LONG CHEAP TALK[1]

BY ROBERT J. AUMANN AND SERGIU HART

With cheap talk, more can be achieved by long conversations than by a single message—even when one side is strictly better informed than the other. ("Cheap talk" means plain conversation—unmediated, nonbinding, and payoff-irrelevant.) This work characterizes the equilibrium payoffs for all two-person games in which one side is better informed than the other and cheap talk is permitted.

KEYWORDS: Cheap talk, communication, long conversation, incomplete information, game theory, signalling, joint lottery, dimartingale, di-span.

1. INTRODUCTION

STRATEGIC INFORMATION TRANSMISSION has been studied in economic theory for over a quarter of a century. Most formal models in this area allow for at most one message from each player. Yet in practice, as in negotiation or bargaining, protracted exchanges of messages are often observed. Does this make economic sense? Can a long exchange convey substantive information that cannot be conveyed by a single message?

Here we examine this question in the context of "cheap talk." The answer is "yes": Long cheap talk may lead to outcomes preferred by all players to those achievable with single messages. We will characterize all the equilibrium outcomes to which it can lead, in any two-person game in which one player is initially better informed than the other.

Cheap talk is just that: cheap—neither costly nor binding; and talk—not some roundabout form of communication, like mediation. Unlike "signalling,"[2] cheap talk—plain conversation—is "payoff-irrelevant;" there is no "credibility cost." The players don't strike, don't get educated, and don't issue guarantees; they simply talk. They may or may not tell the truth, and may or may not believe each other. To be sure, that there is no credibility *cost* does not mean that there is no credibility; depending on the circumstances, the cheap talk may itself create positive motivation for the players to believe each other (Examples 2.1 and 2.2 below). It is this motivation that forms the crux of our analysis.

In a sense, cheap talk is communication in its purest and simplest form: purest in that there is no direct impact on the payoff, and simplest in that there is no intermediary.

The literature on cheap talk addresses two issues, almost diametrically opposite. One is how cheap talk *restricts* the set of outcomes—equilibrium *refinement*; see Section 10. The other is how cheap talk *expands* the set of outcomes.

[2]As in Spence (1974) and Zahavi (1975); for a survey, see Kreps and Sobel (1994).

One way, of course, is by revealing substantive information; see Example 2.1. Another is by agreeing on a compromise; for example, in the Battle of the Sexes (Example 2.3) the players can, with cheap talk, reach the compromise payoff $(4, 4)$, which is not feasible without communication. The current study concerns this second issue—what equilibria are *added* by cheap talk.

In the model of this paper, Nature chooses one of several two-person games in strategic form (bimatrices), using a commonly known probability distribution; the chosen one is called the *true* bimatrix. The row player, Rowena, is informed of the true bimatrix, but the column player, Colin, is not (this is the *information phase*). The players then talk to each other for as long as they wish (the *talk phase*). When they finish talking, each one takes a single action (i.e., chooses a row or column) in the true bimatrix; this is the *action phase*. The players receive the payoffs that the true bimatrix prescribes for their actions. The entire process described in this paragraph is called the *cheap talk game.*

Informationally, this model is the simplest possible: Only Colin has anything substantive to learn.[3]

We characterize all Bayesian Nash equilibrium payoffs in all such games. The idea is to think of the talk phase as a series of revelations by Rowena, interspersed with agreements between Rowena and Colin. The agreements and revelations in the talk phase, and the choices in the action phase, must be optimal given one another; and it must be optimal for the players to keep the agreements and to reveal truthfully. Examples—some of which require long cheap talk—are provided in Section 2.

The revelations and agreements *must* come in the order prescribed, and there may be arbitrarily many of them. There are no shortcuts: The first revelation motivates an agreement, this motivates another revelation, this motivates another agreement, and so on. Any attempt to combine several revelations, or several agreements—to reveal too much too soon, or to agree to too much too soon—destroys the equilibrium. Just possibly, there is a lesson here for conducting negotiations.

In proving our result, we make use of the geometric concepts of *diconvexity* and *dimartingale* (Aumann and Hart (1986)), concepts that also provide a tool for actually computing the equilibrium payoffs.

The paper is organized as follows: Section 2 illustrates the phenomenon under study; inter alia, it presents games requiring one, two, three, and an infinity of cheap talk "stages," each stage being either an agreement or a revelation. The formal model is presented in Section 3 and discussed in Section 4; though largely verbal, this presentation is entirely precise. Section 5 first formulates the "behavioral" result to which we alluded above—the characterization

---

[3]"Substantive" means that the information is about parameters that are not under the players' control—like the payoff function—as distinguished from information about the players' intentions.

of cheap talk equilibria in terms of alternations of agreements and revelations (Theorem A)—then outlines an informal "demonstration" of this result. Section 6 explains the notion of "dimartingale," the chief ingredient in the geometric characterization of the equilibrium payoffs, which enables also an "explicit" calculation of the equilibrium payoffs. Section 7 formulates this geometric characterization (Theorem B) and then "demonstrates" it informally. As will be explained in Section 3, in cheap talk games the formal notion of "strategy" is not quite straightforward; Section 8 provides the required conceptual and technical clarifications. Section 9 discusses extensions and variations. The relevant literature on cheap talk and related issues, which is quite large, is reviewed in Section 10. Finally, the Appendix provides formal proofs.

2. SOME EXAMPLES

EXAMPLE 2.1—*Credible Signalling*. Rowena and Colin are playing one of two bimatrix games, T or B (see Figure 1), with probability $\frac{1}{2}$ each; Rowena knows which, Colin does not. Colin must choose left (L) or right (R); Rowena has no choice. Before choosing, Colin may talk with Rowena, who may, if she wants, tell him the true game; but she need not tell the truth, and Colin has no direct way of verifying her statements.

Here it is clear that Rowena should tell Colin the truth, and that Colin should believe her. To express this formally, think of "Nature" choosing T or B with probability $\frac{1}{2}$ each, and informing Rowena—but not Colin—of the game chosen. We distinguish two games. In the *silent* game, Rowena and Colin may not talk after Rowena gets her private information and before Colin acts; in the *cheap talk* game, they may. The only equilibrium payoff (in fact, the only payoff) in the silent game is then $(2, 2)$, whereas in the talking game, there is an equilibrium with payoff $(4, 4)$—precisely that described above.

Such situations are quite familiar; for example, Colin may take Rowena to dinner, and ask her to recommend a restaurant.

EXAMPLE 2.2—*Signalling that is Not Credible*. The framework is as in the previous example, but T and B are as in Figure 2. Here Rowena would like Colin to play R no matter what the true game is. So Rowena would like Colin to believe that the true game is B, which would motivate him to play R. So she

$$
T \quad
\begin{array}{c}
\begin{array}{cc} L & R \end{array} \\
\begin{array}{|c|c|} \hline 4,4 & 0,0 \\ \hline \end{array}
\end{array}
\quad \tfrac{1}{2}
$$

$$
B \quad
\begin{array}{c}
\begin{array}{cc} L & R \end{array} \\
\begin{array}{|c|c|} \hline 0,0 & 4,4 \\ \hline \end{array}
\end{array}
\quad \tfrac{1}{2}
$$

FIGURE 1.—Credible signalling.

$$\begin{array}{cc} & L \quad\ R \\ T & \boxed{\begin{array}{c|c} 4,6 & 6,0 \end{array}}\ \ \tfrac{1}{2} \end{array}$$

$$\begin{array}{cc} & L \quad\ R \\ B & \boxed{\begin{array}{c|c} 0,0 & 4,4 \end{array}}\ \ \tfrac{1}{2} \end{array}$$

FIGURE 2.—Non-credible signalling.

may tell him that the true game is *B*, no matter what it really is. But he wasn't born yesterday; it's not so easy to pull the wool over his eyes. He knows that no matter what the true game really is, Rowena is motivated to say that it is *B*. So he ignores what she says; it is as if she had remained silent. And indeed, the equilibrium payoffs in the talking game are precisely as in the silent game: the only one is $(2, 3)$.

If the true game really is *B*, it is to the advantage of *both* that she tell him the truth, and that he believe her; they would get $(4, 4)$ instead of $(0, 0)$. Nevertheless, he cannot believe her.

This situation, too, is very familiar. Think of Rowena as a job applicant and Colin as a potential employer. Rowena knows whether or not she is qualified, but wants the job in any case. So when she tells him that she is qualified, he cannot believe her, *even if she really is*. The signal, even when true, is not credible.

EXAMPLE 2.3—*Compromising*. Here the framework is slightly different: There is only one game, the familiar Battle of the Sexes (Figure 3). Rowena chooses up (*U*) or down (*D*); Colin chooses *L* or *R*. The silent game has three equilibrium payoffs: $(6, 2)$, $(2, 6)$, and $(1\tfrac{1}{2}, 1\tfrac{1}{2})$. In the talking game, the players can agree on a $\tfrac{1}{2}$–$\tfrac{1}{2}$ lottery to determine whether to play *UL* or *DR*; once the outcome of the lottery is known, they are motivated to stick to it. So $(4, 4)$, too, is an equilibrium payoff of the talking game. Note that the lottery must be *jointly controlled;* both players must be convinced that the probabilities are indeed $\tfrac{1}{2}$–$\tfrac{1}{2}$, and both must observe its outcome. (See Section 4.3.)

In this case substantive information is *not* conveyed; all the substantive information is commonly known at the beginning of play, so there is nothing to convey. The cheap talk serves a different purpose—to reach a compromise.

$$\begin{array}{c|c|c} & L & R \\ \hline U & 6,2 & 0,0 \\ \hline D & 0,0 & 2,6 \end{array}$$

FIGURE 3.—Compromising.

	LL	L	C	R	RR	
T	$0, 10$	$1, 8$	$0, 5$	$1, 0$	$0, -8$	$\frac{1}{2}$

	LL	L	C	R	RR	
B	$0, -8$	$1, 0$	$0, 5$	$1, 8$	$0, 10$	$\frac{1}{2}$

FIGURE 4.—Partial revelation.

EXAMPLE 2.4—*Credible Signalling with Partial Revelation*. Returning to the framework of Examples 2.1 and 2.2, consider the game of Figure 4. If Rowena does not reveal which game is being played, Colin's probabilities for T and B remain $\frac{1}{2}$–$\frac{1}{2}$; so he is motivated to choose C, which yields 0 to Rowena (see Figure 5, which graphs Colin's expected payoff for each of his choices as a function of the probability q that he assigns to T; the thick line denotes his best-reply payoff). If she does reveal the game being played, Colin is motivated to choose LL or RR, as the case may be, so she again gets 0. To get 1, she must motivate Colin to choose L or R; this she can do by a *partial revelation*, i.e., by hinting what the correct game is, while leaving some doubt in his mind. For example, she could use a noisy channel, which transmits the correct information with probability $\frac{3}{4}$. It then follows from Bayes' rule that after he receives the communication from Rowena, Colin's q is either $\frac{3}{4}$ or $\frac{1}{4}$, according to whether T or B is the real game. In the first case, he is motivated to choose L, in the second R (see Figure 5); either way, Rowena gets 1. Thus cheap talk can, in this case, convey substantive information; but only by *partial* revelation.

In technical terms, the silent game has only one equilibrium, whose payoff is $(0, 5)$; whereas the talking game has also an equilibrium with payoff $(1, 6)$ ($= \frac{3}{4}(1, 8) + \frac{1}{4}(1, 0)$), which Pareto dominates $(0, 5)$.

EXAMPLE 2.5—*Signalling and then Compromising*. Suppose that in the game[4] of Figure 6, the players wish to achieve an expected payoff of $(4, 4)$. This could be done by playing as in Example 2.3 ("Compromising")—i.e., $\frac{1}{2}UL + \frac{1}{2}DR$—in the game T, and playing A in B. But then, Colin would have to know which is the true game. Can Rowena credibly convey this information to him?

She can; but she must tell him the true game *first*. Only *afterwards*, if it is T, should they perform the lottery that determines whether to play UL or DR. Indeed, suppose the lottery is performed first, and calls for DR in T; and suppose the true game is indeed T. Then Rowena will be motivated to say that it is B, which calls for Colin to play A; this yields to Rowena a payoff of 3, rather than the 2 she would have gotten under the compromise.

---

[4]We thank Elchanan Ben-Porath for pointing out an error in a previous version of this example.
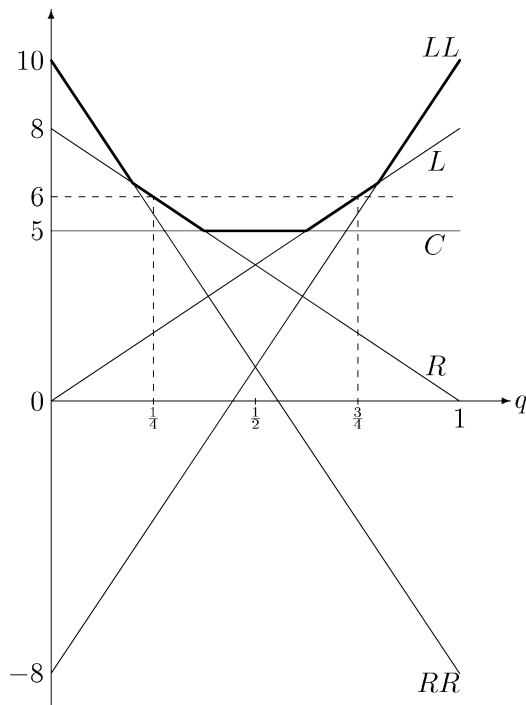
FIGURE 5.—Colin's payoffs in Example 2.4.

We have here the beginnings of a negotiation over time: a process requiring two stages, consisting of a revelation of information *followed* by a compromise.

EXAMPLE 2.6—*Compromising and then Signalling.* This game (Figure 7) differs from that of Figure 4 in that not only L and R, but also LL and RR, yield positive payoffs to Rowena. The silent game still has a unique equilibrium payoff, namely $(0, 5)$. In the talking game, partial revelation (as in Example 2.4) yields $(3, 6)$, which Pareto dominates $(0, 5)$, as before. But now there is also

		L	R	A	
T	U	$6, 2$	$0, 0$	$3, 0$	$\frac{1}{2}$
	D	$0, 0$	$2, 6$	$3, 0$	

		L	R	A	
B	U	$0, 0$	$0, 0$	$4, 4$	$\frac{1}{2}$
	D	$0, 0$	$0, 0$	$4, 4$	

FIGURE 6.—Signalling and then compromising.

	LL	L	C	R	RR	
T	$1, 10$	$3, 8$	$0, 5$	$3, 0$	$1, -8$	$\frac{1}{2}$

	LL	L	C	R	RR	
B	$1, -8$	$3, 0$	$0, 5$	$3, 8$	$1, 10$	$\frac{1}{2}$

FIGURE 7.—Compromising and then signalling.

a *fully* revealing equilibrium whose payoff Pareto dominates $(0, 5)$: Rowena simply tells Colin which game is being played (using a noiseless channel), and Colin responds accordingly, which yields a payoff of $(1, 10)$.

Rowena prefers the partial revelation, Colin the full revelation. So they can compromise: decide on full or partial revelation by tossing a fair coin. This yields $(2, 8)$. Clearly, the compromise must come *before* the signal, since it determines the nature of the signal. So, as in the previous example, the conversation has two stages, but the order of signalling and compromising is reversed.

EXAMPLE 2.7—*Signalling, Compromising, and Signalling.* The talking version of this game (Figure 8) has an equilibrium with three "talk stages": first a partial revelation, then a compromise governed by a jointly controlled lottery, and finally either a full or a partial revelation (depending on the lottery's outcome). First, Rowena reveals whether or not the true game is[5] *BB*. If it is not, then the game is that of Figure 7, which calls for a compromise followed by either a full or a partial revelation, and yields an expectation of $(2, 8)$ (see the previous example). If it is, then Colin plays A, which again yields $(2, 8)$.

If the compromise came before the first revelation, and called for Rowena to reveal fully at the third talk stage—thus yielding her 1—then at the first talk stage, she would have been motivated to say that the true game is *BB*. Colin would then play A, which yields her 2 in both T and B—more than the 1 she

	LL	L	C	R	RR	A	
T	$1, 10$	$3, 8$	$0, 5$	$3, 0$	$1, -8$	$2, 0$	$\frac{1}{3}$

	LL	L	C	R	RR	A	
B	$1, -8$	$3, 0$	$0, 5$	$3, 8$	$1, 10$	$2, 0$	$\frac{1}{3}$

	LL	L	C	R	RR	A	
BB	$0, 0$	$0, 0$	$0, 0$	$0, 0$	$0, 0$	$2, 8$	$\frac{1}{3}$

FIGURE 8.—Signalling, compromising, and signalling.

[5]Rowena's message is thus "two B" or "not two B."

	LL	L	C	R	RR	
T	6, 10	10, 9	0, 7	4, 4	3, 0	$\frac{1}{2}$

	LL	L	C	R	RR	
B	3, 0	4, 4	0, 7	10, 9	6, 10	$\frac{1}{2}$

FIGURE 9.—An unboundedly long conversation.

gets according to plan. Nor can the compromise come after the final revelation, for the reason given in the previous example. Thus it must come at the precise time at which it is scheduled—between the two revelations—neither sooner nor later. So here the conversation has three stages.

EXAMPLE 2.8—*An Unboundedly Long Conversation*. This example (Figure 9), due to Forges (1990a), has a nice interpretation in terms of a job assignment scenario. "Unboundedly long" does not mean "infinitely long": The conversation ends almost surely, but one cannot say beforehand by when. In the silent game, Colin is motivated to choose C, so that Rowena gets 0. By sending a single—fully revealing—message, she can motivate Colin to choose LL or RR, as the case may be, and so improve her payoff to 6. The question arises whether the talking game has an equilibrium that yields her more than 6.

If the conversation may not exceed some given length, the answer is "no"; but if we permit unboundedly long conversations, the answer is "yes": There is an equilibrium yielding $(7, 9\frac{2}{7})$. Since the example is already in the literature, we do not go further into it here. Aumann and Maschler (1995, p. 311) provide an extended discussion of this striking example in the context of repeated games; their discussion applies, almost without change, to the context of cheap talk.

