) for applications of jointly controlled lotteries to auctions and oligopoly, respectively.

In other words, the set of all equilibrium payoffs of $\Gamma_\infty(p)$ is convex.¹⁵ But Aumann, Maschler and Stearns (1968) find examples of equilibrium payoffs which cannot be achieved as convex combinations of joint plan equilibria, namely, equilibrium payoffs that require two stages of signaling from the informed player.

Aumann, Maschler and Stearns (1968) end up with two open questions:

- (1) What is the full characterization of the Nash equilibrium payoffs of the infinitely repeated game $\Gamma_\infty(p)$?
- (2) Does a joint plan equilibrium always exist?

Hart (1985) answers the first one.¹⁶ To do this, he first observes that, as suggested above, a very special kind of random variable, with values (p_s, x_s, β_s) , can be associated with a joint plan equilibrium of $\Gamma_\infty(p)$. First, (p, x, β) is the expectation of the random variable, which can thus be interpreted as the first stage of a martingale.¹⁷ Second, condition (ii) above states that player 1's interim payoff is constant. This further property accounts for the incentive compatibility conditions of the informed player, who cannot rely on a safe experiment but has to send his signal by himself.

Suppose now that, having reached some posterior probability p_s , the players perform a jointly controlled lottery to decide on how to play in the future. Since the informed player does not make use of his information in a jointly controlled lottery, the probability distribution p_s over the states cannot change. But the payoffs of both players do change, as they depend on the outcome of the jointly controlled lottery.

Hart (1985) proves that all equilibrium payoffs of $\Gamma_\infty(p)$ can be achieved by alternating, possibly ad infinitum, signaling stages – at which the posterior

¹⁵It has become customary in the literature on infinitely repeated games to assume at the outset that a public randomization device is available. In Aumann, Maschler and Stearns (1968), the assumption is indeed without loss of generality.

¹⁶During the 1970's, the work on infinitely repeated games with incomplete information was mostly pursued by Mertens and Zamir, who investigated further the zero-sum case, possibly with lack of information on both sides (see Section 10). Research on non-zero-sum infinitely repeated games with incomplete information started again during the spring 1980 at the Institute for Advanced Study of Jerusalem. Hart (1985) first appeared in 1982, as CORE discussion paper 8203.

¹⁷Once again, this is the “splitting lemma”, namely, the not so demanding constraint on posteriors.

changes but player 1’s interim payoff remains constant – and jointly controlled lotteries – at which the posterior remains constant but player 1’s interim payoff changes.

More precisely, Hart (1985) defines a dimartingale¹⁸ $(\tilde{p}_t, \tilde{x}_t, \tilde{\beta}_t)$ as a martingale¹⁹ such that, at every stage t , either $\tilde{x}_{t+1} = \tilde{x}_t$ or $\tilde{p}_{t+1} = \tilde{p}_t$. Hart (1985)’s characterization states that (x, β) is an equilibrium payoff of $\Gamma_\infty(p)$ if and only if there exists a dimartingale $(\tilde{p}_t, \tilde{x}_t, \tilde{\beta}_t)$ whose expectation is (p, x, β) and which becomes nonrevealing at the limit (namely, the limit $(\tilde{p}_\infty, \tilde{x}_\infty, \tilde{\beta}_\infty)$ is such that $(\tilde{x}_\infty, \tilde{\beta}_\infty)$ is a nonrevealing equilibrium payoff of $\Gamma_\infty(\tilde{p}_\infty)$). Aumann and Hart (1986)²⁰ offer a thorough treatment of dimartingales as mathematical objects and construct a geometric example – the “four frogs” – of a converging dimartingale that does not reach its limit within a bounded number of stages.

4 Repeated games of pure information transmission

An early version of Hart (1985)’s characterization of Nash equilibrium payoffs in $\Gamma_\infty(p)$ was already available in the summer 1981. My thesis advisor, J.-F. Mertens, suggested then the characterization of correlated equilibrium payoffs in $\Gamma_\infty(p)$ as a promising research topic. By applying Aumann (1974)’s definition, a correlated equilibrium of the game $\Gamma_\infty(p)$ is a Nash equilibrium of an extension of $\Gamma_\infty(p)$ in which, before the beginning of the game, the players privately observe correlated signals. These signals are generated by an extraneous device so that player 1 gets the same signal whatever his type.

Given the possible complexity of equilibrium behavior in $\Gamma_\infty(p)$, I first investigated games that were simpler than the ones of Hart (1985), namely, games in which the informed player’s actions have no impact on the players’ payoffs. In this case, the one-shot game reduces to the decision problem $DP_0(p)$ introduced in Section 2. Recall that the payoff $V^k(a_2)$ of the uninformed decision-maker, player 2, depends on his action a_2 and the state of nature k , which here, is known to player 1. The payoff $U^k(a_2)$ of the latter

¹⁸The terminology was then “bimartingale” (see Aumann and Hart (2003)).

¹⁹A martingale is a random process with constant conditional expectation, i.e., denoting the past events up to stage t by H_t , $E((\tilde{p}_{t+1}, \tilde{x}_{t+1}, \tilde{\beta}_{t+1}) | H_t) = (\tilde{p}_t, \tilde{x}_t, \tilde{\beta}_t)$.

²⁰As for Hart (1985), the results were available much earlier in a preprint.

depends on the same parameters. The players have the same common prior p over the states of nature.

If the decision-maker does not receive any further information on the state of nature k , he just chooses his action a_2 to maximize his expected payoff $\sum_k p^k V^k(a_2)$. This is also the starting point in Bayesian persuasion. But here there is no reason to assume that the decision-maker picks the informed player's preferred action when he is himself indifferent. We will keep track of all *interim* expected payoffs $(U^k(a_2))_{k \in K}$ that are associated to optimal decisions a_2 of player 2 in $DP_0(p)$. These are player 1's nonrevealing equilibrium payoffs in the infinitely repeated game $\Gamma_\infty(p)$, in which player 1's actions are just costless signals.

In other words, a nonrevealing equilibrium does exist in $\Gamma_\infty(p)$ for every prior p , a property that does not hold in the more general model of Aumann, Maschler and Stearns (1968). What really simplifies the analysis of games of pure information transmission is that, in these games, the players' individually rational levels in $\Gamma_\infty(p)$ and in the one-shot decision problem $DP_0(p)$ are identical. This is quite intuitive, since the informed player's actions have no effect on the payoffs. To see this formally, recall from Section 2 that the informed player's individually rational level in $\Gamma_\infty(p)$ is $cav f_{NR}(p)$, where $f_{NR}(p)$ is the best payoff player 1 can expect without revealing information in $\Gamma_\infty(p)$, if player 2 behaves as an opponent. In the current model, $f_{NR}(p) = \min_{a_2} \sum_k p^k U^k(a_2)$ is concave (as a minimum of linear functions), so that revealing information can only hurt player 1 when he defends himself against player 2.

Summing up, if we keep viewing the above game $\Gamma_\infty(p)$ – of pure information transmission – as an infinitely repeated game, nonrevealing (or “babbling”) equilibrium payoffs exist whatever the prior p and are characterized in a straightforward way. These two properties need not hold in the more general model considered by Aumann, Maschler and Stearns (1968) and Hart (1985).

However, even if infinitely repeated games of pure information transmission look elementary, their equilibria can have a quite complex structure. This is illustrated in Forges (1984).²¹ First of all, some of these equilibrium payoffs cannot be reached unless two stages of signaling are performed. As pointed out above, the phenomenon was already identified by Aumann, Maschler and Stearns (1968), but in games with payoff-relevant actions for the informed

²¹CORE Discussion paper 8220.

player. More surprisingly, the “four frogs” example of Aumann and Hart (1986) can be generated in an infinitely repeated game of pure information transmission. In other words, to achieve some equilibrium payoffs, the players have to alternate signaling and jointly controlled lotteries, without being able to determine an upper bound on the number of stages that will be needed to reach a nonrevealing equilibrium.

5 Two-stage signaling and long cheap talk

The examples in Forges (1984) can easily be re-interpreted as illustrations of the role of two-stage signaling and “long cheap talk” in the basic decision problem $DP_0(p)$. To see this, let us describe the n stage cheap talk game $CT_n(p)$ based on $DP_0(p)$ as follows: player 1’s type k is first chosen according to the prior probability p ; then, at every stage $t = 1, \dots, n$, player 1 and player 2 simultaneously send a message to each other; finally, player 2 makes a decision a_2 and the payoffs are $U^k(a_2), V^k(a_2)$ for player 1 and player 2 respectively.