3. THE MODEL: FORMULATION

Let $\Gamma = (\Gamma^1, \ldots, \Gamma^\kappa)$ be a collection of two-person games in strategic form (bimatrices), all of the same size; $p = (p^1, \ldots, p^\kappa)$ a probability vector; and M a "keyboard"—a finite set with members a, b, ..., called "keystrokes." We assume that $|M| > 1$; i.e., there are at least two keystrokes. The model has three phases, each ending before the next starts. In the *information phase,* a bimatrix $\Gamma^{\mathbf{k}}$ is picked at random from the collection Γ in accordance with the probability vector p; the row player Rowena is informed of $\mathbf{k}$, the column player Colin is not.[6] The *talk phase* has infinitely many time periods $1, 2, \ldots$. At each period t,

---

[6]Boldface letters signify random variables. The value of $\mathbf{k}$ is *Rowena's type*, or the *state of nature*, or simply the *state*.

each player sends a message to the other, consisting of a single keystroke; the two messages are sent simultaneously. The messages for the next period, $t + 1$, are sent only after the messages of period t are received. Finally, in the *action phase,* the players simultaneously choose a row and a column, and are paid off in accordance with the entry in $\Gamma^{\mathbf{k}}$ in the chosen row and column.

Perfect recall is assumed throughout the three phases; i.e., at any stage, each player knows what he previously did as well as what he previously knew. The entire model is assumed commonly known between the players.

This defines a *cheap talk game*, or simply *talking game*, in extensive form. In form it is a little unusual, in that time has order type $\omega + 1$; that is, there is an infinite sequence of time periods, with one additional period *after* the whole sequence. Though conceptually this offers no difficulty, formally defining "strategy" in such a context is a little delicate; we postpone doing so until Section 8, after we have discussed the model and results on a more substantive level.

4. THE MODEL: DISCUSSION

4.1. *Interpretation of the Keys*

Use of the term "keyboard" in describing the talk phase is deliberate. The keystrokes are empty vessels, given meaning only by the particular equilibrium of which they are part. Thus in Example 2.1, the following is an equilibrium: At the first stage of the talk phase, Rowena sends the message a if the true game is T, and b if the true game is B; Colin sends the message a. At each subsequent stage, each player sends the message a. In the action phase, Colin chooses L if Rowena sent a at the first stage of the talk phase, R if she sent b. The thing to note is that there is nothing intrinsic about the interpretation of the keys; the equilibrium works just as well if we associate b with T and a with B.

The size of the keyboard—the number of keys—is unimportant; any size—2 or larger—yields the same results. In fact, the equilibrium payoffs would not change even if the keyboard had different sizes at different stages, or even if it were infinite at some or all stages; this follows from our proof (see the Appendix).[7]

4.2. *Conversation Length*

Some readers may raise eyebrows at the infinite length of the talk phase. The reason is to provide maximum latitude for negotiation. Any artificial restriction on the length of the conversation would distort the outcome; inter alia, terminal effects could propagate backwards, as in the finitely repeated Prisoner's

---

[7]However, in more general contexts—when there is initial incomplete information on both sides—the set of equilibrium payoffs may depend on the size of the keyboard (Amitai (1996a)).

Dilemma, or in overlapping generations models with finite horizon. The players should be allowed to talk for as long as they wish, and this is impossible with a finite talk phase.

Note that in all the examples of Section 2, the conversation is finite.[8,9] For cheap talk with *bounded* conversation length, see Section 9(c).

4.3. *Simultaneous Messages, Compromises, Joint Lotteries*

Though common courtesy calls for people to speak one at a time, the model allows for simultaneous messages. There are several reasons for this. Conceptually, a model with alternating speakers could, in principle, introduce undesirable and artificial asymmetries. For example, there is sometimes a significant asymmetry between making an offer and responding to it. Complete symmetry is maintained in the model by the expedient of simultaneous messages.

Technically, simultaneous messages enable the construction of *jointly controlled lotteries*—or simply *joint lotteries*—even when there is no random device with jointly observed outcomes (Aumann, Maschler, and Stearns (1968)). Such lotteries are essential for compromising. Thus suppose that in the Battle of the Sexes (Example 2.3), the players wish to randomize between the outcomes $(6, 2)$ and $(2, 6)$ with $\frac{1}{2}$–$\frac{1}{2}$ probabilities. They can do so by each sending, in the first stage of the talk phase, the messages a and b with probability $\frac{1}{2}$ each; and then, in the action phase, playing UL if the messages were the same (aa or bb), and DR if they were different (ab or ba). Indeed, the probability of each of these two eventualities is $\frac{1}{2}$, and neither player can by himself change it; moreover, once it is determined what the players are to play—UL or DR, as the case may be—either player would lose by deviating. So the described strategies are indeed in equilibrium.[10]

For talk *without* simultaneous messages—which we call *polite talk*—see Section 9(b).

5. CHARACTERIZING THE EQUILIBRIUM PAYOFFS BEHAVIORALLY

5.1. *Formulation*

An *equilibrium* of the talking game is an ordinary Nash (or Bayesian Nash) equilibrium. That is, Colin's strategy maximizes his expected payoff given Rowena's strategy and the probability vector p; and the strategy of each of

---

[8]Or better, *effectively* finite: After a certain stage, the messages do not affect the outcome.

[9]In Example 2.8 the length is unbounded, but still finite with probability 1.

[10]Composing such $\frac{1}{2}$–$\frac{1}{2}$ joint lotteries can generate any λ–λ' lottery between two outcomes X and Y (where $0 < \lambda < 1$ and $\lambda' = 1 - \lambda$). Indeed, express λ as a binary number $0.\varepsilon_1\varepsilon_2\ldots$ (with $\varepsilon_n \in \{0, 1\}$), and conduct a sequence of $\frac{1}{2}$–$\frac{1}{2}$ joint lotteries until the first time aa or bb obtains; if this occurs at stage n, then the outcome is determined by the nth digit of λ (it is X if $\varepsilon_n = 1$ and Y if $\varepsilon_n = 0$).

Rowena's types k maximizes that type's payoff (i.e., the payoff in Γ^k) given Colin's strategy. Payoffs are $(\kappa + 1)$-dimensional objects, denoted (a, β); here a is a κ-vector, whose components are the payoffs to Rowena's κ types, and β is Colin's expected payoff (expectation over k).

Call a pair of strategies in the talking game *canonical* if it has the following form: In the talk phase, even-numbered periods are devoted to joint lotteries, odd-numbered periods to unilateral signals from Rowena to Colin. Specifically, in each even period, the players conduct a joint lottery with $\frac{1}{2}$–$\frac{1}{2}$ probabilities (using the messages a and b; see Section 4.3); denote the outcome s if the messages are the same, d if they are different. In each odd period t, Rowena sends the messages a, b with respective probabilities $\pi, 1 - \pi$; here π may depend on k, on t, on the messages (a or b) that Rowena actually sent in odd periods before t, and on the lottery outcomes (s or d) actually obtained in even periods before t. In the action phase, the players choose mixed actions, which may depend on all the messages that Rowena actually sent at odd periods of the talk phase, and on all the lottery outcomes actually obtained in even periods; in addition, Rowena's choice may depend on k. Call an equilibrium *canonical* if the strategy pair constituting it is canonical.

THEOREM A: *In a cheap talk game, a vector (a, β) is the payoff to an equilibrium if and only if it is the payoff to a canonical equilibrium.*

This formalizes the view of the talk phase as a series of revelations by Rowena, interspersed with agreements between Rowena and Colin. Though the "agreements" are technically lotteries, the informal interpretation is broader. Bargaining may be viewed as a lottery; a priori, either side may come out "ahead."[11] Thus what this paper suggests is that extended negotiations—with revelations interspersed between "pure" bargaining sessions—occur naturally in the context of cheap talk, even when one side is strictly better informed than the other.

Theorem A characterizes the equilibrium payoffs in "behavioral" terms. One would like also to have an explicit characterization, which gives a clear picture of what the equilibrium payoffs actually are. In the ensuing sections (6–7, after the demonstration of Theorem A in 5.2) we provide such a characterization, in terms of the theory of dimartingales (Aumann and Hart (1986)).

5.2. *Informal Demonstration of Theorem A*

For simplicity we take the keyboard to consist of just two keystrokes, a and b. The formal proof does not use this restriction.

---

[11]This is not to say that all or even most bargaining is conducted by tossing a coin; rather, that ex ante, the outcome is not clear to either side—even when no substantive information is revealed during the bargaining process.

"If" is immediate. To demonstrate "only if," let $\gamma' = (\sigma', \tau')$ be an equilibrium of the talking game. We wish to construct a canonical equilibrium $\gamma = (\sigma, \tau)$ with the same payoffs. At each period of the talk phase, γ' calls for Rowena and Colin to send *simultaneous* messages, which depend on the history up to then; assume for the moment that Colin chooses a or b with probabilities $\frac{1}{2}$–$\frac{1}{2}$. We mimic this in γ by a message from Rowena *followed* by a joint lottery; thus each period in γ' is replaced by two periods in γ. Rowena's t'th period message in γ' becomes her $(2t-1)$'th message in γ, and Colin's t'th period message in γ' is replaced in γ by a joint lottery in period $2t$ (with, say, s corresponding to a and d to b). This induces a correspondence between γ-histories and γ'-histories, finite as well as infinite.[12] The correspondence between finite histories defines γ in the talk phase, and that between infinite histories defines γ in the action phase. In particular, γ specifies the same (mixed) actions for a given infinite γ-history that γ' specifies for the corresponding γ'-history.

Clearly, γ is canonical and has the same payoffs as γ'. To see that it is an equilibrium, note first that nothing that Colin can do can change the outcomes of the joint lotteries. Thus Colin can influence the outcome at the action phase only; but there his choice was a best reply to σ', so it is also a best reply to σ. Therefore τ is indeed a best reply to σ.

Also nothing that Rowena can do can change the probabilities of the joint lotteries. *Her* only influence is in the messages she sends in the odd periods and the mixed actions she takes in the action phase. But then again, since these are the same in σ as in σ', it follows that σ is a best reply to τ.

Here we have assumed for simplicity that Colin's signals in γ' all have $\frac{1}{2}$–$\frac{1}{2}$ probabilities. The argument may be extended to arbitrary probabilities; see footnote 10.

6. DIMARTINGALES

Recall that a *martingale* is a sequence $\mathbf{z} = (\mathbf{z}_0, \mathbf{z}_1, \mathbf{z}_2, \ldots)$ of random variables, with values in some Euclidean space $\mathbb{Z}$, such that for each t, the expectation of $\mathbf{z}_{t+1}$ given $(\mathbf{z}_0, \mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_t)$ is $\mathbf{z}_t$. Intuitively, one may think of a particle splitting, each resulting particle splitting again, and so on, the center of mass of each particle always remaining the same. One may also think of a tree, with an element z_v of $\mathbb{Z}$—called $\mathbf{z}$'s *value at* v—associated with each node v, and a probability distribution on v's sons; the martingale condition is that $\mathbf{z}$'s value at v be equal to the expectation of $\mathbf{z}$'s value at v's sons. A classical instance of a martingale is the sequence of posterior probabilities of an event when one conditions on more and more information. More generally, the sequence of conditional expectations of some fixed random variable, when conditioned on more

---

[12]For example, the 3-period history $(a, a), (a, b), (b, b)$ in γ' corresponds to all 6-period histories in γ of the form $(a, \cdot), s, (a, \cdot), d, (b, \cdot), d$—for instance, $(a, b), (a, a), (a, a), (b, a), (b, b), (a, b)$. Note that in γ, Colin's messages in odd periods are irrelevant.

and more information, is a martingale. An important property of martingales, which we use repeatedly, is the *Martingale Convergence Theorem*: A bounded martingale converges almost surely. The *expectation* of a martingale is the expectation of any z_t; by the definition of martingale, all these expectations are the same.

Now let $\mathbb{A}, \mathbb{B}, \mathbb{Q}$ be Euclidean spaces, from now on fixed, and let $\mathbb{Z} = \mathbb{A} \times \mathbb{B} \times \mathbb{Q}$; thus each point in $\mathbb{Z}$ has the form (a, β, q). A *dimartingale*[13] is a $\mathbb{Z}$-valued bounded martingale $((\mathbf{a}_0, \boldsymbol{\beta}_0, \mathbf{q}_0), (\mathbf{a}_1, \boldsymbol{\beta}_1, \mathbf{q}_1), (\mathbf{a}_2, \boldsymbol{\beta}_2, \mathbf{q}_2), \ldots)$, where $\mathbf{a}_{t+1} = \mathbf{a}_t$ when t is even, $\mathbf{q}_{t+1} = \mathbf{q}_t$ when t is odd, and $(\mathbf{a}_0, \boldsymbol{\beta}_0, \mathbf{q}_0)$ is a constant (which is therefore also the expectation of the dimartingale). Thus the q- and a-coordinates never split simultaneously; rather, they alternate: The a-coordinate splits in odd periods, the q-coordinate in even periods. The β-coordinate may split in any period.

The *di-span*[14] of a set G in $\mathbb{Z}$ is defined as the set of all expectations of dimartingales whose limits are almost surely in G. To get a feeling for the di-span, consider the corresponding notion for ordinary (rather than di-) martingales. Thus, (provisionally) define the "span" of G as the set of all expectations of bounded martingales whose limits are almost surely in G. Then it may be seen
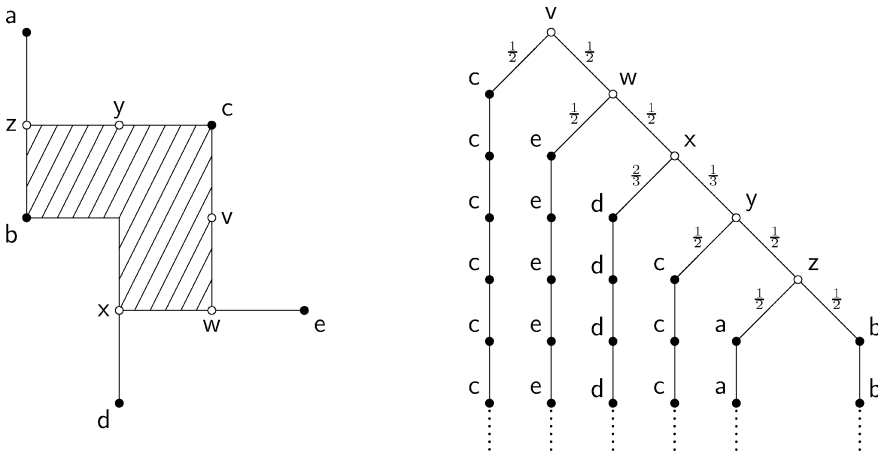

FIGURE 10.—*Left*: The set $G = \{a, b, c, d, e\}$ (the solid dots) and its di-span (the hatched area and the line segments). *Right*: A dimartingale for v, which reaches G in 5 steps.

[13]This term was coined by Aumann and Maschler (1995), though the concept was introduced already in Hart (1985). A *bi*martingale, as in Aumann and Hart (1986), is a dimartingale without the β-coordinate. The definitions and results of Aumann and Hart (1986) translate in a straightforward manner from the "bi" to the "di" setup. For example, a set $G \subset \mathbb{A} \times \mathbb{B} \times \mathbb{Q}$ is *diconvex* if its sections corresponding to constant a—i.e., $\{(\beta, q) \in \mathbb{B} \times \mathbb{Q} : (a, \beta, q) \in G\}$—and to constant q—i.e., $\{(a, \beta) \in \mathbb{A} \times \mathbb{B} : (a, \beta, q) \in G\}$—are all convex sets.

[14]Short for "diconvex span (relative to the a- and q-directions)," the term used in Aumann and Maschler (1995).
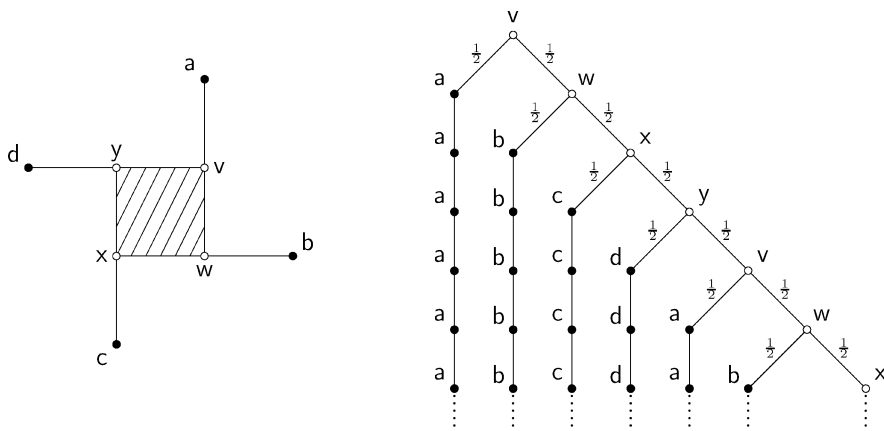
FIGURE 11.—*Left*: The set $G = \{a, b, c, d\}$ (the solid dots) and its di-span (the square and the four line segments). *Right*: A dimartingale for v, which reaches G with probability 1.

that the span of G is the same as its convex hull. Since a dimartingale is a special case of a martingale, it follows that the di-span of a set is included in its convex hull. The di-span is illustrated in Figures 10 and 11 in two instances, both with $\mathbb{A}$ (horizontal axis) and $\mathbb{Q}$ (vertical axis) one-dimensional and no $\mathbb{B}$. In Figure 10 one needs up to 5 consecutive operations of "diconvexification" (convexifying either horizontally or vertically).[15] In Figure 11, diconvexification does not help—G is already diconvex—and the di-span is obtained using dimartingales that are unbounded in length (this so-called "four frogs" configuration is Example 2.5 in Aumann and Hart (1986); it is the basis of Forges' (1990a) example—our Example 2.8).

7. CHARACTERIZING THE EQUILIBRIUM PAYOFFS GEOMETRICALLY

7.1. *Formulation*

In Section 3, we defined the *talking game* determined by a collection $\Gamma = (\Gamma^1, \ldots, \Gamma^\kappa)$ of bimatrices and a probability vector $p = (p^1, \ldots, p^\kappa)$. In order to characterize the equilibria of the talking game, we need to refer to a game defined in precisely the same way, except that the talk phase is eliminated, so that the action phase comes immediately after the information phase. This is called the *silent game* (determined by Γ and p); in it, the informed player cannot send any signals to the uninformed player, and the players cannot reach compromise agreements. As in the talking game, payoffs are $(\kappa+1)$-dimensional objects, denoted (a, β), where a is a κ-vector, whose components are the payoffs to the κ types of Rowena, and β is Colin's expected payoff.

---

[15]In this example, the di-span of G is the same as the *diconvex hull* of G—the smallest diconvex set containing G.

An *equilibrium* of the silent game is an ordinary Nash (or Bayesian Nash) equilibrium: Colin chooses a mixture of columns, and for each k Rowena chooses a mixture of rows, such that Colin's choice maximizes his expected payoff given p and given Rowena's choices, and for each k, Rowena's choice maximizes her payoff in Γ^k given Colin's choice.

To characterize the equilibria in the talking game determined by Γ and p, one must consider probability vectors $q = (q^1, \ldots, q^\kappa)$ other than p. The reason is that the signals that Colin gets from Rowena during the talk phase may change his initial probabilities for the various Γ^k. In fact, it is convenient to consider all possible q—the entire simplex of κ-dimensional probability vectors. We proceed as follows:

- For each q, denote by $\mathcal{E}(q)$ the set of all equilibrium payoffs in the silent game determined by Γ and q.
- When some of the probabilities q^k vanish, we will, for reasons explained later, modify the set of equilibrium payoffs by allowing the corresponding types of Rowena to get *more* than what the equilibrium payoff provides. Denote the set of modified payoffs by $\mathcal{E}^+(q)$; thus a $(\kappa + 1)$-vector $(\hat{a}, \beta)$ is in $\mathcal{E}^+(q)$ if and only if there is an (a, β) in $\mathcal{E}(q)$ with $\hat{a}^k \geq a^k$ for all k and $\hat{a}^k = a^k$ when $q^k > 0$.
- Finally, the graph of the correspondence $\mathcal{E}^+$ is denoted $\operatorname{gr} \mathcal{E}^+$; it lives in a $(\kappa + 1 + \kappa)$-dimensional Euclidean space $\mathbb{A} \times \mathbb{B} \times \mathbb{Q}$ with points (a, β, q).

THEOREM B: *Let $p^k > 0$ for all $k = 1, \ldots, \kappa$. Then (a, β) is the payoff to an equilibrium in the talking game determined by Γ and p if and only if (a, β, p) is in the di-span of* $\operatorname{gr} \mathcal{E}^+$.

To recapitulate: To get the set of equilibrium payoffs in the talking game determined by Γ and p, consider, for *all* probability vectors q, the equilibrium payoffs of the *silent* game determined by Γ and q. On the boundary of the q-simplex only, modify these payoffs by allowing types of Rowena with vanishing q-probability to get *more*. This defines the *modified equilibrium payoff correspondence*; consider the di-span of its *graph*. The section of this di-span corresponding to $q = p$ is the set of equilibrium payoffs in the talking game determined by Γ and p. See Figure 12 for a schematic representation: The graph of the (modified) equilibrium payoff correspondence of the silent game is the curve vewxyz; in particular, the silent game at p has a unique equilibrium payoff, given by e. The di-span of this graph is the hatched area, and the section of the di-span at $q = p$—the segment ef—gives the equilibrium payoffs of the talking game at p.

In addition to its intrinsic interest, this geometric characterization is useful on a "practical" level: It enables the explicit calculation of the equilibrium payoffs (using di-convexity and di-separation arguments; see Aumann and Hart (1986)).
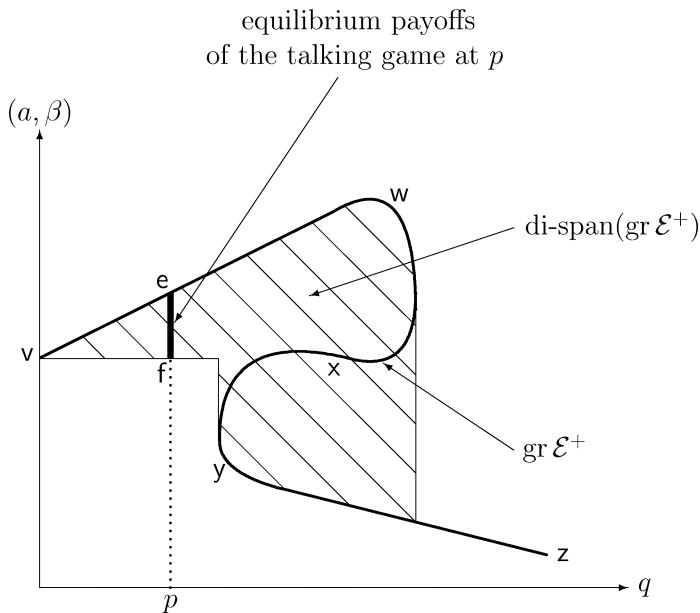
equilibrium payoffs
of the talking game at p

FIGURE 12.—Equilibrium payoffs.

7.2. *Informal Demonstration of Theorem B*

7.2.1. *"Only If": From Equilibrium to Martingale*

The demonstration is a matter of some delicacy. We first describe the rough idea, then point out the difficulties and how to overcome them. The difficulties are conceptual rather than technical; they lie at the heart of the matter.

As in the demonstration of Theorem A, we take the keyboard to consist of the two letters a and b only; this restriction will not be used in the formal proof in the Appendix.

Let γ be an equilibrium of the talking game determined by Γ and p; set $\gamma = (\sigma, \tau) = (\sigma^1, \ldots, \sigma^\kappa, \tau)$, where σ^k is the strategy of Rowena's type k (henceforth Rowena^k). Let the expected payoff of γ be (a, β). We must construct a dimartingale whose expectation is the $(\kappa + 1 + \kappa)$-dimensional vector (a, β, p) and whose limit is almost surely in gr $\mathcal{E}^+$.

By Theorem A, we may take γ to be canonical, and thus to induce an infinite random stream of "messages"—signals from Rowena alternating with joint lottery outcomes. We model the process as a binary tree. Each of the 2^t vertices at level t corresponds to a history up to and including the t'th period. The root is thus the beginning of the game (before any messages—the "empty history"). It has two sons, corresponding to the two signals that Rowena may send in the first period. Each of these sons has two sons, generated by the two outcomes of the joint lottery performed in the second period. Continue in this way, alter-

nating signals with lotteries according to the strategies in γ. For each vertex v, let β_v be Colin's payoff at[16] v; let a_v^k be the payoff of Rowena^k at v; let p_v^k be Colin's probability at v that Rowena is of type k; and set $a_v = (a_v^1, \ldots, a_v^\kappa)$ and $p_v = (p_v^1, \ldots, p_v^\kappa)$. Denote the two sons of v by w and w', and the whole process by $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$.

To make $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ into a martingale, we must specify a probability distribution on the sons of each node v, so that the value of $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ at v is the expectation of its values at its sons. At lottery nodes, there is no problem; simply use the lottery probabilities $\frac{1}{2}$–$\frac{1}{2}$. But matters are more complicated at signalling nodes. Several possibilities come to mind. For each type $k = 1, 2, \ldots, \kappa$, we can use the probabilities π_w^k that at v, Rowena^k picks w. This is appropriate for Rowena^k, because the expectation of her payoffs at w and w' equals her payoff at v. But it is inappropriate for types of Rowena other than k, since π_w^k may be different for different k. And it is also inappropriate for Colin's payoff β_v, and for his probability vector p_v; for these, we should use $\pi_w := p_v^1 \pi_w^1 + \cdots + p_v^\kappa \pi_w^\kappa$, which is the total probability (and also Colin's probability) at v that Rowena picks w. So there are $\kappa + 1$ mutually inconsistent possibilities, none of which seem to fulfill all requirements.

To overcome this difficulty, consider a signalling node v at which Rowena^k chooses each signal with positive probability. Then we must have $a_w^k = a_{w'}^k$; for if $a_w^k > a_{w'}^k$, say, then Rowena^k could gain by picking w with probability 1. So

$$(7.1) \qquad a_w^k = a_{w'}^k = a_v^k,$$

since the payoff at v is an average of the payoffs at w and w'. But then $a_v^k = \lambda a_w^k + (1 - \lambda) a_{w'}^k$ for *any* probability λ; in particular, for $\lambda = \pi_w$. We conclude that *if* each type of Rowena sends each signal with positive probability at each signalling node, and if at signalling nodes v we assign probability π_w to w being picked, then $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ does become a martingale.

Moreover, it is a *di*martingale. Indeed, we have just seen that if v is a signalling node at which both sons have positive probability, then $a_w = a_{w'}$. If, on the other hand, v is a lottery node, then Rowena's message at v is independent of k, and so has no informational content; therefore $p_w = p_{w'} = p_v$, confirming the dimartingale property.

By the Martingale Convergence Theorem, the martingale $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ converges almost surely; denote the limit $(\mathbf{a}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty)$. At the action phase, the players are faced with the *silent* game in which Colin's probability for k is $\mathbf{p}_\infty^k$. So at the action phase they must, almost surely, play an equilibrium in this silent game, since γ is an equilibrium in the talking game. So $(\mathbf{a}_\infty, \boldsymbol{\beta}_\infty)$ must be an equilibrium payoff in this silent game; that is, $(\mathbf{a}_\infty, \boldsymbol{\beta}_\infty)$ is in $\mathcal{E}(\mathbf{p}_\infty)$. So $(\mathbf{a}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty)$ is in $\mathrm{gr}\,\mathcal{E}$. We have already seen that $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ is a dimartingale; its expectation is (a, β, p) (what the players expect at stage 0). Thus (a, β, p)

---

[16]Here "payoff" means "expected payoff," and "at v" means "conditional on v being reached."

is the expectation of a dimartingale whose limit is almost surely in gr $\mathcal{E}$; that is, (a, β, p) is in the di-span of gr $\mathcal{E}$. So a fortiori, it is in the di-span of gr $\mathcal{E}^+$, as was to be shown.

So far, so good. But what if σ^k calls, at a signalling node v, for Rowena^k to pick w, say, for *sure*? The above reasoning then breaks down: She is indeed motivated to choose the signal with the higher payoff, and does choose it, for sure. So we do not get a dimartingale, or for that matter, even a martingale.

But that is not the worst of it. It is not clear what at all is meant by "choosing the signal with the higher payoff." What *is* the payoff? Since w' is impossible under σ^k, how is $a_{w'}^k$ *defined?* If Colin hears the signal leading to w', he cannot conclude that Rowena has "deviated." He does not know her type, and w' may well be possible under σ for types other than k; so Colin will continue to do what τ prescribes. But how will Rowena herself continue? Her strategy σ^k prescribes responses to possible deviations of Colin, but not to her own deviations.[17]

These are not minor technical difficulties. The case where Rowena's signal makes a difference to *her* is basic in cheap talk; it is the issue in our very first example (2.1). This cannot be swept under the rug.

To deal with it, we revise the definition of a_v^k. Any node v is possible under Colin's strategy τ, so what Colin does after[18] v is determined by τ. Redefine a_v^k as the expected payoff of Rowena^k at v if from then on, she plays a best reply to Colin's strategy. This definition of a_v^k is meaningful whether or not v is k-possible (i.e, possible under σ^k); in the former case, it coincides with the previous definition (i.e., as Rowena^k's expected payoff). Indeed, σ^k is then itself a best reply to τ, by the definition of equilibrium; so the best-reply payoff is the same as the expected payoff. Thus as shown above (see (7.1)), if v is a signalling vertex both of whose sons are k-possible, then $a_w^k = a_{w'}^k = a_v^k$ with the revised definition as well.

Now return to the case when w, say, is k-possible, but w' is not. Then

$$(7.2) \qquad a_v^k = a_w^k \geq a_{w'}^k;$$

for otherwise Rowena^k would be motivated to pick w' instead of w, and then continue with a best reply to τ. But there is nothing to prevent strict inequality in (7.2), so $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ is *not* a martingale.

The key to resolving the difficulty is the remark that though w' is a perfectly good node in the tree, it cannot occur under σ^k. If Colin hears the signal

---

[17]To be sure, the standard definitions of "strategy" do prescribe what a player should do in situations that she herself has excluded. But that is a matter of convenience only; when discussing ordinary (rather than perfect) equilibria, as here, such prescriptions have no optimality properties, and play no role in the analysis.

[18]The significant part here is what τ prescribes for the action phase; for the talk phase, it is $\frac{1}{2}$–$\frac{1}{2}$ at each stage.

leading to w', he *knows* that Rowena's type is not k, so $p^k_{w'} = 0$; that is, in equilibrium Rowena never actually gets $a^k_{w'}$. So if we modify the value of $a^k_{w'}$, we do not change the expectation of the whole process; it remains a^k. Specifically, we can raise the value of $a^k_{w'}$ to the level of a^k_w; the expectation of the process is still a^k, *and* we get the required dimartingale property (as in (7.1)). Of course, we must make corresponding modifications on the entire martingale.

The result, denoted $\hat{\mathbf{a}}$ and called Rowena's *virtual* payoff (as distinguished from the *real* payoff $\mathbf{a}$), should be a bounded martingale with

(7.3) $\hat{a}^k_v \geq a^k_v$

and

(7.4) $\hat{a}^k_v = a^k_v$ when $p^k_v > 0$

for all k and v. The actual construction of $\hat{\mathbf{a}}$ will be left to the formal proof in the Appendix.

Define a *history* h as an infinite path in the tree. By the Martingale Convergence Theorem, $(\hat{\mathbf{a}}, \boldsymbol{\beta}, \mathbf{p})$ converges almost surely; its limit $(\hat{\mathbf{a}}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty)$ associates a vector $(\hat{a}_h, \beta_h, p_h)$ with (almost) each history h, namely the limit of the vectors $(\hat{a}_v, \beta_v, p_v)$ for v in h. At the action phase, the players are faced with the *silent* game in which Colin's probability for k is p^k_h. Let a^k_h be the payoff of Rowena^k if γ is played in the action phase; as β_h is Colin's payoff, (a_h, β_h) is in $\mathcal{E}(p_h)$. Now it may be seen that (7.3) yields $\hat{a}^k_h \geq a^k_h$ (a.s.). If $p^k_h > 0$, then $p^k_v > 0$ for each v in h, so by (7.4), $\hat{a}^k_v = a^k_v$ for each v in h, so $\hat{a}^k_h = \lim \hat{a}^k_v = \lim a^k_v = a^k_h$. Thus $\hat{a}^k_h \geq a^k_h$ always, and $\hat{a}^k_h = a^k_h$ when $p^k_h > 0$. Since (a_h, β_h) is in $\mathcal{E}(p_h)$, it follows that $(\hat{a}_h, \beta_h)$ is in $\mathcal{E}^+(p_h)$; i.e., $(\hat{a}_h, \beta_h, p_h)$ is in gr $\mathcal{E}^+$. Since this is so for (almost) all histories h, we deduce that the expectation of the dimartingale $(\hat{\mathbf{a}}, \boldsymbol{\beta}, \mathbf{p})$ is in the di-span of gr $\mathcal{E}^+$. We have already seen that the expectation of $(\boldsymbol{\beta}, \mathbf{p})$ is (β, p); that the expectation of $\hat{\mathbf{a}}$ is a follows from (7.4). So (a, β, p) is in the di-span of gr $\mathcal{E}^+$. This completes the demonstration of "only if" in Theorem B.

To summarize: Start with a random process $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ described by a binary tree, as in the fourth paragraph of this section; assign *Colin's* probabilities to the edges. If each type of *Rowena* assigns positive probability to each node, then the process is a dimartingale; so the expectation (a, β, p) of the process is in the di-span of gr $\mathcal{E}$, and a fortiori in that of gr $\mathcal{E}^+$, as was to be shown. If not, revise the definition of $\mathbf{a}^k$ to be the payoff to a best reply of Rowena^k to Colin's strategy τ. Let v be a k-possible signalling node with exactly one k-possible son w. Then $\mathbf{a}^k$ at the other son w' cannot be larger than at w. If it is equal, leave it as is. If smaller, lift it until it is equal; the result is called the *virtual* payoff of Rowena^k at w'. Then, lift the payoff at the descendants of w' so as to obtain a dimartingale, with the virtual payoff $\hat{\mathbf{a}}^k$ equal to the "real" payoff $\mathbf{a}^k$ at k-possible nodes, and not smaller than the real payoff at k-impossible nodes; thus the expectation of $(\hat{\mathbf{a}}, \boldsymbol{\beta}, \mathbf{p})$ is (a, β, p), and is in the di-span of gr $\mathcal{E}^+$, as was to be shown.

7.2.2. *"If": From Martingale to Equilibrium*

Here we are given a dimartingale $(\hat{\mathbf{a}}, \boldsymbol{\beta}, \mathbf{p})$ whose expectation (a, β, p) is in the di-span of $\mathrm{gr}\,\mathcal{E}^+$, and construct an equilibrium $\gamma = (\sigma^1, \ldots, \sigma^\kappa, \tau)$, with payoff (a, β), of the talking game determined by Γ and p.

We again use the tree representation. Call a node v a "lottery node" or a "signalling node" according to whether $\hat{\mathbf{a}}$ or $\mathbf{p}$ splits there. Assume for simplicity that the tree is binary and that the probabilities on the edges emanating from lottery nodes are $\frac{1}{2}$–$\frac{1}{2}$. We first describe the strategies in the talk phase. At lottery nodes, both players play $\frac{1}{2}$–$\frac{1}{2}$. At signalling nodes v, only Rowena sends a signal—or equivalently, picks a son w—and her choice may depend on her type k. If v is k-possible, define the probability that σ^k picks w at v as $\pi_w^k := p_w^k \pi_w / p_v^k$, where π_w is the probability that the dimartingale assigns to the edge leading from v to w; note that the terms on the right side are given as part of the martingale. One may verify that this is the "right" definition of π_w^k; that is, Colin's probability at w of Rowena^k is precisely p_w^k (use Bayes's rule inductively). If v is k-impossible, it does not matter how σ^k is defined there.

At the action phase, the players have behind them a history h consisting of nodes in the talk phase. Denote the limit of $(\hat{a}_v, \beta_v, p_v)$ for v in h by $(\hat{a}_h, \beta_h, p_h)$; by the Martingale Convergence Theorem it exists, and by the hypothesis is in[19] $\mathrm{gr}\,\mathcal{E}^+$. So there is an equilibrium e_h of the silent game determined by Γ and p_h, with payoff (a_h, β_h), where $\hat{a}_h^k \geq a_h^k$ for all k and $\hat{a}_h^k = a_h^k$ when $p_h^k > 0$. Define γ in the action phase by stipulating that the players play e_h after h.

It can be checked that the payoff of γ is indeed[20] (a, β). To show that γ is an equilibrium, we must show that neither player can profitably deviate in either the action or the talk phase. For the action phase this follows from e_h being a silent game equilibrium. At lottery nodes in the talk phase, neither player can alone change anything, and so cannot gain by deviating. Suppose finally that v is a signalling node, with sons w and w'. Colin has nothing to do at v, and so cannot gain by deviating.