We have seen above that the nonrevealing equilibrium payoffs of $\Gamma_\infty(p)$ coincide with the equilibrium payoffs of the decision problem $DP_0(p)$, in which the uninformed player just maximizes his expected payoff at the prior p . When $n = 1$, $CT_1(p)$ amounts to a sender-receiver game in which the informed player sends a message to the decision-maker.²² It is not difficult to see that the joint plan equilibrium payoffs of $\Gamma_\infty(p)$ coincide with the equilibrium payoffs of the sender-receiver game $CT_1(p)$. More interestingly, the equilibrium payoffs of $\Gamma_\infty(p)$ achieved with two stages of signaling and a jointly controlled lottery in between can be achieved as equilibrium payoffs of the three stage cheap talk game $CT_3(p)$.

Given the previous straightforward interpretation, Example 1 in Forges (1984) shows that the equilibrium payoff of the informed player in $CT_3(p)$, whatever his type, is higher than any payoff he can get in the sender-receiver game $CT_1(p)$. In other words, the informed player benefits from two stages of signaling, whatever his type.

As already mentioned above, another example of Forges (1984) shows that Aumann and Hart (1986)’s “four frogs” can be generated from a game. This example thus demonstrates an equilibrium payoff of $\Gamma_\infty(p)$ in which the

²²In $CT_1(p)$, jointly controlled lotteries do not generate more payoffs than the uninformed player’s mixed strategies.

players do not reach a nonrevealing equilibrium within a bounded number of stages $\Gamma_\infty(p)$. Re-interpreted in the framework of the cheap talk games based on $DP_0(p)$, this example proposes an outcome that could not be achieved in $CT_n(p)$, whatever the finite number of cheap talk stages n . However, this outcome can be achieved with finitely many stages of cheap talk, provided that no deadline n is fixed in advance, i.e., in a long (but finite) cheap talk game.

Following a suggestion of R. Aumann in the spring 1986, I rewrote this example to give the players' types and actions a concrete meaning. In Forges (1990b), the decision-maker is an employer, the informed player is a job candidate whose type is not the traditional high or low quality but rather reflects different tastes for diversified tasks. The informed player is shown to dramatically improve his expected payoff, whatever his type, by engaging in long, rather than bounded, cheap talk. Simon (2002) proposes a variant of Forges (1990b) in which a long cheap talk equilibrium ex ante Pareto-dominates any equilibrium that can be achieved within a bounded number stages of cheap talk.

Aumann and Hart (2003) consider the effects of allowing long, possibly unbounded, bilateral cheap talk in a static game that is more general than the above decision problem $DP_0(p)$. There is still a single informed player but both players have payoff-relevant actions. The basic model is thus similar to the one-shot game of Aumann, Maschler and Stearns (1968) and Hart (1985). The players can exchange costless messages for as long as they want before choosing their actions. It is understood that the players can send their messages simultaneously, so that this is not "polite talk," as pointed out by Aumann and Hart (2003). They characterize the set of all equilibrium payoffs that can be achieved in this way in terms of the dimartingales introduced in Hart (1985) and Aumann and Hart (1986). The model differs nevertheless from the infinitely repeated game $\Gamma_\infty(p)$. For instance, existence of nonrevealing equilibria is guaranteed at the outset but ever lasting cheap talk is hard to justify.²³

Forges and Koessler (2008) show that Forges (1990b) and Aumann and Hart (2003)'s basic insights survive in a variant of the decision problem

²³Ever lasting cheap talk corresponds to a dimartingale that converges without reaching its limit in finitely stages. The characterization does not rule out this possibility even if no example illustrates this pattern in a game. Indeed, in the examples of Forges (1984, 1990b), the underlying dimartingale attains its limit within a finite but not bounded number of stages.

$DP_0(p)$, in which the informed player’s type is certifiable.²⁴ This means that player 1 cannot lie on his type but can hide part of his information. Forges and Koessler (2008) characterize the set of equilibrium payoffs that can be achieved by long bounded cheap talk, using the concepts of diconvexification and dimartingale. They propose an example in which delaying information certification benefits the informed player, whatever his type, compared to all equilibria of the game with a single, unilateral signaling stage.

6 Implementation of a mediator by cheap talk

At the beginning of Section 4, I explained that my aim (in the fall 1981) was to characterize the set $\mathcal{C}(\Gamma_\infty(p))$ of correlated equilibrium payoffs of every infinitely repeated game $\Gamma_\infty(p)$ of pure information transmission. To characterize this set, Forges (1985)²⁵ considers another solution concept, the communication equilibrium²⁶, in the basic one-shot decision problem $DP_0(p)$.

A (canonical) communication device δ for $DP_0(p)$ consists of probability distributions $\delta(\cdot | k)$, $k \in K$, over the decision-maker actions. Adding a communication device δ to the basic decision problem $DP_0(p)$ gives rise to a game with communication, in which player 1 is invited to report a type k' to the communication device which then recommends an action a'_2 to player 2 according to $\delta(\cdot | k')$. The communication device δ defines an equilibrium if in the game with communication induced by δ , it is an equilibrium for player 1 to report his type truthfully and for player 2 to play the recommended action ($k' = k$, $a'_2 = a_2$).²⁷ A communication device is usually interpreted as a mediator (see Forges (1986), Myerson (1986) and Myerson (1991), Section 6.3). We thus denote as $\mathcal{M}(DP_0(p))$ the set of all communication equilibrium payoffs of $DP_0(p)$. This set is tractable, as it is described by a finite set of linear inequalities (see (3.5) to (3.8) in Forges (1985)).

Forges (1985) establishes that, if $\Gamma_\infty(p)$ is a game of pure information

²⁴Following a popular terminology at the time, Forges and Koessler (2008) refer to games with certifiable types as to “persuasion games”, not to be confused with “Bayesian persuasion” (see, e.g., Özdoğan (2016) for a classification of the different models).

²⁵CORE Discussion Paper 8218.

²⁶The terminology adopted in Forges (1985) is “noisy channel” equilibrium (see also Myerson (1982, 1986, 1991) and the discussion below).

²⁷The difference between a communication device and an experiment in Bayesian persuasion is that, in the latter, the information-designer does not have to report the state of nature.

transmission, $\mathcal{C}(\Gamma_\infty(p)) = \mathcal{M}(DP_0(p))$. This shows that correlated equilibrium payoffs of $\Gamma_\infty(p)$ are characterized in terms of solutions of the static decision problem $DP_0(p)$, as in the Folk theorem for games with complete information. In particular, no dimartingale is needed to achieve any correlated equilibrium payoff, as opposed to what may happen for Nash equilibrium payoffs (see Section 5, recall that the examples of Forges (1984, 1990b) are games of pure information transmission²⁸).

The proof of the previous result proceeds in two steps. First it is shown that a very large superset of $\mathcal{C}(\Gamma_\infty(p))$, the set of all communication equilibrium payoffs of $\Gamma_\infty(p)$, in which players make inputs and receive outputs *at every stage*, is included in $\mathcal{M}(DP_0(p))$. The rough intuition for this part is the same as in the revelation principle (but some care is needed to handle payoffs within an infinite horizon). The same kind of argument shows that every equilibrium payoff achieved in some possible noncooperative extension of $DP_0(p)$ enabling the players to communicate (for as long as they want, with or without a mediator) must be in $\mathcal{M}(DP_0(p))$.

The proof of the reverse inclusion, namely, that $\mathcal{M}(DP_0(p))$ is included in $\mathcal{C}(\Gamma_\infty(p))$, is fully constructive and uses strategies in $\Gamma_\infty(p)$ that are so simple that they can be described in the sender-receiver game $CT_1(p)$ based on $\mathcal{M}(DP_0(p))$, which was introduced in the previous section. More precisely, lemma 2 in Forges (1985) states that every communication equilibrium payoff of $DP_0(p)$ can be achieved as a correlated equilibrium payoff of a sender-receiver game based on $DP_0(p)$, in which the informed player sends a single message in a finite set S to the decision-maker. The previous lemma can thus be interpreted independently of the infinitely repeated game framework, namely, in the static decision problem $DP_0(p)$.²⁹

Not surprisingly, the previous characterization of correlated equilibrium payoffs does not hold for Nash equilibrium. This is confirmed by an example in Forges (1985), in which a communication equilibrium payoff (in

²⁸The basic motivation for these examples was indeed to illustrate the difference between Nash equilibria and correlated equilibria in games of pure information transmission.

²⁹The size of the set of messages, although finite, depends on the underlying communication equilibrium. The result is thus easier to formulate in $\Gamma_\infty(p)$. To implement all communication equilibrium payoffs in a single sender-receiver game, one can assume that the set S of messages allowed in $CT_1(p)$ is countable. Then the lemma can be stated as (i) $\mathcal{M}(DP_0(p)) \subseteq \mathcal{C}(CT_1(p))$ (ii) the informed player's strategy used to implement a particular communication equilibrium only uses finitely many messages. By denoting as $CT_\infty(p)$ the long cheap talk game (see Aumann and Hart (2003)), Forges (1985)'s result can be stated as: for every $n = 2, 3, \dots$, $\mathcal{C}(CT_1(p)) = \mathcal{C}(CT_n(p)) = \mathcal{C}(CT_\infty(p)) = \mathcal{M}(DP_0(p))$.

$\mathcal{M}(DP_0(p))$) cannot be achieved as a Nash equilibrium payoff of the associated repeated game $\Gamma_\infty(p)$, or, equivalently, as a Nash equilibrium payoff of $DP_0(p)$ preceded by long cheap talk. In particular, this communication equilibrium payoff cannot be achieved as a Nash equilibrium of any cheap talk game $CT_n(p)$, whatever the number of stages n , or $CT_\infty(p)$, defined as in Aumann and Hart (2003). The same phenomenon arises in Forges (1990b)'s example. Yet Forges (1985) shows that if the players can observe signals (that are correlated with each other but independent of the state of nature), they can implement a mediator by a single stage of information transmission.³⁰

Quite some time after the previous characterization result, implementation of a mediator by cheap talk became an active research topic, specially in situations involving more than two players, both in game theory and computer science (see, e.g., Forges (2009), Forges (2010) and Halpern (2008) for surveys). If there are at least four players, Forges (1990a) shows that the communication equilibrium outcomes of every Bayesian game can be implemented as Nash equilibrium outcomes of the extended game in which the players, knowing their type, can exchange costless messages before making decisions. Communication equilibrium outcomes are implemented by cheap talk in two steps: first, as correlated equilibrium outcomes (i.e., as in Forges (1985)) and then, using a result of Bárány (1992)³¹, as Nash equilibrium outcomes (see also Ben-Porath (2003) and Gerardi (2004)). As seen in the previous paragraph, with only two players, such a result cannot be true, even if the underlying Bayesian game reduces to the decision problem $DP_0(p)$.

Vida (2007) (see Vida and Forges (2013)) implements a mediator by correlated cheap talk (i.e., extends Forges (1985)'s result) in arbitrary (static) two-person games, in which both players have private information and payoff-relevant actions. To state the result precisely and connect it to the previous ones, suppose that, before engaging in long cheap talk, the players can privately observe signals that are correlated with each other but independent of the players' types. By appealing to a general form of the revelation principle, every equilibrium payoff of the corresponding extended game can be achieved by means of a communication equilibrium of the one-shot game. Vida (2007) and Vida and Forges (2013) establish that essentially all com-

³⁰Forges (1988) proposes a characterization of correlated equilibria and communication equilibria in the infinitely repeated game proposed by Aumann, Maschler and Stearns (1968) and studied by Hart (1985), in which player 1 must choose an action at every stage.

³¹CORE Discussion Paper 8718.

munication equilibrium payoffs of the one-shot game can be implemented in this way, namely, as correlated equilibrium payoffs of the long cheap talk game. They also show that the number of stages of cheap talk at equilibrium is finite, but not necessarily uniformly bounded.

Remark: correlated equilibrium in games with incomplete information

As recalled above, Aumann (1974) proposed the correlated equilibrium as a solution concept for games in strategic form. In the previous section, we applied Aumann (1974)'s notion to the strategic form of sender-receiver games based on $DP_0(p)$. Myerson (1982) considers the problem of an uninformed principal who can commit himself but cannot monitor the decisions of informed agents. The principal maximizes his utility function but acts otherwise as a communication device between the agents. Myerson (1982, 1991) views the communication equilibrium as the natural extension of Aumann (1974)'s correlated equilibrium to games with incomplete information.

Forges (1993) points out that the strategic form correlated equilibrium is a legitimate definition of correlated equilibrium as well. The results described above show that this definition makes sense, in particular, in cheap talk games. One could argue that implementing communication equilibrium by correlated equilibrium amounts to replacing a mediator by another one. Nevertheless, the informed player does not have to report his type to the mediator who runs a correlated equilibrium. Hence the latter mediator is appropriate if the players have privacy concerns.

Forges (1993) proposes another legitimate definition of correlated equilibrium in games with incomplete information, the "Bayesian solution". This solution concept can be defined as a communication equilibrium in which the mediator is omniscient. Let us apply the definition to $DP_0(p)$. Whether player 1 is informed or not does not matter. The omniscient mediator recommends an action to player 2, which the latter follows at equilibrium. Bayesian persuasion (see Section 2) corresponds to a Bayesian solution which maximizes player 1's ex ante expected payoff. The previous solution concepts will be further illustrated in the next two sections. Bergemann and Morris (2016a)'s Bayes correlated equilibrium is akin to the Bayesian solution but by contrast with the extensions of Aumann (1974) surveyed above, does not even maintain the assumption that the mediator cannot have access to information that the players do not collectively possess.

7 Comparison of information transmission schemes

Let us go back to the decision problem $DP_0(p)$, in which the payoff of the decision-maker (player 2), $V^k(a_2)$, depends on his action $a_2 \in A_2$ and an unknown state of nature k , distributed according to p . Player 1's prior is p as well and his payoff is $U^k(a_2)$. In this basic model, player 1 is a dummy and player 2 maximizes his expected payoff $\sum_k p^k V^k(a_2)$. Recall from Section 2 that $u_{NR}(p)$ denotes the best ex ante expected payoff that player 1 can obtain in $DP_0(p)$. We have seen three kinds of games that extend $DP_0(p)$ by making a signal available to player 2 before he chooses his action.³²

1. In a sender-receiver game (equivalent to a *cheap talk* game $CT_1(p)$), player 1 is informed of the state of nature k and sends a message s in some set S to player 2.³³ Then the latter chooses an action $a_2(s)$ that maximizes his expected payoff at the posterior probability p_s over K . Player 1 faces strong incentive compatibility conditions, namely, he must be indifferent between any two messages that he sends with positive probability (see Section 5).
2. In a *mediated* game, player 1 sends a message $s \in S$ to a reliable mediator who in turn selects, according to some probability distribution depending on s , another message to player 2. By the revelation principle, every equilibrium outcome of any mediated game can be achieved as a canonical communication equilibrium, which is described by type-dependent probability distributions $\delta(\cdot | k)$, $k \in K$, selecting a recommendation – in A_2 – to player 2. The corresponding set of equilibrium payoffs is $\mathcal{M}(DP_0(p))$ (see Section 6).
3. In the “Bayesian persuasion” game (see Section 2), everything happens as if, before being informed of the state of nature, player 1 could hire an omniscient mediator – or at least a reliable mediator who would learn the true state of nature k – to recommend an action in A_2 to player 2 using a probability distribution $\delta(\cdot | k)$, $k \in K$. In this

³²In this section, we focus on the case where player 2 receives a single message, directly from player 1 or via a mediator.

³³Different games can potentially be generated by varying the (size of the) set S of messages.

scenario, player 2’s equilibrium conditions are exactly as in a canonical communication equilibrium, but player 1 does not face any incentive compatibility constraints. Player 1’s objective is to maximize his own ex ante expected payoff (under player 2’s obedience constraint).

In the third game, player 1’s optimum is well-defined and easy to compute: it is $cavu_{NR}(p)$ (see Section 2). As observed in the previous section, $cavu_{NR}(p)$ is the best ex ante expected payoff player 1 can expect from a Bayesian solution, i.e., by relying on an omniscient mediator. We might also ask about player 1’s best ex ante expected payoff when the mediator cannot check the state of nature, or when no mediator at all is available, namely, in the other two extensions of $DP_0(p)$.