It remains to show that Rowena^k cannot get more than a^k by deviating at the signalling nodes. The crux of the argument here is that, no matter what probabilities Rowena^k uses at the signalling nodes, the sequence $\hat{\mathbf{a}}^k$ always remains a martingale. Indeed, at any signalling node v one has $\hat{a}_v^k = \hat{a}_w^k = \hat{a}_{w'}^k$ by the dimartingale property, so the expectation $\lambda\hat{a}_w^k + (1 - \lambda)\hat{a}_{w'}^k$ equals $\hat{a}_v^k$ for *any* probability λ that Rowena^k may use at v (instead of π_w^k). Therefore, in particular, the expectation of $\hat{\mathbf{a}}_\infty^k$ is $\hat{\mathbf{a}}_0^k = a^k$ for any signalling strategy of Rowena^k. Now the payoff of Rowena^k at the action phase is a_h^k, which is $\leq \hat{a}_h^k$; therefore what she can get is bounded from above by the expectation of $\hat{\mathbf{a}}_\infty^k$—which, as we have just seen, equals a^k no matter how she behaves at the signalling nodes.

---

[19]For almost all histories; henceforth we omit this caveat.

[20]The real payoffs a_v^k can differ from $\hat{a}_v^k$ only when it does not matter; that is, for nodes v with $p_v^k = 0$—which means that the probability of Rowena^k reaching v is 0.

8. STRATEGIES

An unusual aspect of the cheap talk game is that time has order type $\omega + 1$; that is, there is an infinite sequence of time periods, with one additional period *after* the whole sequence. That makes the matter of formally defining strategies—to which we now turn—not entirely straightforward.

Some readers may prefer to skip this section—which they may do with little or no loss in understanding the paper. Strictly speaking, though, it is part of the *statement* of our results, not only of the proofs; that is why we do not relegate it to an appendix. One cannot talk about equilibria if one does not know what a strategy is.

Since we are primarily—indeed exclusively—interested in mixed strategies, we start at once by defining *them*, bypassing the formal definition of pure strategy. Denote by I and J the action sets of Rowena and Colin in the action phase (rows and columns of Γ); let Ξ_R and Ξ_C be copies of the probability space consisting of the unit interval $[0, 1]$, with the Borel sets as events and Lebesgue measure (i.e., the uniform distribution) as the probability measure. A t-period *history*, $t = 0, 1, 2, \ldots$, is a sequence consisting of t pairs of keystrokes. An *infinite* or *complete* history is an infinite sequence of pairs of keystrokes. Denote the set of all t-period histories by H_t, of all infinite histories by[21] H_∞. A *mixed strategy* τ of Colin consists of a sequence of measurable functions $\tau_1, \tau_2, \ldots$, and a measurable function τ_∞, where $\tau_t : \Xi_C \times H_{t-1} \to M$ for $t = 1, 2, \ldots$, and $\tau_\infty : \Xi_C \times H_\infty \to J$; here H_{t-1} has the discrete structure (all sets are measurable), and H_∞ as well as $\Xi_C \times H_\infty$ and $\Xi_C \times H_{t-1}$ have the usual product structure.[22] Rowena's mixed strategies σ are defined similarly: $\sigma_t : \Xi_R \times H_{t-1} \times K \to M$ for $t = 1, 2, \ldots$, and $\sigma_\infty : \Xi_R \times H_\infty \times K \to I$, where K is the collection of states $\{1, \ldots, \kappa\}$ (corresponding to the collection of games $\{\Gamma^1, \ldots, \Gamma^\kappa\}$; recall that Rowena knows the state).

Operationally, the players pick mixed strategies σ and τ, simultaneously, before the beginning of play. Then, chance picks (ξ_R, ξ_C, k) in $\Xi_R \times \Xi_C \times K$ at random, using the uniform distribution on each of Ξ_R and Ξ_C and the distribution p on K, all three being independent (Rowena is informed of ξ_R and k only, and Colin of ξ_C only). At period 1 of the talk phase, the messages sent by Rowena and Colin are $\sigma_1(\xi_R, k)$ and $\tau_1(\xi_C)$, respectively. At period t, the messages are $\sigma_t(\xi_R, h_{t-1}, k)$ and $\tau_t(\xi_C, h_{t-1})$, where h_{t-1} is the $(t-1)$-stage history consisting of the messages sent by both players in the previous $t-1$ stages. In the action phase, the actions taken are $\sigma_\infty(\xi_R, h_\infty, k)$ and $\tau_\infty(\xi_C, h_\infty)$, where h_∞ is the complete history of all messages sent by both players in the talk phase. The bimatrix Γ^k determines the payoff.

---

[21]I.e., $H_t = (M \times M)^t$; note that there is exactly one 0-period history, namely the empty sequence.

[22]For H_∞, this is the smallest σ-field containing all *finite rectangles*: sets of the form $S_1 \times \cdots \times S_t \times (M \times M) \times (M \times M) \times \cdots$ for some finite t, where $S_r \subset M \times M$ for all $r \leq t$.

Set $\Xi := \Xi_R \times \Xi_C \times K$. For fixed σ and τ, each point ξ in Ξ determines payoffs (real numbers) $\mathbf{a}(\xi)$ and $\mathbf{b}(\xi)$ to Rowena and Colin respectively. The total payoff for the mixed strategy pair (σ, τ) is the expectation of $(\mathbf{a}, \mathbf{b})$—its integral over Ξ. For this to be defined, we need the following:

PROPOSITION 8.1: $(\mathbf{a}, \mathbf{b})$ *is a measurable mapping from Ξ to $\mathbb{R}^2$.*

The proof of this proposition may be found in Appendix A.5.

In brief, the $(\omega + 1)$th period—the action phase—is based on infinite histories, of which there are a continuum. A *pure* strategy must specify a function from this continuum to the player's actions; not just any function, though, but one that is measurable in the appropriate sense.[23] A *mixed* strategy, therefore, should be a probability distribution over such functions. Defining distributions over function spaces is not straightforward; see Aumann (1961, 1963). The approach adopted here resembles that of Aumann (1964); for the underlying idea, see Section 3 of that reference.

9. DISCUSSION

This section discusses a number of extensions and variations of our model.

(a) *Complete Information.* In the special case where there is complete information (i.e., $\kappa = 1$, as in Example 2.3), cheap talk adds only the "weighted averages" of the Nash equilibria of the silent game. In other words, here cheap talk is equivalent to adding lotteries with *publicly* observed outcomes to the silent game.[24] Geometrically, the set of equilibrium payoffs of the talking game is just the convex hull of the set of equilibrium payoffs of the silent game.[25]

(b) *Polite Talk.* Consider the case where at each period *only one player* can send a message; call this *polite talk*. The parallel of Theorem B for polite talk is based on dimartingales $((\mathbf{a}_t, \boldsymbol{\beta}_t, \mathbf{q}_t))_{t=0,1,2,\dots}$ that satisfy $\mathbf{a}_{t+1} = \mathbf{a}_t$ for even t and $(\boldsymbol{\beta}_{t+1}, \mathbf{q}_{t+1}) = (\boldsymbol{\beta}_t, \mathbf{q}_t)$ for odd[26] t.

---

[23]For example, if Colin plays $\frac{1}{2}$–$\frac{1}{2}$ at each stage, then we want each pure strategy s of Rowena to induce a payoff; this will be so if and only if s is measurable with respect to the product σ-field on H_∞ (the set of infinite histories). Here, the main thrust of measurability is *not* informational (in the sense of "what can a strategy depend on"); indeed, the product σ-field separates points, which means that a strategy is allowed to treat any two histories differently. Rather, measurability is needed so that payoffs can indeed be "measured" (see the discussion in Aumann (2001)).

[24]Lotteries with *privately* observed outcomes yield the correlated equilibria.

[25]A direct proof of this known result is as follows: In an equilibrium of the talking game, after (almost) every possible history of talk, the mixed actions must form an equilibrium of the silent game, and the talking phase induces a probability distribution over histories, and thus also over silent equilibria. The converse follows using jointly controlled lotteries.

[26]Compare this with the definition of Section 7. The addition here is that in the periods in which only Colin talks, his expected payoffs following all his possible messages must be equal (exactly as is the case for Rowena's payoffs when she talks), which makes the β-coordinate stay constant as well (in addition to the q-coordinate).

3, 4	−5, −5	−5, −5	−5, −5
−5, −5	4, 2	−5, −5	−5, −5
−5, −5	−5, −5	2, 1	−5, −5
−5, −5	−5, −5	−5, −5	1, 3

FIGURE 13.—Unbounded polite talk.

When there is complete information,[27] this yields the bi-span of the set of Nash equilibrium payoffs of the silent game. For a simple example, consider the game of Figure 13: The set of silent equilibrium payoffs that lie in the positive orthant consists of the four diagonal entries of the game matrix (the points $a = (3, 4)$, $b = (4, 2)$, $c = (2, 1)$, and $d = (1, 3)$), to which polite talk adds the square and the four line segments (see Figure 11 Left).[28] For example, the payoff $v = (2, 2)$ is obtained as follows (cf. Figure 11 Right): In the first period, Rowena says "stop" or "continue" with equal probabilities $\frac{1}{2}$–$\frac{1}{2}$. If she said "stop," the talk ends and they play the first row and the first column (i.e., the silent equilibrium with payoffs $a = (3, 4)$). If she said "continue," then in the second period Colin chooses (with equal probabilities) either "stop"—after which they play the second row and the second column (i.e., b)—or "continue." And so on. It may be verified that each player in his or her turn is indifferent between "stop" and "continue," and so one gets an equilibrium of the polite talk game.[29]

(c) *Bounded Talk*, or *Talk with a Deadline*. If there is an a priori bound on the number of talk periods (a "deadline"), the di-span of Theorem B is replaced by the *diconvex hull* of $\operatorname{gr}\mathcal{E}^+$ (Section 6), which is in general a strictly smaller set (for example, in Figure 11, the diconvex hull of G is just G itself).

The same applies to polite talk, even when there is complete information. In the example of Figure 13, bounded polite talk does not yield any new equilibrium payoffs in the positive orthant beyond those of the silent game (i.e., $\{a, b, c, d\}$).

To gain some intuition on the difference between bounded and unbounded talk, think of the finitely vs. the infinitely repeated discounted Prisoner's Dilemma. The finite repetition has a *last period*, from which everything "unravels" backward, so that only a repetition of the one-shot equilibrium is possible. In contrast, the infinitely repeated game has no last period. So in each period, the players can look to the future; what they expect to gain there may balance

---

[27]And thus a and β are each one-dimensional, and there is no q-coordinate.

[28]The silent game has additional equilibrium payoffs, all of them lying in the negative orthant. Separation arguments as in Aumann and Hart (1986) imply that there is no interaction between these and the four points in the positive orthant (since there are no points in the other two orthants); therefore any polite talk equilibrium payoff in the positive orthant can be based only on a, b, c and d.

[29]An interesting sidelight is that though the number of talk periods is unbounded, its expectation is just 2.

out any immediate loss from playing a strategy that is nonoptimal myopically (i.e., from the viewpoint of a single period).

With cheap talk it is similar. For example, in the polite talk game discussed at (b) above, consider an equilibrium with positive payoffs. Each possible talk history[30] must result in an equilibrium at the action phase, which must be one of the four diagonal boxes of the game matrix of Figure 13 (with payoffs `a`, `b`, `c`, and `d`). If the talk phase were bounded in length, then at the last talk period, each possible message would determine one of these outcomes. But all such outcomes would have to yield the same payoff to the player who talks at that period—which could happen only if the outcomes were identical.[31] So the message at the last period would be irrelevant; working backwards, we conclude that all possible histories would lead to the same outcome—so bounded polite talk adds no new (positive) equilibrium payoffs.

This argument depends crucially on there being a last period. If there is no last period, then at each period the player talking there can either stop or continue, both of which lead to the same payoff for him. But if the talk is bounded, then at some point he can no longer choose "continue," and the whole thing unravels.

For talk that ends at a *random* time—specifically, with a certain fixed probability after each period—see Amitai (1996b).

(d) *Publicly Observed Signals*, or *Sunspots*. If, in addition to messages between the players, there are also publicly observed signals, then jointly controlled lotteries are no longer needed. However, as our analysis shows, it is important for such "sunspots" to be available in each period and not just at the beginning of the game, so they can be interspersed with Rowena's signals as needed.

10. THE LITERATURE

For a good nontechnical introduction to some of the main issues of cheap talk, see the survey of Farrell and Rabin (1996).

An early study of the issue examined here—equilibrium expansion by cheap talk—is Crawford and Sobel (1982) (see also Green and Stokey (1980)). Their models are similar to ours, but less general in several ways; most importantly, they permit only one (cheap) message to be sent, from Rowena to Colin. This sounds reasonable on the face of it, as only Rowena has something substantive to say; but as we have seen, it involves significant loss of generality. This is precisely what the current paper explores.

The opposite issue—that of equilibrium restriction, or *refinement*, under cheap talk—has been dealt with extensively by several authors. One approach is to assume a common language (that is, messages have intrinsic given mean-

---

[30]One with positive probability.
[31]Since the four outcomes have different payoffs to each of the two players.

ing); see Farrell (1993), Rabin (1990), Aumann (1990), Matthews, Okuno-Fujiwara, and Postlewaite (1991), Blume and Sobel (1995), Rabin and Sobel (1996), and Zapater (1997). Another approach uses evolutionary stability considerations; see Robson (1990), Wärneryd (1993), Blume, Kim, and Sobel (1993), Kim and Sobel (1995), and Banerjee and Weibull (2000).

The current study grew out of the theory of repeated games of incomplete information with undiscounted payoffs; see Aumann, Maschler, and Stearns (1968) and the large literature cited in Aumann and Maschler (1995), in particular Hart (1985) and Forges (1992). Very loosely speaking, the early stages of repeated game equilibria correspond to the talk phase,[32] the later ones to the action phase. For various reasons, cheap talk is simpler and "cleaner" than repetition; to cite just two examples, one does not have the problem of defining the repeated game payoff as some kind of limiting average, and one avoids the complexities of individual rationality in the repeated context.

In our model, one side is strictly better informed than the other, and so has nothing to learn. When there is incomplete information on *both* sides—both have something to learn—the situation is a good deal more complex; see Amitai (1996a), which is based on the work presented here. Amitai (1996b) contains another extension of our model; see Section 9(c).

Another direction of generalization retains the "cheapness," but replaces the "talk" by something less direct, like a mediator. This is the subject of *general communication systems* or *mechanisms*, consisting of an impartial "*mediator*" who receives *input* messages from the players and sends them *output* messages (for a survey, see Myerson (1994)). All the messages are private (each message is known only to its sender and receiver). The mediator may use randomizations when determining the outputs (as functions of the inputs). However, the mediator is *not* a player: it is a device that follows its (commonly known) program to the letter. A general *communication game* is a game to which a communication system has been added, without affecting the payoffs; its equilibria are called *communication equilibria*. In this framework, the cheap talk equilibria are precisely that subset of the communication equilibria that use only plain public communication. For a restricted class of communication systems—"public mediated talk"—see Lehrer and Sorin (1997).

In games of complete information—when Rowena has just one type—general mechanisms lead to all *correlated equilibria* (Aumann (1974)), whereas long cheap talk leads only to all weighted averages of Nash equilibria (see Section 9(a))—a much smaller set. In games of incomplete information, where inputs from the players may be relevant, the set of communication equilibria is determined by finitely many inequalities (the "incentive-compatibility" constraints, which are obtained using the "Revelation Principle"[33]).

---

[32]In an undiscounted repetition, no finite number of stages can affect the payoff, so they may be used for communication.

[33]Which asserts that every communication equilibrium may be obtained by a *direct mechanism*: one in which the players report their private information to the mediator who then recommends

When there are more than two players, the restriction to direct "talk" is less effective in circumscribing the set of equilibria: With at least three, or four, or five players—depending on the specific conditions—mediators can be emulated by cheap talk. Thus in the complete information case, essentially all correlated equilibrium payoffs in the original game are Nash equilibria in an appropriately specified cheap talk game; and in the case of incomplete information, all equilibrium outcomes of general mechanisms are Nash equilibria of an appropriately specified cheap talk game (Barany (1992), Forges (1990b), Ben-Porath (1998, 2002), Gerardi (2000)). Rather than a sequence of simple agreements and revelations, as here, this literature uses complex, sophisticated communication protocols in which "telephone conversations" (private one-way communication between pairs of players through a secure channel) take the place of a mediator. The presence of more than two players enables the construction of protocols in which no single player can profitably affect the outcome.

A related literature is that of computer science,[34] specifically cryptography, with its "public key" techniques ("one-way" or "trapdoor" functions, "oblivious transfer," and so on). In situations in which the players are computationally restricted and the appropriate cryptographic mechanisms exist, mediators may be emulated even when there are only two players (Yao (1986), Goldreich, Micali, and Wigderson (1987), Goldwasser and Levin (1991), Urbano and Vila (2002)). Moreover, this can be done also when the players have unlimited computational abilities, provided, again, that there are more than two players (Ben-Or, Goldwasser, and Wigderson (1988), Rabin and Ben-Or (1989)).

The spirit of this "emulation" literature is quite different from that of the current work. Here we try to model what may happen in real-life, informal, face-to-face cheap talk, where players react naturally to revelations and agreements. The emulation literature is resourceful and impressive, but contrived; real-life face-to-face negotiations (unlike, perhaps, via the internet) would hardly use such protocols.

There are a growing number of applications of cheap talk to specific models; see the references in Farrell and Rabin (1996) and, recently, Baliga and Morris (2002), Baliga and Sjöström (2001), and Battaglini (2002).

Finally, two recent works are worthy of particular notice. Simon (2002, Example 4, p. 98) exhibits a game in which bounded cheap talk adds nothing to the silent game equilibrium payoffs, whereas unbounded cheap talk yields equilibria that are *Pareto superior* to all silent game equilibria (though the paper is about repeated games, this example applies without change to cheap talk). Krishna and Morgan (2002) exhibit games where just *two* periods of

---

to them what to play, and, where, in equilibrium, everyone reports truthfully and follows the recommendation.

[34]We thank Michael Ben-Or for guiding us here. A recent survey and further references are included in Canetti (2000).

cheap talk yield equilibria that are Pareto superior to all equilibria of the one-period cheap talk.

Center for Rationality and Interactive Decision Theory, and Department of Mathematics, The Hebrew University of Jerusalem, Israel; raumann@math.huji.ac.il,

and

Center for Rationality and Interactive Decision Theory, Department of Economics, and Department of Mathematics, The Hebrew University of Jerusalem, Israel; hart@huji.ac.il; http://www.ma.huji.ac.il/~hart.

A. APPENDIX

This appendix is devoted to the formal proof of our results.

A.1. *Preliminaries*

We start by recalling the notations, definitions, and statements. There are two players, player 1 (Rowena) and player 2 (Colin). $K = \{1, 2, \ldots, \kappa\}$ is the set of *states* (of nature), or, equivalently, of *types* of player 1. I and J are the finite action sets of player 1 and player 2, respectively. For each state $k \in K$, the game Γ^k is given by two payoff matrices A^k and B^k, of size $I \times J$ (i.e., $A^k(i, j)$ and $B^k(i, j)$ are the payoffs to player 1 and player 2, respectively, when player 1 chooses $i \in I$, player 2 chooses $j \in J$, and the state is $k \in K$). The (common) prior probability distribution on K is[35] $p \in \Delta(K)$.

A.1.1. *The Silent Game*

The underlying (or "silent") game $\Gamma(p)$ consists of two phases, as follows:

• *Information phase:* The state $\mathbf{k} \in K$ is chosen according to the probability distribution p. Player 1 is then told the chosen ("true") $\mathbf{k}$, while player 2 is not.[36]

• *Action phase:* Player 1 chooses an action $i \in I$ and, simultaneously, player 2 chooses an action $j \in J$; they receive payoffs $A^{\mathbf{k}}(i, j)$ and $B^{\mathbf{k}}(i, j)$, respectively ($\mathbf{k}$ is the true state).

A mixed strategy of player 1 in $\Gamma(p)$ is a K-tuple of mixed actions, i.e., $x = (x^k)_{k \in K}$, with $x^k \in \Delta(I)$ for all $k \in K$; a mixed strategy of player 2 in $\Gamma(p)$ is a mixed action $y \in \Delta(J)$. Such a pair (x, y) constitutes a (Bayesian–Nash) *equilibrium* in $\Gamma(p)$ if [37]

$$a^k := A^k(x^k, y) = \max_{\widetilde{x} \in \Delta(I)} A^k(\widetilde{x}, y), \quad \text{for all } k \in K; \quad \text{and}$$

$$\beta := \sum_{k \in K} p^k B^k(x^k, y) = \max_{\widetilde{y} \in \Delta(J)} \sum_{k \in K} p^k B^k(x^k, \widetilde{y})$$

---

[35]For a finite set X, the set of probability distributions on X is $\Delta(X) := \{q \in \mathbb{R}_+^X : \sum_{x \in X} q_x = 1\}$, the $(|X| - 1)$-dimensional unit simplex.

[36]Our analysis extends to any case where player 1 knows everything that player 2 knows, and perhaps more (in epistemic terms: the information partition of player 1 is finer than that of player 2). Indeed, any such case may be reduced to a combination of situations where, as in this paper, player 1 knows the state of nature and player 2 does not (see Mertens, Sorin, and Zamir (1994) or Hart (1985, first paragraph on page 120)).

[37]Note that the action of type k of player 1 is a best reply even when that type's probability p^k is 0. This is convenient for the "modified" definition below.

(we have extended A^k and B^k bilinearly to pairs of mixed actions). The resulting *equilibrium payoffs* are the vector $a = (a^k)_{k \in K} \in \mathbb{R}^K$ to player 1 and the scalar $\beta \in \mathbb{R}$ to player 2. Let $\mathcal{E}(p)$ be the set of equilibrium payoffs of $\Gamma(p)$. Note that it is nonempty for all p (replace each type of player 1 by an agent, and apply the existence theorem of Nash (1951) to the resulting $(\kappa + 1)$-player game).

When some coordinates of p vanish, we need the following modification (Hart (1985)): If $p^k = 0$, then the corresponding coordinate a^k of a may be arbitrarily increased. Thus we define for each $p \in \Delta(K)$ the set of *modified equilibrium payoffs* $\mathcal{E}^+(p)$ as the set of all $(a, \beta) \in \mathbb{R}^K \times \mathbb{R}$ such that there exist $(x^k)_{k \in K} \in [\Delta(I)]^K$ and $y \in \Delta(J)$ satisfying:

(A.1) $a^k \geq A^k(x^k, y) = \max_{\widetilde{x} \in \Delta(I)} A^k(\widetilde{x}, y), \quad \text{for all } k \in K;$

(A.2) $a^k = A^k(x^k, y) \quad \text{if } p^k > 0, \quad \text{for all } k \in K; \quad \text{and}$

(A.3) $\beta = \sum_{k \in K} p^k B^k(x^k, y) = \max_{\widetilde{y} \in \Delta(J)} \sum_{k \in K} p^k B^k(x^k, \widetilde{y}).$

The *graph* of the modified equilibrium payoff correspondence is

$$\mathrm{gr}\,\mathcal{E}^+ := \left\{ (a, \beta, p) \in \mathbb{R}^K \times \mathbb{R} \times \Delta(K) : (a, \beta) \in \mathcal{E}^+(p) \right\}.$$

A.1.2. *The Talking Game*

Given the "keyboard" M with $|M| > 1$, the *talking game* (or, *cheap talk extension*) $\Gamma_C(p)$ is obtained by adding a "talk phase" to the silent game $\Gamma(p)$ before the "action phase," but after the initial "information phase." Formally, $\Gamma_C(p)$ is a game of *perfect recall*, consisting of three phases:

- *Information phase:* As in $\Gamma(p)$.
- *Talk phase:* At each $t = 1, 2, \ldots,$ player 1 chooses a message $m_t^1 \in M$ and player 2 chooses a message $m_t^2 \in M$ (these choices are made independently); then they are both told (m_t^1, m_t^2).
- *Action phase:* As in $\Gamma(p)$.

In Section 8 we formally defined (mixed) strategies σ and τ of the two players. Together with the initial choice of the state $\mathbf{k}$ (according to p), a pair of strategies (σ, τ) generates a sequence of pairs of messages and a pair of actions, which results finally in payoffs $\mathbf{a}$ and $\mathbf{b}$ (all these are random variables, on the space Ξ). The expected payoffs of (σ, τ) are $(a, \beta) \equiv (a_{\sigma,\tau}, \beta_{\sigma,\tau})$ where $a^k := E[\mathbf{a} \mid \mathbf{k} = k]$, $a := (a^k)_{k \in K}$, and $\beta := E[\mathbf{b}]$. A pair of strategies (σ, τ) constitutes a (Bayesian–Nash) *equilibrium* in $\Gamma_C(p)$ if

$$a_{\sigma,\tau}^k = \max_{\widetilde{\sigma}} a_{\widetilde{\sigma},\tau}^k, \quad \text{for all } k \in K; \quad \text{and}$$

$$\beta_{\sigma,\tau} = \max_{\widetilde{\tau}} \beta_{\sigma,\widetilde{\tau}}.$$

Denote by $\mathcal{E}_C(p)$ the set of *equilibrium payoff vectors* (a, β) of $\Gamma_C(p)$. Note that $\mathcal{E}(p) \subset \mathcal{E}_C(p)$ for all p (any equilibrium of the silent game is an equilibrium of the talking game—called a "babbling equilibrium"—in which the players ignore all the messages of the talk phase).

A.1.3. *Behavior Strategies*

Our cheap talk game is a game of perfect recall. In such games, there is no need for a player to correlate his choices at different information sets. It suffices to consider only *behavior* strategies, which specify a probability distribution at each decision node; any mixed strategy has an equivalent behavior strategy that yields the same payoffs, no matter what the opponents do, and vice versa (see Kuhn (1953), Aumann (1964), and, specifically for our case where time has order $\omega + 1$, Mertens, Sorin, and Zamir (1994, Exercise II.1.d.8(c), page 76)). In our case, a *behavior*

strategy τ^b for Colin is a sequence of functions $\tau_t^b : H_{t-1} \to \Delta(M)$ for $t = 1, 2, \ldots$, and a measurable function $\tau_\infty^b : H_\infty \to \Delta(J)$; for Rowena, we have $\sigma_t^b : H_{t-1} \times K \to \Delta(M)$ for $t = 1, 2, \ldots$, and $\sigma_\infty^b : H_\infty \times K \to \Delta(I)$.

In the proof (in Section A.4), we will use only the "easy" part of this equivalence: Given a behavior strategy, one constructs an equivalent mixed strategy, by performing all the randomizations (in $\Delta(M)$ and $\Delta(I)$ or $\Delta(J)$) *before* the play starts, rather than at each stage $t = 1, 2, \ldots, \infty$ (see Aumann (1964, Sections 7 and 9)).

A.1.4. *The Diconvex Span*

The *diconvex span* of $\operatorname{gr} \mathcal{E}^+$, denoted $\operatorname{di-span}(\operatorname{gr} \mathcal{E}^+)$, is the set of all triples (a, β, p) in $\mathbb{R}^K \times \mathbb{R} \times \Delta(K)$ for which there exists an $(\mathbb{R}^K \times \mathbb{R} \times \Delta(K))$-valued bounded martingale $(\mathbf{z}_t)_{t=0,1,\ldots} \equiv ((\hat{\mathbf{a}}_t, \boldsymbol{\beta}_t, \mathbf{p}_t))_{t=0,1,\ldots}$ (on some probability space $(Y, \mathcal{F}, \Pi)$) that is adapted to a nondecreasing sequence $(\mathcal{F}_t)_{t=0,1,\ldots}$ of finite fields[38] and satisfies, Π-a.s.,

(A.4) $\qquad \mathbf{z}_0 = (a, \beta, p);$

(A.5) $\qquad \mathbf{z}_\infty \in \operatorname{gr} \mathcal{E}^+; \quad$ and

(A.6) $\qquad \hat{\mathbf{a}}_{t+1} = \hat{\mathbf{a}}_t \quad$ for even t, $\quad$ and $\quad \mathbf{p}_{t+1} = \mathbf{p}_t \quad$ for odd t,

where $\mathbf{z}_\infty \equiv (\hat{\mathbf{a}}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty) := \lim_{t\to\infty} \mathbf{z}_t$ (it exists a.s. by the Martingale Convergence Theorem).

A.2. *Organization of the Proof*

The proof is structured as follows: In the next two sections we fix $p \in \Delta(K)$, and assume that $p^k > 0$ for all $k \in K$. In Section A.3 we prove that every equilibrium of the talking game generates an appropriate dimartingale (Proposition A.8). Then, in Section A.4, we start from any such dimartingale and construct an equilibrium of the talking game (Proposition A.16). Putting this together yields Theorem B. Theorem A now follows,[39] since the equilibrium constructed in Section A.4 is a canonical equilibrium. Finally, if some coordinates of p vanish, we first replace K by the support of p, namely[40] $\{k \in K : p^k > 0\}$.

The proof here parallels the proof of the Main Theorem of Hart (1985).[41] However, a good part of the work there was needed to handle the infinite streams of payoffs, and to obtain uniform approximations; here one needs instead to manage infinite histories.

Let L be a strict bound on all possible payoffs, i.e., $L > |A^k(i, j)|, |B^k(i, j)|$ for all $i \in I, j \in J$ and $k \in K$.

A.3. *From Equilibrium to Martingale*

Let (σ, τ) be an equilibrium of $\Gamma_C(p)$, with payoffs $a \in \mathbb{R}^K$ for player 1 and $\beta \in \mathbb{R}$ for player 2. We will construct an appropriate dimartingale.

It will be convenient to work with the probability space $\Omega := K \times (M \times M)^\infty \times I \times J$, endowed with the usual product σ-field. Ω is the set of all "realizations" $\omega = (k, m_1^1, m_2^1, \ldots, m_t^1, m_t^2,$

---

[38]That is, $\mathbf{z}_t$ is $\mathcal{F}_t$-measurable for each t. By Remark 4.11 in Aumann and Hart (1986), there is no loss of generality in taking the fields to be finite.

[39]One can also prove Theorem A directly—as indicated in Section 5.2—without resorting to dimartingales.

[40]In fact, our arguments show that $\operatorname{gr} \mathcal{E}_C^+ = \operatorname{di-span}(\operatorname{gr} \mathcal{E}^+)$, where $\mathcal{E}_C^+$ is the *modified* equilibrium payoff correspondence for the talking game (obtained from $\mathcal{E}_C$ by arbitrarily increasing payoffs of types with probability 0).

[41]See Aumann and Maschler (1995, Postscripts to Chapter Five, pp. 294–311) for an outline of that proof.

$\dots, i, j)$, consisting of a state of nature k, a sequence of messages (an "infinite history") $(m_t^1, m_t^2)_{t=1,2,\dots} \in (M \times M)^\infty$, and a pair of actions $i \in I$ and $j \in J$. All the random variables— $\mathbf{k}, \mathbf{m}_t^1, \mathbf{m}_t^2, \mathbf{h}_t, \mathbf{h}_\infty, \mathbf{i}, \mathbf{j}$ and so on—are now defined on[42] Ω. Denote by $P \equiv P_{\sigma, \tau, p}$ the probability distribution on Ω generated by σ, τ and p, with $E \equiv E_{\sigma, \tau, p}$ the corresponding expectation operator. Thus, for example, $P[\mathbf{k} = k] = p^k$ and

$$P[\mathbf{m}_t^1 = m \mid \mathbf{h}_{t-1} = h_{t-1}, \mathbf{k} = k] = \lambda(\{\xi_R \in \Xi_R : \sigma_t(\xi_R, h_{t-1}, k) = m\}),$$

where λ is Lebesgue measure. All statements in this proof should be understood to hold P-almost surely.

In this part of the proof we will consider also histories consisting of the first $t+1$ messages of player 1 together with the first t messages of player 2: For each nonnegative integer $s = 0, 1, 2, \dots,$ put

$$g_s := \begin{cases} h_t \equiv (m_1^1, m_1^2, \dots, m_t^1, m_t^2), & \text{if } s = 2t \text{ is } even; \\ (h_t, m_{t+1}^1) \equiv (m_1^1, m_1^2, \dots, m_t^1, m_t^2, m_{t+1}^1), & \text{if } s = 2t+1 \text{ is } odd, \end{cases}$$

and let $\mathbf{g}_s$ be the corresponding random variable (thus, for instance, $\mathbf{g}_{2t} = \mathbf{h}_t$). The set of all such histories g_s is $G_s := M^s$; an infinite history h_∞ will also be denoted g_∞ (and so $G_\infty \equiv H_\infty$). Let $\mathcal{G}_s$ be the finite field on Ω generated by[43] $\mathbf{g}_s$. Then $(\mathcal{G}_s)_s$ is an increasing sequence of fields converging to $\mathcal{G}_\infty \equiv \mathcal{H}_\infty$, the σ-field generated by $\mathbf{g}_\infty \equiv \mathbf{h}_\infty$.

We start with the (posterior) probabilities on the state $\mathbf{k}$. For each $s = 0, 1, \dots, \infty$ (note: including $s = \infty$) and each k in K, let $\mathbf{p}_s^k$ be the conditional probability that the state is k, given a history $\mathbf{g}_s$ of talk:

$$\mathbf{p}_s^k := P[\mathbf{k} = k \mid \mathbf{g}_s]$$

(again, $P = P_{\sigma, \tau, p}$). Put $\mathbf{p}_s := (\mathbf{p}_s^k)_{k \in K} \in \Delta(K)$.

PROPOSITION A.1: *The sequence* $(\mathbf{p}_s)_{s=0,1,\dots}$ *is a bounded martingale adapted to the sequence of finite fields* $(\mathcal{G}_s)_{s=0,1,\dots}$ *satisfying, P-almost surely,*
   (i) $\mathbf{p}_0 = p$;
   (ii) $\mathbf{p}_s \to \mathbf{p}_\infty$ *as* $s \to \infty$; *and*
   (iii) $\mathbf{p}_{s+1} = \mathbf{p}_s$ *for all odd s.*

PROOF: The martingale property is immediate since $(\mathcal{G}_s)$ is an increasing sequence of fields converging to the σ-field $\mathcal{G}_\infty$—which implies (ii). $\mathcal{G}_0$ is the trivial field, hence (i). Finally, the addition from $\mathbf{g}_s$ to $\mathbf{g}_{s+1}$ for odd s, say $s = 2t - 1$, is player 2's message $\mathbf{m}_t^2$, which is independent of $\mathbf{k}$; therefore the conditional probability of $\mathbf{k}$ does not change from s to $s + 1$.   *Q.E.D.*

Next consider the payoff of player 2. For each $s = 0, 1, \dots, \infty$, define

$$\boldsymbol{\beta}_s := E[B^{\mathbf{k}}(\mathbf{i}, \mathbf{j}) \mid \mathbf{g}_s].$$

The following is immediate:

PROPOSITION A.2: *The sequence* $(\boldsymbol{\beta}_s)_{s=0,1,\dots}$ *is a bounded martingale adapted to the sequence of finite fields* $(\mathcal{G}_s)_{s=0,1,\dots}$ *satisfying, P-almost surely,*
   (i) $\boldsymbol{\beta}_0 = \beta$; *and*
   (ii) $\boldsymbol{\beta}_s \to \boldsymbol{\beta}_\infty$ *as* $s \to \infty$.

---

[42]Formally, they are projections to the appropriate coordinates: If $\omega = (k, m_1^1, m_1^2, \dots, i, j) \in \Omega$ then $\mathbf{k}(\omega) = k$; $\mathbf{m}_1^1(\omega) = m_1^1$; $\mathbf{h}_1(\omega) = (m_1^1, m_1^2)$; $\mathbf{i}(\omega) = i$; and so on.

[43]The smallest σ-field on Ω such that $\mathbf{g}_s$ is measurable; it is generated by all sets of the form $\{\omega : \mathbf{g}_s(\omega) = g_s\}$ for some $g_s \in G_s$ (since G_s is finite).

For player 1, the considerations are more intricate. Given the strategy τ of player 2, we define $\mathbf{a}_s^k$, for each k in K and each s in $\mathbb{N}$, as the highest payoff that player 1 can achieve against τ in state k after the history $\mathbf{g}_s$:

$$\mathbf{a}_s^k := \sup_{\widetilde{\sigma}} E_{\widetilde{\sigma}, \tau, p}\big[A^k(\mathbf{i}, \mathbf{j}) \mid \mathbf{g}_s\big],$$

where the supremum is over all strategies $\widetilde{\sigma}$ of player 1 such that $P_{\widetilde{\sigma}, \tau, p}[\mathbf{g}_s \mid \mathbf{k} = k] > 0$. It is well defined for any history g_s in G_s that is not ruled out by τ; in particular, for all g_s with $P[\mathbf{g}_s = g_s] \equiv P_{\sigma, \tau, p}[\mathbf{g}_s = g_s] > 0$. Note that $\mathbf{a}_s^k$ depends only on the continuation of τ after $\mathbf{g}_s$; neither p, nor $\widetilde{\sigma}(\cdot, \cdot, \ell)$ for $\ell \neq k$, nor $\widetilde{\sigma}_t(\cdot, h_t, k)$ for $t < s/2$, matter.