Regarding the second class of games, Salamanca (2016/2019) shows that player 1’s ex ante best communication equilibrium payoff in $DP_0(p)$ – let us denote it as $u^*(p)$ – is the value of a fictitious persuasion problem constructed on the players’ virtual utilities (see Myerson (1991)). In the fictitious problem, by construction, there are no incentive constraints so that concavification can be applied as in the third game. The expression of $u^*(p)$ is more complex than in a Bayesian persuasion problem because the value function to be concavified changes with the underlying prior belief p .

Lipnowski and Ravid (2019) consider a particular class of sender-receiver games, in which player 1 has “transparent motives”, namely, his payoff is type-independent ($U^k(a_2) = U(a_2)$, for every $a_2 \in A_2$). Under this further assumption, player 1’s incentive compatibility conditions reduce to equalities in \mathbb{R} (rather than in \mathbb{R}^k , recall condition (ii) in Section 2 and the analysis in Section 5). Lipnowski and Ravid (2019) show that the best expected payoff player 1 can expect from cheap talk in $DP_0(p)$ is $qucavu_{NR}(p)$, where $qucavu_{NR}$ denotes the smallest quasi-concave function above u .

Furthermore, they establish that player 1 cannot expect a higher expected payoff by using several (possibly infinitely many) stages of cheap talk. In other words, if player 1 has transparent motives, focusing on a single stage of cheap talk is w.l.o.g. to characterize player 1’s best ex ante expected payoff. However, Lipnowski and Ravid (2019) show that several stages of cheap talk can improve player 2’s expected payoff.

Example 3

We consider a particular decision problem $DP_0(p)$, proposed by Salamanca (2019). As in Example 1, the informed player has two types, $K = \{1, 2\}$ and the decision-maker has four actions $A_2 = \{j_1, j_2, j_3, j_4\}$. The latter's utility function is the same as in Example 1. However, and this is a main feature of the example, the preferences of the informed player do not depend on his type, as in Lipnowski and Ravid (2019).

$$\begin{aligned} (U^1, V^1)(\cdot) &= \begin{array}{cccc} j_1 & j_2 & j_3 & j_4 \\ 2, 0 & 0, 3 & 3, 4 & 1, 5 \end{array} \\ (U^2, V^2)(\cdot) &= \begin{array}{cccc} j_1 & j_2 & j_3 & j_4 \\ 2, 5 & 0, 4 & 3, 3 & 1, 0 \end{array} \end{aligned}$$


Whatever his type, player 1's most preferred decision is j_3 , his second best is j_1 , then comes j_4 and finally j_2 . Recall that p^1 denotes the probability of type 1. The optimal decisions of player 2 are, as in Example 1, j_1 if $0 \leq p^1 \leq \frac{1}{4}$, j_2 if $\frac{1}{4} \leq p^1 \leq \frac{1}{2}$, j_3 if $\frac{1}{2} \leq p^1 \leq \frac{3}{4}$ and j_4 if $\frac{3}{4} \leq p^1 \leq 1$.³⁴

The figure below, borrowed from Salamanca (2019), describes the mappings $u_{NR}(p) \leq qucavu_{NR}(p) \leq u^*(p) \leq cavu_{NR}(p)$. All the inequalities are strict at $p = \frac{2}{5}$.

Let us look more closely what happens when the prior probability of type 1 is $\frac{2}{5}$. Then, in the absence of further information, player 2 chooses j_2 , the worst decision from player 1's point of view. Player 1 gets $u_{NR}(\frac{2}{5}) = 0$, player 2 gets $\frac{18}{5} = 3.6$.

Assume first that player 1, knowing his type, can modify the prior belief of player 2 by having an informal conversation with him. Let m_1 and m_2 be two messages available to player 1. Let him send m_1 with probability 1 when his type is 1 and with probability $\frac{2}{3}$ when his type is 2. Given this strategy of player 1, player 2 believes that both types are equally likely when he receives message m_1 . He is then indifferent between j_2 and j_3 . Let him choose j_2 with probability $\frac{1}{3}$ and j_3 with probability $\frac{2}{3}$. When player 2 receives message m_2 , he believes that player 1's type is 2 for sure and decides on j_1 . If player 2 uses the previous strategy (i.e., $\tau(m_1) = (\frac{1}{3}j_2, \frac{2}{3}j_3)$, $\tau(m_2) = j_1$), then player 1 is indifferent between the two messages and can send them with the previous

³⁴The player's utility function can be given a concrete interpretation, in the vein of Forges (1990).



type-dependent probabilities. This cheap talk equilibrium improves player 1's expected payoff (he gets 2) as well as player 2's one (he gets $\frac{59}{15} \simeq 3.93$).

Let us turn to the second scenario, in which player 1 still knows his type but uses an intermediary to talk to player 2, namely, let us consider a particular communication equilibrium. If player 1 reports “type 1” to the mediator, the latter recommends j_3 (resp., j_4) with probability $\frac{3}{4}$ (resp., $\frac{1}{4}$) to player 2. If player 1 reports “type 2” to the mediator, the latter recommends j_1 (resp., j_3) with probability $\frac{1}{2}$ (resp., $\frac{1}{2}$). If player 2 follows the recommendation, player 1 cannot benefit from lying. If player 1 reports his type truthfully, player 2 believes that he faces type 1 (resp., type 2) for sure when he is recommended j_4 (resp., j_1). When j_3 is recommended, both types are equally likely. In all cases, player 2 is happy to follow the recommendation. At this communication equilibrium, player 1 gets an expected payoff of 2.5 and player 2 gets 4.1. Both players get more than in the cheap talk equilibrium above. Using his characterization, Salamanca (2019) obtains a full description of the mapping $u^*(p)$ in this example. In particular, $2.5 = u^*(\frac{2}{5})$, the previous communication equilibrium achieves player 1's best possible ex ante expected payoff.

As a consequence of Forges (1985)'s result (see Section 6), the previous communication equilibrium must be implementable as a correlated equilib-

rium of the sender-receiver game. Salamanca (2019) shows that in this example, implementation turns out to be extremely simple. Only two messages, say ℓ and r , are necessary. The correlation device recommends two possible strategies $\sigma : \{1, 2\} \rightarrow \{\ell, r\}$ with positive probability to player 1: $r\ell$ (type 1 sends r , type 2 sends ℓ) and rr (both types send r). Two strategies $\tau : \{\ell, r\} \rightarrow \{j_1, \dots, j_4\}$ are also recommended with positive probability to player 2: j_1j_3 ($\tau(\ell) = j_1, \tau(r) = j_3$) and j_1j_4 ($\tau(\ell) = j_1, \tau(r) = j_4$). The correlation device selects the pair of strategies as follows:

	j_1j_3	j_1j_4
$r\ell$	$\frac{1}{4}$	$\frac{1}{4}$
rr	$\frac{1}{2}$	0

Player 1's equilibrium conditions are straightforward. For player 2, one checks that:

$$\begin{aligned} \Pr(\text{type 1} \mid j_1j_3, \sigma(\tilde{k}) = \ell) &= 0, \text{ hence, } \tau(\ell) = j_1 \text{ is optimal.} \\ \Pr(\text{type 1} \mid j_1j_3, \sigma(\tilde{k}) = r) &= \frac{1}{2}, \text{ hence, } \tau(r) = j_3 \text{ is optimal.} \\ \Pr(\text{type 1} \mid j_1j_4, \sigma(\tilde{k}) = \ell) &= 0, \text{ hence, } \tau(\ell) = j_1 \text{ is optimal.} \\ \Pr(\text{type 1} \mid j_1j_4, \sigma(\tilde{k}) = r) &= 1, \text{ hence, } \tau(r) = j_4 \text{ is optimal.} \end{aligned}$$

Let us come to the third scenario. Player 1 does not know his type but can make it check by some reliable test or can rely on an omniscient mediator. Assume that the test can have two possible results, r_1 or r_2 , that the result is necessarily r_1 for type 1 and r_1 with probability $\frac{2}{3}$, r_2 with probability $\frac{1}{3}$ for type 2 (hence the test is highly biased). When player 2 observes r_2 , he is sure to face type 2 so that j_1 is an optimal decision. When player 2 observes r_1 , he concludes that both types are equally likely. Hence player 2's posteriors are exactly as in the cheap talk equilibrium above. The difference is that, now, player 1 does not have to fulfill any incentive constraint. Hence we can have player 2 deciding on j_3 (player 1's first best) when his posterior belief is $\frac{1}{2}$. Player 1's expected payoff is now $2.8 = \text{cavu}_{NR}(\frac{2}{5})$. Player 2 only gets $\frac{19}{5} = 3.8$.

Let us sum up what this section has taught us. In games of pure information transmission (i.e., games based on $DP_0(p)$), a ‘‘concavification approach’’

applies to compute the informed player’s best ex ante expected communication equilibrium payoff $u^*(p)$. Not surprisingly, $u^*(p)$ lies below the regular concavification of $u_{NR}(p)$, the informed player’s best ex ante expected payoff at a Bayesian solution, i.e., under Bayesian persuasion. If in addition, the sender’s motives are transparent, then a “concavification approach” applies also to compute the informed player’s best ex ante expected equilibrium payoff under cheap talk, which turns out to be the quasi-concavification of $u_{NR}(p)$.

8 Cheap talk and mediation in Crawford and Sobel’s leading example

In the information transmission framework adopted up to now, in which finitely many types and actions are feasible, various examples have been proposed to illustrate that nonrevealing equilibrium payoffs can be improved by a single stage of cheap talk, that long cheap talk can do better than one stage signaling and that mediated talk can do even better than long cheap talk. The papers briefly described below compare the effects of different forms of information transmission within a single parametrized example, the “uniform quadratic” case of Crawford and Sobel (1982). The solution concepts reviewed above apply to this example, even if its setup does not satisfy our finiteness assumptions.

In the “uniform quadratic” decision problem, the informed player’s type k is uniformly distributed over the unit interval $[0, 1]$ and the uninformed player’s actions also belong to $A_2 = [0, 1]$. The utility functions are $U^k(a_2) = -[a_2 - (k + b)]^2$ for the informed player and $V^k(a_2) = -[a_2 - k]^2$ for the uninformed one. These utility functions, which are parametrized by a bias $b > 0$, capture the fact that for every type k , player 1’s uniquely defined ideal decision $k + b$ differs from player 2’s ideal decision by the bias b . Conceptually, this decision problem is similar to the model $DP_0(p)$ considered in Sections 2 and 4 (formally, the main difference is that types and actions take values in a continuum). As we will see below, the comparison of the equilibrium outcomes associated with different information transmission schemes, as the bias b varies, is enlightening.

Crawford and Sobel (1982) consider the plain sender-receiver game, in which the informed player sends a single message (in $[0, 1]$) to the uninformed

one who then makes a decision. They show that all equilibria of the sender-receiver game are equivalent to partition equilibria and that no revelation of information can occur if $b > \frac{1}{4}$. The form of the utility functions implies that, when the decision-maker chooses his action optimally after having updated his beliefs, the players' ex ante payoffs only differ by a constant. As a consequence, the ex ante equilibrium expected payoffs of the sender-receiver game can be Pareto-ranked.³⁵ The players' ex ante equilibrium expected payoffs increase as the partition gets finer.³⁶

Krishna and Morgan (2004) extend the sender-receiver game into a cheap talk game with several stages, which is similar to the game $CT_n(p)$ introduced in Section 5. They describe explicit, plausible, "conversations" in which a second stage of information transmission depends on the outcome of a jointly controlled lottery. As long as $b \leq \frac{1}{\sqrt{8}}$ ($\simeq 0.35$), such conversations enable the players to Pareto-improve on the equilibrium payoffs of the plain sender-receiver game. To be more precise, let us say that an equilibrium is monotonic if the receiver's decision increases with the sender's type. Krishna and Morgan (2004) identify two different classes of equilibria: monotonic ones that yield Pareto superior payoffs for $b \leq \frac{1}{8}$ and nonmonotonic ones with the same property for $\frac{1}{8} \leq b \leq \frac{1}{\sqrt{8}}$. In particular, for $\frac{1}{4} \leq b \leq \frac{1}{\sqrt{8}}$, bilateral cheap talk makes the revelation of some information possible.

Goltsman, Hörner, Pavlov, and Squintani (2009) go on by considering the mediated game, namely, the communication equilibrium payoffs (the analog of the set $\mathcal{M}(DP_0(p))$ of Section 6, see also scenario 2 in Section 7). They show that no useful information can be transmitted at a communication equilibrium if $b \geq \frac{1}{2}$, so that we may henceforth assume that $b < \frac{1}{2}$. Goltsman et al. (2009) identify the optimal ex ante expected communication equilibrium payoff and show that it can be achieved by a particular mediation procedure proposed earlier by Blume, Board and Kawamura (2007). Goltsman et al. (2009) also establish that equilibria achieved with bounded cheap talk do as well as communication equilibria if and only if $b \leq \frac{1}{8}$ and that, in this case, Krishna and Morgan (2004)'s cheap talk equilibrium achieves the best possible communication equilibrium payoff. Using the terminology of Section 6, when the conflict of interest is sufficiently low, the players can implement a mediator by cheap talk in a straightforward way.

³⁵The same property holds in all the extensions of the game that are considered below.

³⁶For the informed player, the interim expected payoff is the relevant one. But here, by contrast with Forges (1990), no ranking is available in terms of these.

Blume (2012) proves that the optimal communication equilibrium payoff identified by Goltsman et al. (2009) can be implemented as a correlated equilibrium payoff of Crawford and Sobel (1982)’s sender-receiver game. Conceptually, this result is similar to Forges (1985) – see Section 6 – applied to a specific communication equilibrium outcome. But the continuum of types calls for a different argument.

So far, cheap talk in the uniform quadratic decision problem has been limited to a few stages. Chen, Goltsman, Hörner and Pavlov (2017) consider a particular form of cheap talk, “straight talk”, which can last for an arbitrary number of stages. In a straight talk equilibrium, only two messages are used, one of which leads to the end of the conversation. Chen et al. (2017) observe that the nonmonotonic equilibrium identified by Krishna and Morgan (2004) for $\frac{1}{8} \leq b \leq \frac{1}{\sqrt{8}}$ is a special case of straight talk (i.e., ending “shortly”) while the unbounded cheap talk in Forges (1990b)’s example is not straight.

Chen et al. (2017) show, among other results, that if $\frac{1}{8} \leq b \leq \frac{1}{\sqrt{8}}$, then, for every number of stages n , the best straight talk equilibrium of length $n+1$ (within a particular class) achieves a higher expected payoff than the best equilibrium straight talk equilibrium of length n (within the same particular class), with no a priori bound on the duration of this conversation. In other words, for appropriate values of b , bilateral, possibly long cheap talk enables the players to improve on the expected equilibrium payoffs they achieve in the sender-receiver game. This suggests that examples like the ones of Forges (1984, 1990b) are not pathological.

9 Existence of a joint plan equilibrium

In Section 6, we established strong connections between the infinitely repeated game $\Gamma_\infty(p)$ and static games of information transmission. Since then, we focused on the latter framework. Let us come back to $\Gamma_\infty(p)$ and investigate other relationships between $\Gamma_\infty(p)$ and static decision problems.

Recall, from Section 3, that the question of the existence of an equilibrium in $\Gamma_\infty(p)$ was not solved by Aumann, Maschler and Stearns (1968). Note also that the question is meaningless in the infinitely repeated games of pure information transmission considered in Sections 4 and 6, in which a nonrevealing equilibrium always exists.

Sorin (1983) established the existence of a joint plan equilibrium (as defined in Section 3) in $\Gamma_\infty(p)$ under the further assumption that the informed

player has two types only. The general problem remained open for a while until Simon et al. (1995) managed to extend Sorin (1983)’s proof to the case of an arbitrary number of types. Simon et al. (1995) present their result as one of the “Borsuk-Ulam type”. A very rough intuition for this interpretation comes from the informed player’s incentive compatibility conditions (condition (ii) in Section 2), which take the form of equalities of vector payoffs.³⁷

It is quite remarkable that Sorin (1983) and Simon et al. (1995) establish the existence of an equilibrium in the infinitely repeated game $\Gamma_\infty(p)$, for every prior probability distribution p , without making use of Hart (1985)’s characterization of all equilibrium payoffs of the game (described in Section 2). Indeed, recalling the definition of a joint plan equilibrium, Sorin (1983) and Simon et al. (1995) prove that if, at some prior p , $\Gamma_\infty(p)$ does not have a nonrevealing equilibrium, then p can be split into posteriors p_s , $s \in S$, such that (i) for every $s \in S$, $\Gamma_\infty(p_s)$ does have a nonrevealing equilibrium *and* (ii) the splitting is incentive compatible for the informed player. In other words, there always exist equilibria in which Hart (1985)’s dimartingale converges in at most one stage.