PROPOSITION A.3: *For every k in K, the sequence $(\mathbf{a}_s^k)_{s=0,1,\ldots}$ is a bounded supermartingale adapted to the sequence of finite fields $(\mathcal{G}_s)_{s=0,1,\ldots}$ satisfying, P-almost surely,*
   (i) $\mathbf{a}_0^k = a^k$;
   (ii) $\mathbf{a}_s^k = E[\mathbf{a}_{s+1}^k \mid \mathbf{g}_s]$ *for all odd s; and*
   (iii) $\mathbf{a}_s^k = \max[\mathbf{a}_{s+1}^k \mid \mathbf{g}_s]$ *for all even s.*

Note: The meaning of condition (iii) should be clear: Conditional expectation is replaced by "conditional maximum." That is, for every g_s in G_s, the value of $\mathbf{a}_s^k$ on the atom $\{\mathbf{g}_s = g_s\}$ of $\mathcal{G}_s$ is the maximum of the values of $\mathbf{a}_{s+1}^k$ on this set; i.e., it is the maximum of the values of $\mathbf{a}_{s+1}^k$ on all the atoms of $\mathcal{G}_{s+1}$ that are included in $\{\mathbf{g}_s = g_s\}$ (these atoms are $\{\mathbf{g}_{s+1} = (g_s, m_{t+1}^1)\}$ for all m_{t+1}^1 in M, where $s = 2t$).

PROOF: When $s = 2t - 1$ is odd, the addition from $\mathbf{g}_s \equiv (\mathbf{h}_{t-1}; \mathbf{m}_t^1)$ to $\mathbf{g}_{s+1} \equiv \mathbf{h}_t$ is $\mathbf{m}_t^2$, whose distribution depends only on τ, and is independent of player 1's strategy $\widetilde{\sigma}$. When $s = 2t$ is even, note that, in order to obtain the best-reply payoff $\mathbf{a}_s^k$ given $\mathbf{g}_s \equiv \mathbf{h}_t$, player 1 must choose his next message $\mathbf{m}_{t+1}^1$ so as to maximize his continuation payoff $\mathbf{a}_{s+1}^k$. This proves (ii) and (iii).

Condition (iii) implies $\mathbf{a}_s^k \geq E[\mathbf{a}_{s+1}^k \mid \mathbf{g}_s]$; together with the equality in (ii), it implies that the sequence $(\mathbf{a}_s^k)$ is a supermartingale, bounded by L (which bounds all possible payoffs). Finally, (i) just says that $\sigma(\cdot, \cdot, k)$ is a best reply to τ in state k (recall that $p^k > 0$).   Q.E.D.

From the supermartingale $(\mathbf{a}_s^k)$ we generate a martingale $(\hat{\mathbf{a}}_s^k)$ as follows: Starting with

$$\hat{\mathbf{a}}_0^k := \mathbf{a}_0^k,$$

assume inductively that $\hat{\mathbf{a}}_s^k$ has been defined. Then[44]

$$(A.7) \qquad \hat{\mathbf{a}}_{s+1}^k := \begin{cases} \hat{\mathbf{a}}_s^k, & \text{if } s \text{ is even;} \\ L - \dfrac{L - \hat{\mathbf{a}}_s^k}{L - \mathbf{a}_s^k}(L - \mathbf{a}_{s+1}^k), & \text{if } s \text{ is odd.} \end{cases}$$

We have the following:

PROPOSITION A.4: *For every k in K, the sequence $(\hat{\mathbf{a}}_s^k)_{s=0,1,\ldots}$ is a bounded martingale adapted to the sequence of finite fields $(\mathcal{G}_s)_{s=0,1,\ldots}$ satisfying, P-almost surely,*
   (i) $\hat{\mathbf{a}}_0^k = a^k$;
   (ii) $\mathbf{a}_s^k \leq \hat{\mathbf{a}}_s^k < L$ *for all $s \geq 0$; and*
   (iii) $\hat{\mathbf{a}}_{s+1}^k = \hat{\mathbf{a}}_s^k$ *for all even s.*

---

[44]It appears simpler to define $\hat{\mathbf{a}}_{s+1}^k := \mathbf{a}_{s+1}^k + (\hat{\mathbf{a}}_s^k - \mathbf{a}_s^k)$ when s is odd; but the resulting martingale could be unbounded—which leads to other complications.

PROOF: To prove (ii), assume inductively that it holds for s (starting with $s = 0$, where it is clearly true). For even s, we then have $\hat{\mathbf{a}}_{s+1}^k = \hat{\mathbf{a}}_s^k \geq \mathbf{a}_s^k \geq \mathbf{a}_{s+1}^k$, the last inequality following from (iii) of Proposition A.3. For odd s, note that the definition in (A.7) is equivalent to

$$\hat{\mathbf{a}}_{s+1}^k = \lambda_s^k \mathbf{a}_{s+1}^k + (1 - \lambda_s^k)L,$$

where $\lambda_s^k \in (0, 1]$ is determined by $\hat{\mathbf{a}}_s^k = \lambda_s^k \mathbf{a}_s^k + (1 - \lambda_s^k)L$; from this (ii) follows for $s + 1$. Finally, the martingale property holds for even s by (A.7) and for odd s by Proposition A.3(ii):

$$E[\hat{\mathbf{a}}_{s+1}^k \mid \mathbf{g}_s] = \lambda_s^k E[\mathbf{a}_{s+1}^k \mid \mathbf{g}_s] + (1 - \lambda_s^k)L = \lambda_s^k \mathbf{a}_s^k + (1 - \lambda_s^k)L = \hat{\mathbf{a}}_s^k. \qquad Q.E.D.$$

Now consider the action phase. Let $y(h_\infty) \in \Delta(J)$ be the mixed action of player 2 following the infinite history of talk h_∞:

$$y(h_\infty)(j) := P_{\varXi}(\{\xi_C : \tau_\infty(\xi_C, h_\infty) = j\})$$

for[45] each $j \in J$, and put $y(h_\infty):=(y(h_\infty)(j))_{j \in J}$. Note that the resulting random variable $\mathbf{y} := y(\mathbf{h}_\infty)$ satisfies P-a.s. $\mathbf{y}(j) = P[\mathbf{j} = j \mid \mathbf{h}_\infty]$. Similarly, let

$$x^k(h_\infty)(i) := P_{\varXi}(\{\xi_R : \sigma_\infty(\xi_R, h_\infty, k) = i\})$$

for each $i \in I$ and $k \in K$, and put $x^k(h_\infty) := (x^k(h_\infty)(i))_{i \in I} \in \Delta(I)$—this is the mixed action of player 1 when the state is k. Then $\mathbf{x}^k := x^k(\mathbf{h}_\infty)$ satisfies $\mathbf{x}^k(i) = P[\mathbf{i} = i \mid \mathbf{h}_\infty, \mathbf{k} = k]$ (P-a.s.). We consider player 2 first.

PROPOSITION A.5: *P-almost surely,*

$$\boldsymbol{\beta}_\infty = \max_{y \in \Delta(J)} \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, y) = \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}).$$

PROOF: Recalling the definition of $\boldsymbol{\beta}_\infty$, we have

$$\begin{aligned}
\boldsymbol{\beta}_\infty &= E\big[B^{\mathbf{k}}(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty\big] = \sum_{k \in K} P[\mathbf{k} = k \mid \mathbf{h}_\infty] E\big[B^{\mathbf{k}}(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k} = k\big] \\
&= \sum_{k \in K} \mathbf{p}_\infty^k \sum_{i \in I} \sum_{j \in J} P[\mathbf{i} = i \mid \mathbf{h}_\infty, \mathbf{k} = k] P[\mathbf{j} = j \mid \mathbf{h}_\infty] B^k(i, j) \\
&= \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}).
\end{aligned}$$

It remains to show that, P-almost surely, $\mathbf{y}$ is an optimal response to $(\mathbf{x}^k)_{k \in K}$ in the silent game $\Gamma(\mathbf{p}_\infty)$. Indeed, the event $D(j) := \{\sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}) < \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, j)\}$ is null for each j in J; otherwise, the following strategy $\widetilde{\tau}$ would have increased the payoff of player 2 in $\Gamma_C(p)$: Let $\widetilde{\tau}_\infty(\xi_C, h_\infty) := j$ for all ξ_C when $h_\infty \in D(j)$ and $\widetilde{\tau} = \tau$ otherwise (thus, $\widetilde{\tau}$ is identical with τ except that, in the action phase, j is always chosen after any history of talk h_∞ belonging to $D(j)$). Q.E.D.

We now come to player 1. The sequence $(\hat{\mathbf{a}}_s^k)$ is a bounded martingale; let $\hat{\mathbf{a}}_\infty^k$ be its P-a.s. limit, and put $\hat{\mathbf{a}}_\infty = (\hat{\mathbf{a}}_\infty^k)_{k \in K}$.

PROPOSITION A.6: *P-almost surely, for every k in K,*
  (i) $\hat{\mathbf{a}}_\infty^k \geq \max_{x \in \Delta(I)} A^k(x, \mathbf{y}) \geq A^k(\mathbf{x}^k, \mathbf{y})$; *and*
  (ii) *if $\mathbf{p}_\infty^k > 0$, then $\hat{\mathbf{a}}_\infty^k = \max_{x \in \Delta(I)} A^k(x, \mathbf{y}) = A^k(\mathbf{x}^k, \mathbf{y})$.*

---

[45] We write $P_{\varXi}$ for the probability measure on $\varXi$ (recall Section 8).

PROOF: Consider the following strategy $\tilde{\sigma}$ of player 1: During the talk phase, it is the average nonrevealing strategy derived from σ; in the action phase, it is, for every state k, an optimal reply to player 2's strategy τ. Formally, take $\tilde{\tilde{\Xi}}_R := \Xi_R \times K$ with the uniform distribution on[46] $\tilde{\tilde{\Xi}}_R$ and the distribution p on K (independently), and define[47] $\tilde{\sigma}_t(\tilde{\xi}_R, h_{t-1}, k) \equiv \tilde{\sigma}_t((\xi_R, \ell), h_{t-1}, k) := \sigma_t(\xi_R, h_{t-1}, \ell)$ for all $k, \ell \in K$, $\xi_R \in \Xi_R$, $h_{t-1} \in H_{t-1}$ and $t \geq 1$; and[48] $\tilde{\sigma}_\infty(\tilde{\tilde{\xi}}_R, h_\infty, k) \in \arg\max_{i \in I} A^k(i, y(h_\infty))$ for all $\tilde{\tilde{\xi}}_R \in \tilde{\tilde{\Xi}}_R$, $h_\infty \in H_\infty$ and $k \in K$. Denote $\tilde{\tilde{P}} := P_{\tilde{\sigma}, \tau, p}$ and $\tilde{E} := E_{\tilde{\sigma}, \tau, p}$. The construction of $\tilde{\sigma}$ implies that all histories of talk are realized with the same probability in all states k, and moreover that this probability is identical to the (overall) P-probability; i.e., $\tilde{P}[\mathbf{g}_s = g_s] = \tilde{P}[\mathbf{g}_s = g_s \mid \mathbf{k} = k] = P[\mathbf{g}_s = g_s]$ for all $s = 0, 1, \ldots, \infty$, all $g_s \in G_s$, and all $k \in K$.

Recall that $\hat{\mathbf{a}}_s^k \geq \mathbf{a}_s^k$ (Proposition A.4(ii)), and that $\mathbf{a}_s^k$ was defined as the highest payoff that player 1 can achieve in state k against τ, given $\mathbf{g}_s$. We therefore have

$$\hat{\mathbf{a}}_s^k \geq \mathbf{a}_s^k \geq \tilde{E}\big[A^k(\mathbf{i}, \mathbf{j}) \mid \mathbf{g}_s\big] = \tilde{E}\big[\tilde{E}\big[A^k(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty\big] \mid \mathbf{g}_s\big]$$
$$= E\big[\tilde{E}\big[A^k(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty\big] \mid \mathbf{g}_s\big] = E\big[\max_{x \in \Delta(I)} A^k(x, \mathbf{y}) \mid \mathbf{g}_s\big],$$

the last equality following from the definition of $\tilde{\sigma}_\infty$, and the last but one from the fact that $\tilde{P}$ and P coincide on $\mathcal{H}_\infty$. As $s \to \infty$, the right-most expression converges to $\max_{x \in \Delta(I)} A^k(x, \mathbf{y})$ (since this is $\mathcal{H}_\infty$-measurable). Together with $\hat{\mathbf{a}}_s^k \to \hat{\mathbf{a}}_\infty^k$, it yields (i).

Next, recall Propositions A.1 and A.4: The sequence $((\hat{\mathbf{a}}_s, \mathbf{p}_s))$ is a bimartingale, and therefore $(\mathbf{p}_s \cdot \hat{\mathbf{a}}_s) \equiv (\sum_{k \in K} \mathbf{p}_s^k \hat{\mathbf{a}}_s^k)$ is a martingale (see Hart (1985, Proposition 3.18): from s to $s+1$ either $\hat{\mathbf{a}}$ is constant—when s is even—or $\mathbf{p}$ is constant—when s is odd). Hence its (a.s.) limit $\mathbf{p}_\infty \cdot \hat{\mathbf{a}}_\infty$ satisfies $E[\mathbf{p}_\infty \cdot \hat{\mathbf{a}}_\infty] = E[\mathbf{p}_0 \cdot \hat{\mathbf{a}}_0] = p \cdot a$. But

$$p \cdot a = E\big[A^{\mathbf{k}}(\mathbf{i}, \mathbf{j})\big] = E\big[E\big[A^{\mathbf{k}}(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty\big]\big]$$
$$= E\Big[\sum_{k \in K} P[\mathbf{k} = k \mid \mathbf{h}_\infty] E\big[A^k(\mathbf{i}, \mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k} = k\big]\Big]$$
$$= E\Big[\sum_{k \in K} \mathbf{p}_\infty^k A^k(\mathbf{x}^k, \mathbf{y})\Big].$$

Putting it together yields

$$E\Big[\sum_{k \in K} \mathbf{p}_\infty^k \hat{\mathbf{a}}_\infty^k\Big] = E\Big[\sum_{k \in K} \mathbf{p}_\infty^k A^k(\mathbf{x}^k, \mathbf{y})\Big].$$

Recalling (i), this implies that $\hat{\mathbf{a}}_\infty^k = A^k(\mathbf{x}^k, \mathbf{y})$ whenever $\mathbf{p}_\infty^k > 0$, thus (ii). Q.E.D.

Together, the last two propositions yield the following:

PROPOSITION A.7: *P-almost surely,*

$$(\hat{\mathbf{a}}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty) \in \operatorname{gr} \mathcal{E}^+.$$

---

[46]Recall that Ξ_R is isomorphic to [0, 1].

[47]This is tantamount to player 1 ignoring the k chosen by Nature, selecting $\ell \in K$ according to p (independently of k), and playing in the talk phase as if the state were ℓ.

[48]"arg max" denotes the set of maximizers; to make the choice specific (and $\mathcal{H}_\infty$-measurable), choose always, say, the smallest such i relative to some fixed order of the finite set I.

PROOF: If $\mathbf{p}_\infty^k > 0$ for all k, then Propositions A.5 and A.6(ii) imply that the strategies $(\mathbf{x}^k)_{k \in K}$ for player 1 and $\mathbf{y}$ for player 2 form an equilibrium in $\Gamma(\mathbf{p}_\infty)$, with payoffs $(\hat{\mathbf{a}}_\infty, \boldsymbol{\beta}_\infty)$. If some coordinates of $\mathbf{p}_\infty$ vanish, then for every k in K with $\mathbf{p}_\infty^k = 0$, replace $\mathbf{x}^k$ by a best reply $\widetilde{\mathbf{x}}^k$ to $\mathbf{y}$ (for instance, take $\widetilde{\mathbf{x}}^k := \widetilde{\sigma}_\infty(\widetilde{\xi}_R, \mathbf{h}_\infty, k)$ for the strategy $\widetilde{\sigma}$ constructed in the proof of Proposition A.6). These changes from $\mathbf{x}^k$ to $\widetilde{\mathbf{x}}^k$ will not affect the result of Proposition A.5, since the corresponding coefficients $\mathbf{p}_\infty^k$ vanish. Therefore $\mathbf{y}$ remains a best reply of player 2. As for player 1, his payoff in state k is $A^k(\mathbf{x}^k, \mathbf{y}) = \max_{x \in \Delta(I)} A^k(x, \mathbf{y}) = \hat{\mathbf{a}}_\infty^k$ if $\mathbf{p}_\infty^k > 0$, and $A^k(\widetilde{\mathbf{x}}^k, \mathbf{y}) = \max_{x \in \Delta(I)} A^k(x, \mathbf{y}) \leq \hat{\mathbf{a}}_\infty^k$ if $\mathbf{p}_\infty^k = 0$. Therefore $(\hat{\mathbf{a}}_\infty, \boldsymbol{\beta}_\infty, \mathbf{p}_\infty)$ indeed belongs to $\operatorname{gr} \mathcal{E}^+$.   Q.E.D.

The next proposition summarizes this part of the proof.

PROPOSITION A.8: *Let* $(a, \beta) \in \mathcal{E}_C(p)$, *where* $p^k > 0$ *for all* k *in* K. *Then* $(a, \beta, p) \in$ di-span$(\operatorname{gr} \mathcal{E}^+)$.

PROOF: The martingale $((\hat{\mathbf{a}}_s, \boldsymbol{\beta}_s, \mathbf{p}_s))_{s=0,1,\dots}$ satisfies all requirements (see Section A.1.4) by Propositions A.1–A.7.   Q.E.D.

A.4. *From Martingale to Equilibrium*

We now assume that a point $(a, \beta, p) \in$ di-span$(\operatorname{gr} \mathcal{E}^+)$ is given with $p^k > 0$ for all $k \in K$, and we will construct a *canonical* equilibrium of the talking game $\Gamma_C(p)$. Thus, let $(Y, \mathcal{F}, \Pi)$, $(\mathcal{F}_t)_{t=0,1,\dots}$, $(\mathbf{z}_t)_{t=0,1,\dots}$ and $\mathbf{z}_\infty$ be as in Section A.1.4.