Adopting a “design point of view”, once we know that the set of all equilibrium payoffs of $\Gamma_\infty(p)$, characterized by Hart (1985), is nonempty, we can maximize one of the players’ ex ante expected payoff over this set. Then the full power of the characterization can become useful, namely, several or even infinitely many stages of signaling may be necessary to reach such a best payoff.

Renault (2000) generalizes Simon et al. (1995) existence theorem to the case where after every stage, instead of observing each other’s actions, the players get a private, random, state-independent signal. In particular, in his Proposition 4.2, he proposes a reformulation of Simon et al. (1995)’s result that abstracts from the repeated games framework. Simon et al. (2008) suggest that this abstract formulation can be useful to obtain existence results in a class of static principal-agent problems. However the agent’s participation constraints in Simon et al. (2008)’s are not easy to interpret in a standard, yet not trivial, principal-agent framework.

Forges, Horst and Salomon (2016) and Forges and Horst (2018) start with

³⁷Chakraborty and Harbaugh (2010) make a direct use of the Borsuk-Ulam theorem to show the existence of informative equilibria in a model in which the preferences of the informed player do not depend on his type (as in Lipnowski and Ravid (2019)) but this type belongs \mathbb{R}^d , for some d .

the same static game as Aumann and Hart (2003) (see Section 5). They explore the effects of a phase of cheap talk followed by commitment on a joint strategy. More precisely, they make use of Renault (2000)'s formulation to establish – under appropriate assumptions – the existence of cooperative solutions, which are incentive compatible and “posterior individually rational” (i.e., individually rational given the transmitted information). These solutions can be implemented as Nash (or even perfect Bayesian) equilibrium outcomes of various extensions of the static game in which the informed player sends a message to the uninformed one and both players make an agreement on how to jointly choose their actions.

Renault (2001) investigates the existence of a Nash equilibrium in n person repeated games with lack of information on one side, namely, in the case where all players but one know the state of nature. If $n \geq 4$, the existence of a completely revealing equilibrium is immediate, hence Renault (2001) focuses on the case $n = 3$. He establishes that a joint plan equilibrium exists when there are two states of nature but shows that this result may fail as soon as there are three states. Existence of an equilibrium (that would not be achieved by a joint plan) is still an open question in this model.

10 More general models

The set up of the previous sections is quite restrictive : there are only two players, only one of them has access to information and in many cases, the latter individual does not have payoff relevant decisions. We shall now turn to more general models, keeping in mind the motivation of applying results from repeated games to static games with information transmission.

We have discussed above the first and the last chapter of Aumann and Maschler (1995) – Aumann and Maschler (1966) and Aumann, Maschler and Stearns (1968) – which both deal with infinitely repeated games with lack of information on one side. The chapters in between contain a number of results on zero-sum infinitely repeated games with lack of information *on both sides*.

The simplest possible game $\Gamma_\infty(p, q)$ generalizing the model of Sections 1 and 2 is described by a set of types K and a set of actions A_1 for player 1, a set of types L and a set of actions A_2 for player 2, independent priors p over K and q over L and payoff functions $U^{k\ell}$ and $V^{k\ell}$ (over $A_1 \times A_2$) for player 1 and player 2 respectively, for every state (k, ℓ) .³⁸ Nature selects a

³⁸All sets (K , L , A_1 and A_2) are assumed to be finite.

pair of types (k, ℓ) at a preliminary stage, every player is only informed of his own type. Then at every subsequent stage, the players simultaneously choose actions, which are then announced to both. As for $\Gamma_\infty(p)$, payoffs in $\Gamma_\infty(p, q)$ are undiscounted, i.e., evaluated as limits of averages.

Aumann and Maschler (1967, Chapter 2 of Aumann and Maschler (1995)) assume that $\Gamma_\infty(p, q)$ is zero-sum, namely, that $V^{k\ell} = -U^{k\ell}$ for every state (k, ℓ) . Defining $u_{NR}(p, q)$ as the value of the game in which both player 1 and player 2 do not reveal any information, they show that the *minmax* of $\Gamma_\infty(p, q)$ is the function $ve x_q cav_p u_{NR}(p, q)$ and that, similarly, the *maxmin* of $\Gamma_\infty(p, q)$ is the function $cav_p ve x_q u_{NR}(p, q)$. As a consequence, $\Gamma_\infty(p, q)$ may not have a value (it may happen that $cav_p ve x_q u_{NR}(p, q) < ve x_q cav_p u_{NR}(p, q)$)³⁹ and the existence of a joint plan equilibrium (recall Section 9) is no longer guaranteed if both players have private information.

Mertens and Zamir (1971-72) consider the value $v_n(p, q)$ of the n times repeated game $\Gamma_n(p, q)$ and show that $\lim_{n \rightarrow \infty} v_n(p, q)$ exists and is the unique solution of a system of functional equations.⁴⁰ Starting from a continuous mapping u over beliefs (i.e., over $\Delta(K) \times \Delta(L)$), the system is described as follows:

$$\begin{aligned} f &= ve x \max \{u, f\} \\ f &= cav \min \{u, f\}. \end{aligned}$$

Under appropriate assumptions, there is a unique solution f to the system, which will be denoted as $f = MZ(u)$. This defines the “Mertens and Zamir operator” $MZ(\cdot)$. With this notation, Mertens and Zamir (1971-72) show that in the model described above

$$\lim_{n \rightarrow \infty} v_n(p, q) = MZ(u_{NR})(p, q)$$

and

$$cav_p ve x_q u_{NR}(p, q) \leq MZ(u_{NR})(p, q) \leq ve x_q cav_p u_{NR}(p, q).$$

This system has been thoroughly studied, starting with Mertens and Zamir (1977) (see Mertens, Sorin and Zamir (2015)).

³⁹Aumann and Maschler (1967), together with R. Stearns, propose a first example (see Example 4.10 in Aumann and Maschler (1995, Chapter 2)).

⁴⁰The same holds for $\lim_{\delta \rightarrow 1} v_\delta(p, q)$, where $v_\delta(p, q)$ is the value of the δ -discounted game $\Gamma_\delta(p, q)$.

Among the various models related to the previous system of functional equations, “splitting games” (Laraki (2002), Sorin (2002), Oliu-Barton (2017), Laraki and Renault (2018)) are particularly relevant to the present survey. The basic version of a splitting game is a zero-sum stochastic game in which the states are probability distributions (p, q) over $K \times L$. At every stage t , as in the models of Section 2, player 1 chooses a “splitting” of the current probability distribution p_t over K namely, a probability distribution over $\Delta(K)$ (to choose p_{t+1}), with expectation p_t . Similarly player 2 chooses a “splitting” of q_t , in $\Delta(\Delta(L))$. In other words, each player controls a martingale. The splittings are used to select the next state and the stage payoffs depend continuously on the state.

Splitting games first appeared in the analysis of infinitely repeated games with incomplete information with lack of information on both sides, building once again on the martingale property of posterior beliefs pointed out in Sections 2 and 3. Splitting games were then studied for themselves, which involved extensions of the previous model and lead to new results for systems of functional equations of the form initially proposed by Mertens and Zamir.

Koessler, Laclau, Renault and Tomala (2019) apply the previous methodology to a genuine information design problem. As above, two strictly competing players interact over stages; player 1 controls p_t in $\Delta(K)$ and player 2 controls q_t in $\Delta(L)$. But p_t and q_t now represent the successive beliefs of a third individual, who has to make a single decision, at the end of the persuasion process. For instance, player 1 and player 2 are competing firms, the decision-maker is a consumer who makes an optimal decision given his final beliefs over the quality of the products. Koessler et al. (2019) consider various possible games, depending on whether the firms move – i.e., choose a “splitting” – simultaneously or not, and whether they face a deadline or not.

Let $u_{NR}(p, q)$ be player 1’s induced payoff, given final beliefs (p, q) of the decision-maker. Koessler et al. (2019) assume that the correspondence of splittings available to player 1 and player 2 is well-behaved. They show that the various possible games have a (Markovian) equilibrium. In particular, if the players move one after the other for $T \geq 2$ stages, the value is $ve x_q cav_p u_{NR}(p, q)$ if player 1 moves last (resp., $cav_p ve x_q u_{NR}(p, q)$) if player 2 moves last).⁴¹ If the players can split for infinitely many stages, then the value is $MZ(u_{NR})(p, q)$. Koessler et al. (2019) show on an example that the

⁴¹If $T = 1$ and say, player 1 is the only one to move, we recover $cav_p u_{NR}(p, q)$ as in Section 2.

number of disclosure periods needed to reach the value may be unbounded, although disclosure stops in finite time with probability one. The pattern may remind us of the long cheap talk equilibria of Section 5.

The previous paper illustrates particularly well how the methodology of repeated games with incomplete information can be used in information design. A promising literature on multistage persuasion is also developing independently of this methodology (see, e.g., Best and Quigley (2017), Makris and Renou (2019)). There are many more papers to mention on information design with multiple information designers and possibly, multiple agents to convince but they focus on one-shot information design (see, e.g., Koessler, Laclau and Tomala (2018), Gentzkow and Kamenica (2017), Mathevet, Prego and Taneva (2019)).

Bergemann and Morris (2016b, 2019) propose a general setting in which finitely many players have to choose an action. The players’ utility is determined by their actions and an unknown state of nature. The outcomes that are achievable (under a given solution concept, e.g., Nash equilibrium⁴²) in this basic environment depend on the underlying information structure. Under a literal interpretation, the latter is chosen by an information designer. But, as Bergemann and Morris (2016b, 2019) explain, a metaphorical interpretation can also be adopted, in which it is the analyst who characterizes the effect of different information structures. Whatever the interpretation, a two steps approach is in order. The first one is to characterize the set of outcomes that are feasible given the information structure (and the underlying solution concept). The second one is to identify the best outcome according to some criterion, e.g., maximizing the designer’s expected utility. This method is illustrated, on a very particular case, in Section 7.

We have indicated that the study of zero-sum repeated games with lack of information on both sides sheds light on Bayesian persuasion problems. Let us turn to two-person *nonzero-sum* repeated games with lack of information on both sides. Very few results are available on this topic, but some of them could be useful to analyze long cheap talk in models extending the ones of Section 5.

Koren (1988/1992) considers infinitely repeated games $\Gamma_\infty(p, q)$ with “private values” or “known own payoffs”, namely, in which $U^{k\ell}(a) = U^k(a)$ and

⁴²All along this survey, we have focused on equilibrium concepts that can be defined from Nash’s fundamental notion. Other solution concepts, like rationalizable strategies, are obviously conceivable.

$V^{k\ell}(a) = V^\ell(a)$ for every pair of actions $a \in A_1 \times A_2$ and every state (k, ℓ) . He shows that, in this case, Nash equilibria are payoff-equivalent to completely revealing equilibria.⁴³ Using this simple characterization, he constructs an example in which, as in the zero-sum case above, there is no equilibrium at all. Both the characterization and the example make use of the – *long term* – individual rationality levels (i.e., minmax) in $\Gamma_\infty(p, q)$, which are computed in the associated zero-sum games.⁴⁴ Salomon and Forges (2015) make the further assumption that long term and short term individual rationality levels coincide (they refer to “uniform punishment strategies”). Under this assumption, the Nash equilibria of $\Gamma_\infty(p, q)$ can be characterized in terms of the one-shot game but the non-existence phenomenon persists.⁴⁵

Amitai (1996b) substantially extends Koren (1988/1992)’s work by characterizing the equilibria of the infinitely repeated game $\Gamma_\infty(p, q)$ without making restrictive assumptions on the payoff functions. In other words, Amitai (1996b) extends Hart (1985) to the case of lack of information on both sides. A companion paper, Amitai (1996a) undertakes a similar task for long cheap talk, i.e., seeks to generalize the results of (a preliminary version of) Aumann and Hart (2003). Summing up very loosely, in both models, equilibria generate a dimartingale, as in the case of lack of information on one side (see Section 3 and 5); the converse holds in static games with cheap talk and in the infinitely repeated game, but, in the latter case, Amitai (1996b) makes an additional assumption, referred to as “tightness”. As the existence of uniform punishment strategies, “tightness” makes long term punishments unnecessary.

⁴³The same holds for the communication equilibria of $\Gamma_\infty(p, q)$ (see Forges (1992)).

⁴⁴Under the known own payoff assumption, a vector payoff $x \in \mathbb{R}^K$ is individually rational for player 1 in $\Gamma_\infty(p, q)$ if and only if, for every $\rho \in \Delta(K)$, $\rho \cdot x \geq \text{cav}u_{NR}(\rho)$ with u_{NR} defined as in the previous sections. A symmetric definition obviously applies for player 2.

⁴⁵Salomon and Forges (2015) consider δ -discounted infinitely repeated games with lack of information on both sides (see also Peski (2008, 2014) on this topic). They investigate the possible extension, in the nonzero-sum case, of Mertens and Zamir (1971-72)’s convergence of v_δ as $\delta \rightarrow 1$.

References

- [1] Amitai, M., 1996a, Cheap talk with incomplete information on both sides, preprint (Hebrew University, Jerusalem).
- [2] Amitai, M., 1996b, Repeated games with incomplete information on both sides, preprint (Hebrew University, Jerusalem).
- [3] Aumann, R., 1974, Subjectivity and correlation in randomized strategies, *Journal of Mathematical Economics* 1, 67–95.
- [4] Aumann, R. and S. Hart, 1986, Bi-convexity and bi-martingales, *Israel Journal of Mathematics* 54, 159-180.
- [5] Aumann, R. and S. Hart, 2003. Long cheap talk, *Econometrica* 71, 1619–1660 (early version: Polite talk isn't cheap, 1993, Hebrew University, Jerusalem).
- [6] Aumann, R. and M. Maschler, 1966, Game-theoretic aspects of gradual disarmament, Reports to the U.S. Arms Control and Disarmament Agency, Princeton: Mathematica ST-80, Ch. V, 1–55
- [7] Aumann, R. and M. Maschler, 1967, Repeated games with incomplete information: A survey of recent results, Reports to the U.S. Arms Control and Disarmament Agency, Princeton: Mathematica ST-116, Ch. III, pp. 287-403.
- [8] Aumann, R., M. Maschler and R. Stearns, 1968, Repeated games of incomplete information: an approach to the non-zero-sum case, Reports to the U.S. Arms Control and Disarmament Agency, Princeton: Mathematica ST-143, Ch. IV, 117–216.
- [9] Aumann, R. and M. Maschler, 1995, Repeated games with incomplete information, Cambridge, MA: MIT Press.
- [10] Bárány, I., 1992, Fair distribution protocols or how the players replace fortune, *Mathematics of Operations Research* 17, 327–340.
- [11] Ben-Porath, E. , 2003, Cheap talk in games with incomplete information, *Journal of Economic Theory* 108, 45–71.

- [12] Bergemann, D., and S. Morris, 2016a, Bayes correlated equilibrium and the comparison of information structures in games, *Theoretical Economics* 11, 487–522.
- [13] Bergemann, D., and S. Morris, 2016b, Information design and Bayesian persuasion, *American Economic Review: Papers & Proceedings* , 106, 586–591.
- [14] Bergemann, D., and S. Morris, 2019, Information design: A unified perspective, *Journal of Economic Literature* 2019, 57, 44–95.
- [15] Best, J. and D. Quigley, 2017, Persuasion in the long run, preprint (University of Oxford).
- [16] Blackwell, D. , 1956, An analog of the minmax theorem for vector payoffs, *Pacific Journal of Mathematics* 6, 1–8.
- [17] Blume, A., 2002, A class of strategy-correlated equilibria in sender–receiver games, *Games and Economic Behavior* 75, 510-517.
- [18] Blume, A., O. Board and K. Kawamura, 2007, Noisy talk, *Theoretical Economics*, 2, 395–440.
- [19] Celik, G. and M. Peters, 2016, Reciprocal relationships and mechanism design, *Canadian Journal of Economics* 49, 374-411.
- [20] Chakraborty, A., and R. Harbaugh, 2010, Persuasion by cheap talk, *American Economic Review* 100, 2361–82.
- [21] Chen, Y., M. Goltsman, J. Hörner and G. Pavlov, 2017, Straight talk, preprint, Yale University.
- [22] Crawford, V. and J. Sobel, 1982, Strategic information transmission, *Econometrica*, 50, 1431–51.
- [23] Forges, F., 1984, Note on Nash equilibria in infinitely repeated games with incomplete information, *International Journal of Game Theory* 13, 179–187.
- [24] Forges, F., 1985, Correlated equilibria in a class of repeated games with incomplete information, *International Journal of Game Theory* 14, 129–150.