Following Hart (1985, Section 5.1), we modify the process so that, for every t, each atom of the finite field $\mathcal{F}_t$ splits into exactly two atoms of $\mathcal{F}_{t+1}$, with conditional probability $\frac{1}{2}$ each. Thus $\Pi(F_t) = 1/2^t$ for all atoms F_t of $\mathcal{F}_t$. Also, by (A.4) we may take $\mathcal{F}_0$ to be the trivial field $\{\emptyset, Y\}$. Without loss of generality we therefore assume that Y is the space of binary sequences $\{0, 1\}^\infty$; the field $\mathcal{F}_t$ is generated by the first t coordinates (thus: an atom of $\mathcal{F}_t$ corresponds to a sequence of t binary digits); $\mathcal{F}$ is the product σ-field; and Π is the uniform distribution (i.e., the coordinates are independent and each one is 0 or 1 with probabilities $\frac{1}{2} - \frac{1}{2}$).

We start with the limit of the dimartingale, and construct from it mixed choices of the players (to be used in the action phase).

PROPOSITION A.9: *There exist random variables* $(\mathbf{x}^k)_{k \in K} \in [\Delta(I)]^K$ *and* $\mathbf{y} \in \Delta(J)$ *on* Y *satisfying*:[49]

  (i) $\hat{\mathbf{a}}_\infty^k \geq A^k(\mathbf{x}^k, \mathbf{y}) = \max_{x \in \Delta(I)} A^k(x, \mathbf{y})$, *for all* $k \in K$ *(a.s.)*;
  (ii) $\hat{\mathbf{a}}_\infty^k = A^k(\mathbf{x}^k, \mathbf{y})$ *if* $\mathbf{p}_\infty^k > 0$, *for all* $k \in K$ *(a.s.)*;
  (iii) $\boldsymbol{\beta}_\infty = \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}) \geq \max_{y \in \Delta(J)} \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, y)$ *(a.s.); and*
  (iv) $\liminf_{t \to \infty} \hat{\mathbf{a}}_t^k \geq \max_{x \in \Delta(I)} A^k(x, \mathbf{y})$, *for all* $k \in K$.

PROOF: Let Y_1 be the event that $\mathbf{z}_t \to \mathbf{z}_\infty$ and $\mathbf{z}_\infty \in \operatorname{gr} \mathcal{E}^+$; thus $\Pi(Y_1) = 1$. On Y_1, (i)–(iv) are readily obtained from (A.1)–(A.3).[50] To guarantee (iv) on the remaining (null) event $Y_0 := Y \setminus Y_1$, we proceed as follows: Let $IR := \{c \in \mathbb{R}^K : \text{there exists } y \in \Delta(J) \text{ such that } c^k \geq \max_{\widetilde{x} \in \Delta(I)} A^k(\widetilde{x}, y)$ for all $k \in K\}$ be the set of *individually rational* payoff vectors for player 1. Then $\hat{\mathbf{a}}_\infty \in IR$ a.s. (by (A.1) since $\mathbf{z}_\infty \in \operatorname{gr} \mathcal{E}^+$ a.s.). This yields $\hat{\mathbf{a}}_t = E[\hat{\mathbf{a}}_\infty | \mathcal{F}_t] \in IR$ everywhere (since $(\hat{\mathbf{a}}_t)_{t=0,1,\dots}$ is a bounded martingale converging a.s. to $\hat{\mathbf{a}}_\infty$, the set IR is convex, and finally $\hat{\mathbf{a}}_t$ has finitely many values), and therefore $\liminf_{t \to \infty} \hat{\mathbf{a}}_t \in IR$ everywhere (since the set IR is closed). Hence we can choose $\mathbf{y}$ on Y_0 so that (iv) holds there too.   Q.E.D.

---

[49](i)–(iii) hold Π-almost everywhere, whereas (iv) holds *everywhere* (in fact, as we will see in the proof of Proposition A.15 below, it suffices to have (iv) satisfied whenever $\lim_{t \to \infty} \hat{\mathbf{a}}_t^k$ exists).

[50]Using a "measurable selection theorem"—see Hildenbrand (1974, Lemma D.II.2.1) or Klein and Thompson (1984, Theorem 14.2.1).

For the talk phase, let M be the given set of messages. We select two distinct elements of M, and denote them a and b. We start by associating a binary digit δ_t to each pair of messages (m_t^1, m_t^2) at time t as follows:[51]

- For odd $t = 1, 3, 5, \ldots$: put $\delta_t := 1$ if $m_t^1 = $ a, and put $\delta_t := 0$ otherwise;
- For even $t = 2, 4, 6, \ldots$: put $\delta_t := 1$ if either $m_t^1 = m_t^2 = $ a, or $m_t^1 \neq $ a and $m_t^2 \neq $ a, and put $\delta_t := 0$ otherwise.

Thus each history of talk $h_t = (m_r^1, m_r^2)_{r=1}^t \in H_t$—where t may be finite or infinite—is mapped into a binary sequence $d_t = (\delta_r)_{r=1}^t$; denote this mapping by Φ. Note that Φ maps H_∞ onto Y and that, for finite t, the sequence $d_t = \Phi(h_t)$ corresponds to an atom of $\mathcal{F}_t$; we will abuse our notation and denote this atom by d_t or $\Phi(h_t)$.

We can now define the pair of strategies σ and τ for the two players in $\Gamma_C(p)$; it is convenient to define them in terms of their behavior equivalents σ^b and τ^b (recall Section A.1.3):

- In the *talk phase* at *even* period t: Both players choose their messages to be[52] a or b with probabilities of $\frac{1}{2}$ each—regardless of history and, for player 1, also of k. Formally,[53]

$$\sigma_t^b(h_{t-1}, k)(m) := \frac{1}{2}; \quad \text{and}$$

$$\tau_t^b(h_{t-1})(m) := \frac{1}{2};$$

for all even $t = 2, 4, 6, \ldots$, all $h_{t-1} \in H_{t-1}$, all $k \in K$ and $m = $ a, b.

- In the *talk phase* at *odd* period t: Player 2's message is[54] a or b with equal probabilities of $\frac{1}{2}$ each. As for player 1, let $h_{t-1} \in H_{t-1}$ be the history, and let $d_{t-1} := \Phi(h_{t-1})$ be the corresponding atom of $\mathcal{F}_{t-1}$. For each $k \in K$, write p_{t-1}^k for the value of $\mathbf{p}_{t-1}^k$ on $\Phi(h_{t-1})$, and $p_{t,a}^k$, $p_{t,b}^k$ for the values of $\mathbf{p}_t^k$ on the two atoms of $\mathcal{F}_t$ into which $\Phi(h_{t-1})$ splits, namely, $\Phi(h_{t-1}; a, m_t^2)$ and $\Phi(h_{t-1}; b, m_t^2)$, respectively (note that $p_{t,a}^k$ and $p_{t,b}^k$ do not depend on m_t^2 by the definition of Φ). If $p_{t-1}^k > 0$, then messages a and b are sent with probabilities $p_{t,a}^k/(2p_{t-1}^k)$ and $p_{t,b}^k/(2p_{t-1}^k)$, respectively. If $p_{t-1}^k = 0$, then a and b are sent with equal probabilities. Formally,

$$\sigma_t^b(h_{t-1}, k)(m) := \frac{p_{t,m}^k}{2p_{t-1}^k}, \quad \text{if } p_t^k > 0;$$

$$\sigma_t^b(h_{t-1}, k)(m) := \frac{1}{2}, \quad \text{if } p_t^k = 0; \quad \text{and}$$

$$\tau_t^b(h_{t-1})(m) := \frac{1}{2};$$

for all odd $t = 1, 3, 5, \ldots$, all $h_{t-1} \in H_{t-1}$, all $k \in K$ and $m = $ a, b.

- In the *action phase:* Let h_∞ be an infinite history in H_∞, with corresponding $d_\infty := \Phi(h_\infty) \in Y$; put $(x^k)_{k \in K} \in [\Delta(I)]^K$ and $y \in \Delta(J)$ for the value on d_∞ of $(x^k)_{k \in K}$ and $\mathbf{y}$, respectively—as

---

[51]When M contains more than 2 elements, we identify with b all $m \in M$ that are different from a. For odd t, $\delta_t = 1$ corresponds to a and $\delta_t = 0$ to b; for even t, $\delta_t = 1$ corresponds to s and $\delta_t = 0$ to d (cf. Section 5.1).

[52]One may wish, for reasons related to perfectness, to ensure that all messages have positive probability—and so, that all deviations are undetectable. This may be done throughout by splitting the probability of b among all $m \neq $ a.

[53]For each $m \in M$, we write $\sigma_t^b(h_{t-1}, k)(m)$ for the m-component of the probability vector $\sigma_t^b(h_{t-1}, k) \in \Delta(M)$.

[54]Player 2's message here is irrelevant.

given by Proposition A.9. We define

$$\sigma_\infty^b(h_\infty, k) := x^k; \quad \text{and}$$

$$\tau_\infty^b(h_\infty) := y;$$

for all $k \in K$ and all $h_\infty \in H_\infty$.

This completes the definition of the two strategies.

PROPOSITION A.10: (σ, τ) *is a canonical pair of strategies*.

PROOF: By definition (see Section 5.1).   *Q.E.D.*

Recall that $\Omega := (M \times M)^\infty \times I \times J \times K$ denotes the space of "realizations"; let $P = P_{\sigma,\tau,p}$ be the probability distribution induced by σ, τ, and p on Ω.

PROPOSITION A.11: *For all $t \geq 0$, all h_t in H_t and all atoms d_t of $\mathcal{F}_t$,*
  (i) $P[\Phi(\mathbf{h}_t) = d_t] = 1/2^t = \Pi(d_t)$; *and*
  (ii) $P[\mathbf{k} = k \mid \mathbf{h}_t = h_t] = \mathbf{p}_t^k(\Phi(\mathbf{h}_t))$ *for all k in K.*

PROOF: Induction on t: Clearly (i) and (ii) hold for $t = 0$ (recall that $\mathbf{p}_0 = p$ by (A.4)). Assume they hold for $t - 1$.

If t is even, then both players choose their messages $\mathtt{a}$ or $\mathtt{b}$ with equal probabilities $1/2$, so $P[\mathbf{m}_t^1 = \mathbf{m}_t^2] = 1/2$, so[55] $P[\Phi(\mathbf{h}_t)] = (1/2)P[\Phi(\mathbf{h}_{t-1})] = (1/2)(1/2^{t-1}) = 1/2^t$. Also, $\mathbf{m}_t^1$ is independent of $\mathbf{k}$, so $P[\mathbf{k} = k \mid \mathbf{h}_t] = P[\mathbf{k} = k \mid \mathbf{h}_{t-1}] = \mathbf{p}_{t-1}^k(\Phi(\mathbf{h}_{t-1}))$ (by (ii) for $t - 1$), which equals $\mathbf{p}_t^k(\Phi(\mathbf{h}_t))$ by (A.6).

If t is odd, then

$$P[\mathbf{k}, \mathbf{m}_t^1 \mid \mathbf{h}_{t-1}] = P[\mathbf{m}_t^1 \mid \mathbf{h}_{t-1}, \mathbf{k}] \cdot P[\mathbf{k} \mid \mathbf{h}_{t-1}]$$

$$= \sigma_t^b(\mathbf{h}_{t-1}, \mathbf{k})(\mathbf{m}_t^1) \cdot P[\mathbf{k} \mid \mathbf{h}_{t-1}]$$

$$= \frac{\mathbf{p}_t^\mathbf{k}(\Phi(\mathbf{h}_t))}{2\mathbf{p}_{t-1}^\mathbf{k}(\Phi(\mathbf{h}_{t-1}))} \mathbf{p}_{t-1}^\mathbf{k}(\Phi(\mathbf{h}_{t-1})) = \frac{1}{2}\mathbf{p}_t^\mathbf{k}(\Phi(\mathbf{h}_t))$$

(we have used the definition of σ_t^b and (ii) for $t - 1$). Summing over all values $k \in K$ of $\mathbf{k}$ yields $P[\mathbf{m}_t^1 \mid \mathbf{h}_{t-1}] = \frac{1}{2}$ (since $\mathbf{p}_t$ is a probability vector), so $P[\Phi(\mathbf{h}_t)] = (1/2)P[\Phi(\mathbf{h}_{t-1})] = (1/2)(1/2^{t-1}) = 1/2^t$. Moreover,

$$P[\mathbf{k} \mid \mathbf{h}_t] = \frac{P[\mathbf{k}, \mathbf{m}_t^1 \mid \mathbf{h}_{t-1}]}{P[\mathbf{m}_t^1 \mid \mathbf{h}_{t-1}]} = \frac{\frac{1}{2}\mathbf{p}_t^\mathbf{k}(\Phi(\mathbf{h}_t))}{\frac{1}{2}} = \mathbf{p}_t^\mathbf{k}(\Phi(\mathbf{h}_t)),$$

which completes the proof.   *Q.E.D.*

Now $\mathcal{H}_\infty$ and $\mathcal{F}$ are the σ-fields on Ω and Y, respectively, generated by the finite fields $\mathcal{H}_t$ and $\mathcal{F}_t$; so, since $\mathbf{p}_t \to \mathbf{p}_\infty$, Proposition A.11 yields the following:

COROLLARY A.12:
  (i) $P \circ \Phi^{-1} = \Pi$ (*i.e.,* $P[\Phi(\mathbf{h}_\infty) \in W] = \Pi(W)$ *for all $W \in \mathcal{F}$*); *and*
  (ii) $P[\mathbf{k} = k \mid \mathbf{h}_\infty] = \mathbf{p}_\infty^k(\Phi(\mathbf{h}_\infty))$ (*P-a.s.*), *for all k in K.*

---

[55]The equality $P[\Phi(\mathbf{h}_t)] = \frac{1}{2}P[\Phi(\mathbf{h}_{t-1})]$ should be understood to mean $P[\Phi(\mathbf{h}_t) = d_t] = \frac{1}{2}P[\Phi(\mathbf{h}_{t-1}) = d_{t-1}]$ for all d_t and d_{t-1} (with $d_t = (d_{t-1}, \delta_t)$). The same applies to the other expressions here, like $P[\mathbf{k}, \mathbf{m}_t^1 \mid \mathbf{h}_{t-1}]$ and so on.

We come now to the payoffs.

PROPOSITION A.13:
  (i) $E[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{k}=k] = a^k$ for all k in K; and
  (ii) $E[B^{\mathbf{k}}(\mathbf{i},\mathbf{j})] = \beta$.

PROOF: By construction of σ_∞ and τ_∞ we have (a.s.)

$$E\big[B^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k\big] = B^k\big(\mathbf{x}^k(\Phi(\mathbf{h}_\infty)), \mathbf{y}(\Phi(\mathbf{h}_\infty))\big);$$

we will write this as $B^k(\mathbf{x}^k, \mathbf{y}) \circ \Phi$ for short (and similarly for the other expressions below). Corollary A.12(ii) and Proposition A.9(iii) imply:

$$E\big[B^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty\big] = \sum_{k \in K} P[\mathbf{k}=k \mid \mathbf{h}_\infty] E\big[B^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k\big]$$

$$= \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}) \circ \Phi = \boldsymbol{\beta}_\infty \circ \Phi.$$

Taking expectation yields $E[B^{\mathbf{k}}(\mathbf{i},\mathbf{j})] = E[\boldsymbol{\beta}_\infty \circ \Phi] = E_\Pi[\boldsymbol{\beta}_\infty]$ (by Corollary A.12(i); we write E_Π for the expectation, with respect to Π, over the space Y). The sequence $(\boldsymbol{\beta}_t)$ is a bounded martingale converging to $\boldsymbol{\beta}_\infty$, so $E_\Pi[\boldsymbol{\beta}_\infty] = \boldsymbol{\beta}_0 = \beta$, proving (ii).

The argument for (i) is similar, but slightly more complex. By (i) and (ii) of Proposition A.9, together with the construction of the strategies σ_∞ and τ_∞, we have for all $k \in K$ (a.s.)

(A.8) $E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k\big] = A^k(\mathbf{x}^k, \mathbf{y}) \circ \Phi \le \hat{\mathbf{a}}_\infty^k \circ \Phi;$ and

(A.9) $E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k\big] = A^k(\mathbf{x}^k, \mathbf{y}) \circ \Phi = \hat{\mathbf{a}}_\infty^k \circ \Phi$ if $\mathbf{p}_\infty^k \circ \Phi > 0$.

From (A.8) we get (using, again, Corollary A.12(i) and the fact that $(\hat{\mathbf{a}}_t)$ is a bounded martingale converging to $\hat{\mathbf{a}}_\infty$):

$$E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{k}=k\big] = E\big[E[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k]\big]$$

(A.10) $\le E\big[\hat{\mathbf{a}}_\infty^k \circ \Phi\big] = E_\Pi\big[\hat{\mathbf{a}}_\infty^k\big] = \hat{\mathbf{a}}_0^k = a^k.$

Together with (A.9) and Corollary A.12(ii) this yields

$$E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j})\big] = E\bigg[\sum_{k \in K} P[\mathbf{k}=k \mid \mathbf{h}_\infty] E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k}=k\big]\bigg]$$

$$= E\bigg[\sum_{k \in K} \mathbf{p}_\infty^k \hat{\mathbf{a}}_\infty^k \circ \Phi\bigg] = E_\Pi[\mathbf{p}_\infty \cdot \hat{\mathbf{a}}_\infty].$$

The sequence $(\mathbf{p}_t \cdot \hat{\mathbf{a}}_t)$ is a bounded martingale (again, since $(\mathbf{z}_t)$ is a dimartingale) which converges a.s. to $\mathbf{p}_\infty \cdot \hat{\mathbf{a}}_\infty$, so $E_\Pi[\mathbf{p}_\infty \cdot \hat{\mathbf{a}}_\infty] = \mathbf{p}_0 \cdot \hat{\mathbf{a}}_0 = p \cdot a$. So $E[A^{\mathbf{k}}(\mathbf{i},\mathbf{j})] = p \cdot a = \sum_{k \in K} p^k a^k$. But

$$E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j})\big] = \sum_{k \in K} P[\mathbf{k}=k] E\big[A^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{k}=k\big] \le \sum_{k \in K} p^k a^k$$

by (A.10), so there must be equality in (A.10) for each k (since $p^k > 0$ for all k); this proves (i).   *Q.E.D.*

Thus the payoffs induced by the pair of strategies (σ, τ) are indeed (a, β). Next, we show that (σ, τ) is an equilibrium point in $\Gamma_C(p)$. We start with player 2.