- [25] Forges, F., 1986, An approach to communication equilibria, *Econometrica* 54, 1375–1385.
- [26] Forges, F., 1988, Communication equilibria in repeated games with incomplete information, *Mathematics of Operations Research*, 13, 191–231.
- [27] Forges, F., 1990a, Universal mechanisms, *Econometrica* 58, 1341–1364.
- [28] Forges, F., 1990b, Equilibria with communication in a job market example, *Quarterly Journal of Economics*, 105, 375–398.
- [29] Forges, F., 1992, Repeated games of incomplete information : non-zero-sum, in : R. Aumann and S. Hart, *Handbook of Game Theory with Economic Applications*, vol. I, Elsevier Science Publishers, North-Holland, chapter 6, 155–177.
- [30] Forges, F., 1993, Five legitimate definitions of correlated equilibrium in games with incomplete information, *Theory and Decision* 35, 277–310.
- [31] Forges, F., 1994, Non-zero-sum repeated games and information transmission, in : N. Meggido, *Essays in Game Theory in Honor of Michael Maschler*, Springer-Verlag, chapter 6, 65-95.
- [32] Forges, F., 2009, Correlated equilibrium and communication in games, R. Meyers (ed.), *Encyclopedia of Complexity and Systems Science*, Springer, 1587–1596.
- [33] Forges, F., 2010, Communication in Bayesian games: overview of work on implementing mediators in game theory, preprint (Université Paris–Dauphine).
- [34] Forges, F. and U. Horst, 2018, Sender-receiver games with cooperation, *Journal of Mathematical Economics* 76, 52–61.
- [35] Forges, F., U. Horst, U. and A. Salomon, 2016, Feasibility and individual rationality in two-person Bayesian games, *International Journal of Game Theory* 45, 11–36.
- [36] Forges, F. and F. Koessler, 2008, Long persuasion games, *Journal of Economic Theory* 143, 1–35.

- [37] Gentzkow, M. and E. Kamenica, 2017, Competition in persuasion, *The Review of Economic Studies* 84, 300–322.
- [38] Gerardi, D., 2004, Unmediated communication in games with complete and incomplete information, *Journal of Economic Theory* 114, 104–31.
- [39] Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani, 2009, Mediation, arbitration and negotiation, *Journal of Economic Theory*, 144, 1397–420.
- [40] Halpern, J., 2008, Computer science and game theory: A brief survey, *Palgrave Dictionary of Economics* (S. N. Durlauf and L. E. Blume, eds.), Palgrave MacMillan.
- [41] Harsanyi, J., 1967, Games with incomplete information played by Bayesian players, *Management Science* 14, 159–182.
- [42] Hart, S., 1985, Nonzero-sum two-person repeated games with incomplete information. *Mathematics of Operations Research* 10, 117–153.
- [43] Kamenica, E., 2019, Bayesian persuasion and information design, *Annual Review of Economics* 11, 249–272.
- [44] Kamenica, E. and M. Gentzkow, 2011, Bayesian persuasion, *American Economic Review*, 101, 2590–615.
- [45] Koessler, F., M. Laclau, T. Tomala, 2018, Interactive information design, preprint (Paris School of Economics).
- [46] Koessler, F., M. Laclau, R. Renault and T. Tomala, 2019, Long (and strictly competitive) information design, preprint (Paris School of Economics).
- [47] Koren, G., 1992, Two-person repeated games where players know their own payoffs, preprint, Courant Institute of Mathematical Sciences, New York University (M.Sc. thesis, Tel-Aviv University, 1988).
- [48] Kreps, D. and J. Sobel, 1994, Signalling, in : R. Aumann and S. Hart, *Handbook of Game Theory with Economic Applications*, vol II, Elsevier Science Publishers, North-Holland, 849–867.

- [49] Krishna, V., and J. Morgan, 2004, The art of conversation: eliciting information from experts through multi-stage communication,” *Journal of Economic Theory* 117,147–79.
- [50] Laraki, R., 2002, The splitting game and applications, *International Journal of Game Theory* 30, 359-376.
- [51] Laraki, R and J. Renault, 2018, Acyclic gambling games, preprint (University of Toulouse), *Mathematics of Operations Research*, forthcoming.
- [52] Lipnowski, E. and D. Ravid, 2019, Cheap talk with transparent motives, preprint (University of Chicago).
- [53] Makris, and L. Renou, 2019, Information design in multistage games, preprint (Queen Mary University, London).
- [54] Mathevet, L., J. Perego, and I. Taneva, 2019, On Information Design in Games, *Journal of Political Economy*, forthcoming.
- [55] Matthews, S. and A. Postlewaite, Preplay communication in two-person sealed bid auctions, *Journal of Economic Theory* 48, 238–263.
- [56] Mertens, J.-F. and S. Zamir, 1971-72, The value of two-person zero-sum repeated games with lack of information on both sides, *International Journal of Game Theory* 1, 39–64.
- [57] Mertens, J.-F. and S. Zamir, 1977, A duality theorem on a pair of simultaneous functional equations, *Journal of Mathematical Analysis and Applications* 60, 550–558.
- [58] Mertens J.-F., S. Sorin and S. Zamir, 2015, *Repeated Games*, Cambridge University Press.
- [59] Myerson, R., 1982, Optimal coordination mechanisms in generalized principal agent problems, *Journal of Mathematical Economics* 10, 67–81.
- [60] Myerson, R., 1986, Multistage games with communication, *Econometrica* 54, 323-358.
- [61] Myerson, R., 1991, *Game Theory: Analysis of Conflict*, Harvard University Press.

- [62] Oliu Barton M., 2017, The splitting game: uniform value and optimal strategies, *Dynamic Games and Applications*.
- [63] Özdoğan, A., 2016, A survey of strategic communication and persuasion, *Boğaziçi Journal Review of Social, Economic and Administrative Studies* 30, 1–21.
- [64] Peski, M., 2008, Repeated games with incomplete information on one side, *Theoretical Economics* 3, 29–84.
- [65] Peski, M., 2014, Repeated games with incomplete information and discounting, *Theoretical Economics* 9, 651–694.
- [66] Renault, J., 2000, On two-player repeated games with lack of information on one side and state-dependent signalling, *Mathematics of Operations Research* 25, 552–572.
- [67] Renault, J., 2001, 3-player repeated games with lack of information on one side, *International Journal of Game Theory* 30, 221–246.
- [68] Salamanca, A., 2016 (revised, 2019), The value of mediated communication, preprint (University of Southern Denmark).
- [69] Salomon, A. and F. Forges, 2015, Bayesian repeated games and reputation, *Journal of Economic Theory* 159, 70–104.
- [70] Simon, R., 2002. Separation of joint plan equilibrium payoffs from the min–max functions. *Games and Economic Behavior* 41, 79–102.
- [71] Simon, R., Spież, S., Toruńczyk, H., 1995. The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk-Ulam type. *Israel Journal of Mathematics* 92, 1–21.
- [72] Simon, R., Spież, S., Toruńczyk, H., 2008. Equilibria in a class of games and topological results implying their existence. *RACSAM* 102, 161–179.
- [73] Sobel, J., 2013. Giving and receiving advice. *Advances in Economics and Econometrics* 1, 305–341.
- [74] Sorin, S., 1983, Some results on the existence of Nash equilibria for non-zero-sum games with incomplete information, *International Journal of Game Theory* 12, 193–205.

- [75] Sorin, S., 2002, A first course on zero-sum repeated games: Springer.
- [76] Vida, P. , 2007, Long Pre-play Communication in Games. Ph.D. thesis, Universitat Autònoma de Barcelona.
- [77] Vida, P. and F. Forges, 2013, Implementation of communication equilibria by correlated cheap talk: the two-player case, Theoretical Economics 8, 95 - 123.