PROPOSITION A.14: *The strategy τ is a best reply of player 2 to the strategy σ of player 1 in $\Gamma_C(p)$.*

PROOF: Let $\widetilde{\tau}$ be any strategy of player 2 in $\Gamma_C(p)$, and write $\widetilde{P}$ for the resulting probability measure $P_{\sigma,\widetilde{\tau},p}$ on Ω, and $\widetilde{E}$ for the expectation with respect to $\widetilde{P}$.

First, we claim that the result of Proposition A.11 holds also with $\widetilde{P}$ instead of P. Indeed, the same inductive proof applies: At even t we still have $\widetilde{P}[\mathbf{m}_t^1 = \mathbf{m}_t^2] = \frac{1}{2}$ regardless of the choice of $\mathbf{m}_t^2$, since $\mathbf{m}_t^1$ is uniform on $\{\mathtt{a}, \mathtt{b}\}$, and also $\widetilde{P}[\mathbf{k} \mid \mathbf{h}_t] = \widetilde{P}[\mathbf{k} \mid \mathbf{h}_{t-1}]$ since $\mathbf{m}_t^1$ is independent of k. At odd t, player 2's choice $\mathbf{m}_t^2$ never enters the computations. So Corollary A.12 applies too: $\widetilde{P} \circ \Phi^{-1} = \Pi$ and $\widetilde{P}[\mathbf{k} = k \mid \mathbf{h}_\infty] = \mathbf{p}_\infty^k \circ \Phi$ for all k in K.

Second, let $\widetilde{\mathbf{y}} := \widetilde{\tau}_\infty(\mathbf{h}_\infty) \in \Delta(J)$ be the (mixed) choice of player 2 in the action phase, according to $\widetilde{\tau}$. Then

$$\widetilde{E}\big[B^{\mathbf{k}}(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty\big] = \sum_{k \in K} \widetilde{P}[\mathbf{k} = k \mid \mathbf{h}_\infty]\widetilde{E}\big[B^k(\mathbf{i},\mathbf{j}) \mid \mathbf{h}_\infty, \mathbf{k} = k\big]$$

$$= \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \widetilde{\mathbf{y}}) \circ \Phi$$

$$\leq \sum_{k \in K} \mathbf{p}_\infty^k B^k(\mathbf{x}^k, \mathbf{y}) \circ \Phi = \boldsymbol{\beta}_\infty \circ \Phi,$$

the inequality following from Proposition A.9(iii). Taking expectation yields

$$\widetilde{E}\big[B^{\mathbf{k}}(\mathbf{i},\mathbf{j})\big] \leq \widetilde{E}[\boldsymbol{\beta}_\infty \circ \Phi] = E_\Pi[\boldsymbol{\beta}_\infty] = \boldsymbol{\beta}_0 = \beta$$

(the second equality follows from $\widetilde{P} \circ \Phi^{-1} = \Pi$). Q.E.D.

It remains to show the following:

PROPOSITION A.15: *The strategy σ is a best reply of player 1 to the strategy τ of player 2 in $\Gamma_C(p)$.*

PROOF: Let $\widetilde{\sigma}$ be any strategy of player 1 in $\Gamma_C(p)$, and write $\widetilde{P}$ for the resulting probability measure $P_{\widetilde{\sigma},\tau,p}$ on Ω, and $\widetilde{E}$ for the expectation with respect to $\widetilde{P}$. Through the mapping Φ, this generates a new probability measure[56] $\widetilde{\Pi} := \widetilde{P} \circ \Phi^{-1}$ on Y.

First, we claim that $(\hat{\mathbf{a}}_t)$ is a bounded martingale also with respect to $\widetilde{\Pi}$, i.e., $E_{\widetilde{\Pi}}[\hat{\mathbf{a}}_{t+1} \mid \mathcal{F}_t] = \hat{\mathbf{a}}_t$ for all $t \geq 0$; indeed, at even t we have $\hat{\mathbf{a}}_{t+1} = \hat{\mathbf{a}}_t$ (recall (A.6)), and at odd t the conditional probabilities from $\mathcal{F}_t$ to $\mathcal{F}_{t+1}$ remain the same (since player 2 chooses his message uniformly in $\{\mathtt{a}, \mathtt{b}\}$). Therefore $(\hat{\mathbf{a}}_t)$ converges $\widetilde{\Pi}$-a.s.

Next, recall that player 2's strategy τ_∞ in the action phase was defined so as to satisfy *everywhere* $\liminf_t \hat{\mathbf{a}}_t^k \circ \Phi \geq \max_{x \in \Delta(I)} A^k(x, \mathbf{y}) \circ \Phi$ for all $k \in K$ (see Proposition A.9(iv)[57]). For each k, player 1's payoff in the action phase is thus bounded from above by $\liminf_t \hat{\mathbf{a}}_t^k \circ \Phi$. Taking the expectation of this bound (with respect to the probability $\widetilde{P}$ induced by the talk phase) yields $\widetilde{E}[\liminf_t \hat{\mathbf{a}}_t^k \circ \Phi] = E_{\widetilde{\Pi}}[\liminf_t \hat{\mathbf{a}}_t^k] = \hat{\mathbf{a}}_0^k$ (the first equality since $\widetilde{\Pi} = \widetilde{P} \circ \Phi^{-1}$, and the second since $(\hat{\mathbf{a}}_t)$ is a bounded martingale converging $\widetilde{\Pi}$-a.s.). But $\hat{\mathbf{a}}_0^k = a^k$, so player 1's expected payoff vector from $\widetilde{\sigma}$ against τ is at most a. Q.E.D.

Putting together Propositions A.10 and A.13–A.15, we get the following result.

---

[56]Deviations of player 2 in the talk phase *cannot* affect the probabilities, whereas those of player 1 *can* (due to the odd periods where "signalling" occurs); this is the reason there was no $\widetilde{\Pi}$ in the Proof of Proposition A.14.

[57]The reason for "everywhere" rather than "almost everywhere" should be clear now: It needs to apply for any probability measure $\widetilde{\Pi}$, not just for Π.

PROPOSITION A.16: *Let $(a, \beta, p) \in$ di-span(gr $\mathcal{E}^+$), where $p^k > 0$ for all $k \in K$. Then there exists a canonical equilibrium of $\Gamma_C(p)$ with payoffs (a, β).*

A.5. *Measurability*

PROOF OF PROPOSITION 8.1: Fix a pair (σ, τ) of mixed strategies. Let a, b be numbers. We must prove that the set Ξ_b^a of all sample points ξ leading to the payoff (a, b) is measurable. Consider the diagram

(A.11) $$\Xi \longrightarrow \Xi \times H_\infty \longrightarrow I \times J \times K \longrightarrow \mathbb{R}^2.$$

The first arrow is the talk phase: Applied to a sample point $\xi \equiv (\xi_R, \xi_C, k)$, the strategies σ and τ generate a history h_∞; thus ξ goes to (ξ, h_∞). The second arrow has two components: The first corresponds to the action phase—depending on the history h_∞ and again on the sample point ξ, the strategies σ and τ generate certain actions—and the second to the information phase (it is the projection $(\xi_R, \xi_C, k) \to k$). The actions and the state, in turn, generate payoffs; this is the last arrow.

The set Ξ_b^a is obtained by starting with (a, b) in $\mathbb{R}^2$ and then going backward in the diagram. If all the arrows represent measurable mappings, then the inverse images are all measurable, and it follows that Ξ_b^a is indeed measurable. The second arrow is defined by the action components σ_∞ and τ_∞, which are measurable by definition, and by the projection mapping from Ξ onto K, which is of course measurable. The last arrow poses no problem, as $I \times J \times K$ is finite, so all its subsets are measurable, so all mappings from it are measurable.

It remains to prove that the first arrow is measurable. Consider the infinite diagram:

$$\Xi \xrightarrow{id, \sigma_1, \tau_1} \Xi \times H_1 \xrightarrow{id, \sigma_2, \tau_2} \Xi \times H_2 \xrightarrow{id, \sigma_3, \tau_3} \Xi \times H_3 \longrightarrow \cdots,$$

where *id* stands for the "identity" map. To understand the diagram, start with a ξ in Ξ. Applying the identity yields ξ in Ξ; applying σ_1 yields a message m_1^1 from Rowena to Colin; applying τ_1 yields a message m_1^2 from Colin to Rowena; putting them together yields a 1-stage history h_1 in H_1. Next, again applying the identity yields ξ in Ξ and h_1 in H_1; applying σ_2 yields a message m_2^1 from Rowena to Colin; applying τ_2 yields a message m_2^2 from Colin to Rowena; combining (m_2^1, m_2^2) with the 1-stage history h_1 previously obtained yields a 2-stage history h_2 in H_2. And so on.

Since the σ_t and τ_t are measurable, each of the arrows in the above diagram represents a measurable mapping. So their concatenations—the induced mappings from Ξ to $\Xi \times H_t$—are also measurable. We are, for the moment, interested only in the second components of these mappings—the induced mappings from Ξ to H_t—which we call η_t. Putting all the η_t together[58] yields a mapping η_∞ from Ξ to H_∞. We claim that η_∞, too, is measurable.

To show this, consider a finite rectangle in H_∞. This has the form $S_t \times (M \times M) \times (M \times M) \times \cdots$, where $S_t \subset H_t$ for some finite t. The inverse image of this finite rectangle under η_∞ is the same as the inverse image of S_t under η_t. Since η_t is measurable, so is this inverse image. So all finite rectangles in H_∞ have measurable inverse images. So all sets "generated" by finite rectangles (i.e., sets in the smallest σ-field containing the finite rectangles) have measurable inverse images. But these are precisely the measurable sets. So all measurable sets in H_∞ have measurable inverse images. But that is the meaning of η_∞ being measurable.

We saw above that the two other arrows in Diagram (A.11) are measurable. So their concatenation is also measurable. This completes the proof.   *Q.E.D.*

---

[58]η_{t+1} is obtained from η_t by adding the messages sent in period $t + 1$.

REFERENCES

AMITAI, M. (1996a): "Cheap-Talk with Incomplete Information on Both Sides," DP-90, Center for Rationality, The Hebrew University of Jerusalem.

——— (1996b): "Cheap-Talk with Random Stopping," DP-91, Center for Rationality, The Hebrew University of Jerusalem.

AUMANN, R. J. (1961): "Borel Structures for Function Spaces," *Illinois Journal of Mathematics*, 5, 614–630.

——— (1963): "On Choosing a Function at Random," in *Ergodic Theory*, ed. by F. B. Wright. New York: Academic Press, 1–20.

——— (1964): "Mixed and Behavior Strategies in Infinite Extensive Games," in *Advances in Game Theory* (*Annals of Mathematics Study* 52), ed. by M. Dresher et al. Princeton: Princeton University Press, 627–650.

——— (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

——— (1990): "Nash Equilibria are Not Self-Enforcing," in *Economic Decision-Making: Games, Econometrics and Optimisation*, ed. by J. J. Gabszewicz, J.-F. Richard, and L. A. Wolsey. Amsterdam: Elsevier, 201–206.

——— (2001): "The Rationale for Measurability," in *Economics Essays, A Festschrift for Werner Hildenbrand*, ed. by G. Debreu, W. Neuefeind, and W. Trockel. Berlin: Springer, 5–7.

AUMANN, R. J., AND S. HART (1986): "Bi-Convexity and Bi-Martingales," *Israel Journal of Mathematics*, 54, 159–180.

AUMANN, R. J., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.

AUMANN, R. J., M. MASCHLER, AND R. E. STEARNS (1968): "Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case," Mathematica, Inc. Report to the U.S. Arms Control and Disarmament Agency ST-140, Princeton, NJ. Published as Chapter V in Aumann and Maschler (1995).

BALIGA, S., AND S. MORRIS (2002): "Co-ordination, Spillovers and Cheap Talk," *Journal of Economic Theory*, 105, 450–468.

BALIGA, S., AND T. SJÖSTRÖM (2001): "Arms Races and Negotiations," *Review of Economic Studies* (forthcoming).

BANERJEE, A., AND J. WEIBULL (2000): "Neutrally Stable Outcomes in Cheap-Talk Coordination Games," *Games and Economic Behavior*, 32, 1–24.

BARANY, I. (1992): "Fair Distribution Protocols or How Players Replace Fortune," *Mathematics of Operations Research*, 17, 327–340.

BATTAGLINI, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70, 1379–1401.

BEN-OR, M., S. GOLDWASSER, AND A. WIGDERSON (1988): "Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract)," *Proceedings of the 20th ACM Symposium on the Theory of Computing* (*STOC*), ACM, 1–10.

BEN-PORATH, E. (1998): "Correlation without Mediation: Expanding the Set of Equilibrium Outcomes by Cheap Pre-Play Procedures," *Journal of Economic Theory*, 80, 108–122.

——— (2002): "Cheap Talk in Games with Incomplete Information," *Journal of Economic Theory* (forthcoming).

BLUME, A., Y.-G. KIM, AND J. SOBEL (1993): "Evolutionary Stability in Games of Communication," *Games and Economic Behavior*, 5, 547–575.

BLUME, A., AND J. SOBEL (1995): "Communication-Proof Equilibria in Cheap-Talk Games," *Journal of Economic Theory*, 65, 359–382.

CANETTI, R. (2000): "Security and Composition of Multiparty Cryptographic Protocols," *Journal of Cryptology*, 13, 143–202.

CRAWFORD, V. P., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

FARRELL, J. (1993): "Meaning and Credibility in Cheap-Talk Games," *Games and Economic Behavior*, 5, 514–531.

FARRELL, J., AND M. RABIN (1996): "Cheap Talk," *Journal of Economic Perspectives*, 10, 103–118.

FORGES, F. (1990a): "Equilibria with Communication in a Job Market Example," *Quarterly Journal of Economics*, 105, 375–398.

——— (1990b): "Universal Mechanisms," *Econometrica*, 58, 1341–1364.

——— (1992): "Repeated Games of Incomplete Information: Non-Zero-Sum," in *Handbook of Game Theory, with Economic Applications*, Vol. 1, ed. by R. J. Aumann and S. Hart. Amsterdam: North-Holland, 155–177.

GERARDI, D. (2000): "Unmediated Communication in Games with Complete and Incomplete Information," Mimeo, Northwestern University.

GOLDREICH, O., S. MICALI, AND A. WIGDERSON (1987): "How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority (Extended Abstract)," *Proceedings of the 19th ACM Symposium on the Theory of Computing* (*STOC*), ACM, 218–229.

GOLDWASSER, S., AND L. LEVIN (1991): "Fair Computation of General Functions in Presence of Immoral Majority," in *Advances in Cryptology – CRYPTO '90*, ed. by A. J. Menezes and S. A. Vanstone. Berlin: Springer, 77–93.

GREEN, J. R., AND N. L. STOKEY (1980): "A Two-Person Game of Information Transmission," Mimeo, Harvard University and Northwestern University.

HART, S. (1985): "Nonzero-Sum Two-Person Repeated Games with Incomplete Information," *Mathematics of Operations Research*, 10, 117–153.

HILDENBRAND, W. (1974): *Core and Equilibria of a Large Economy*. Princeton, NJ: Princeton University Press.

KIM Y.-G., AND J. SOBEL (1995): "An Evolutionary Approach to Pre-Play Communication," *Econometrica*, 63, 1181–1193.

KLEIN, E., AND A. C. THOMPSON (1984): *Theory of Correspondences*. New York: Wiley.

KREPS, D. M., AND J. SOBEL (1994): "Signalling," in *Handbook of Game Theory, with Economic Applications*, Vol. 2, ed. by R. J. Aumann and S. Hart. Amsterdam, North-Holland, 849–867.

KRISHNA, V., AND J. MORGAN (2002): "The Art of Conversation," Mimeo.

KUHN, H. W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games II* (*Annals of Mathematics Study* 28), ed. by H. W. Kuhn and A. W. Tucker. Princeton, NJ: Princeton University Press, 193–216.

LEHRER, E., AND S. SORIN (1997): "One-Shot Public Mediated Talk," *Games and Economic Behavior*, 20, 131–149.

MATTHEWS, S., M. OKUNO-FUJIWARA, AND A. POSTLEWAITE (1991): "Refining Cheap-Talk Equilibria," *Journal of Economic Theory*, 55, 247–273.

MERTENS, J.-F., S. SORIN, AND S. ZAMIR (1994): "Repeated Games," DP 9420, 9421, 9422, CORE, Université Catholique de Louvain.

MYERSON, R. B. (1994): "Communication, Correlated Equilibria, and Incentive Compatibility," in *Handbook of Game Theory, with Economic Applications*, Vol. 2, ed. by R. J. Aumann and S. Hart. Amsterdam: North-Holland, 827–847.

NASH, J. (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286–295.

RABIN, M. (1990): "Communication between Rational Agents," *Journal of Economic Theory*, 51, 144–170.

RABIN, M., AND J. SOBEL (1996): "Deviations, Dynamics, and Equilibrium Refinements," *Journal of Economic Theory*, 68, 1–25.

RABIN, T., AND M. BEN-OR (1989): "Verifiable Secret Sharing and Multiparty Protocols with Honest Majority (Extended Abstract)," *Proceedings of the 21st ACM Symposium on the Theory of Computing* (*STOC*), ACM, 73–85.

ROBSON, A. (1990): "Efficiency in Evolutionary Games: Darwin, Nash and the Secret Handshake," *Journal of Theoretical Biology*, 144, 379–396.

SIMON, R. S. (2002): "Separation of Joint Plan Equilibrium Payoffs from the Min-Max Functions," *Games and Economic Behavior*, 41, 79–102.

SPENCE, A. M. (1974): *Market Signaling*. Cambridge, MA: Harvard University Press.

URBANO, A., AND J. VILA (2002): "Computational Complexity and Communication: Coordination in Two-Player Games," *Econometrica*, 70, 1893–1927.

WÄRNERYD, K. (1993): "Cheap Talk, Coordination, and Evolutionary Stability," *Games and Economic Behavior*, 5, 532–546.

YAO, A. (1986): "How to Generate and Exchange Secrets," *Proceedings of the 27th Symposium on the Foundations of Computer Science* (*FOCS*), IEEE, 162–167.

ZAHAVI, A. (1975): "Mate Selection: A Selection for a Handicap," *Journal of Theoretical Biology*, 53, 205–214.

ZAPATER, I. (1997): "Credible Proposals in Communication Games," *Journal of Economic Theory*, 72, 173–197.