

Extensive Form Correlated Equilibrium: Definition and Computational Complexity

Bernhard von Stengel

Department of Mathematics, London School of Economics,
Houghton St, London WC2A 2AE, United Kingdom
email: stengel@maths.lse.ac.uk

Françoise Forges

CEREMADE, University of Paris – Dauphine,
Place du Marechal de Lattre de Tassigny, 75775 Paris cedex 16, France
email: Francoise.Forges@dauphine.fr

March 18, 2006

Abstract: This paper defines the *extensive form correlated equilibrium* (EFCE) for extensive games with perfect recall. The EFCE concept extends Aumann's strategic-form correlated equilibrium. Before the game starts, a correlation device generates a move for each information set. This move is recommended to the player only at the time of reaching the information set. The condition of perfect recall in two-player extensive games without chance moves leads to strong restrictions on the players' information sets, which are of some interest on their own. These are used to characterize the set of EFCE by means of a polynomial number of consistency and incentive constraints for correlating sequences of moves. In contrast, strategic-form correlated equilibria for two-player games without chance moves give rise to NP-hard optimization problems. Similarly, maximizing the pay-off of an EFCE, or of a strategic-form correlated equilibrium, is NP-hard for two-player games with chance moves.

Keywords: Correlated equilibrium, extensive game, polynomial-time computable.

JEL classification: C72, C63.

AMS subject classification: 91A18, 91A05, 91A28, 68Q17.

Contents

1	Introduction	3
2	The EFCE concept	5
2.1	Definition of EFCE	6
2.2	Reduced strategies suffice	7
2.3	Example: A signaling game	8
2.4	Relationship to other solution concepts	9
2.5	Discussion and open problems	11
3	Computational complexity	12
3.1	Review of the sequence form	13
3.2	Correlation plans and marginal probabilities	14
3.3	Example of generating move recommendations	16
3.4	Information structure of two-player games without chance moves	18
3.5	Using the consistency constraints	21
3.6	Incentive constraints	25
3.7	Hardness results	27
	Acknowledgments	30
	References	31

1 Introduction

Aumann (1974) defined the concept of *correlated equilibrium* for games in strategic form. Before the game starts, a device selects private signals from a joint probability distribution and sends them to the players. In the “canonical” interpretation, these signals are strategies that players are recommended to play.

The strategic-form correlated equilibrium can be applied to a game in extensive form. This approach is analogous to expressing a Nash equilibrium as a profile of mixed strategies, which are randomizations on the sets of pure strategies. We assume each player has perfect recall. Then, by Kuhn’s theorem (1953), a mixed strategy can be replaced by a behavior strategy, which defines a randomization over the moves at each information set of the player. A behavior strategy is much less complex than a mixed strategy because it can be specified by a probability for every move, whereas a mixed strategy requires typically an exponential number of probabilities, one for each pure strategy.

This paper proposes a new concept of correlated equilibrium for extensive games, called *extensive form correlated equilibrium* or EFCE. In the same way as in a Nash equilibrium mixed strategies are replaced by behavior strategies, in an EFCE recommendations of pure strategies are replaced by recommendations of *moves* at information sets. Like in a strategic-form correlated equilibrium, the recommendations to the players are generated before the game starts. However, the recommended move is not revealed to the player before the respective information set is reached. As recommendations become local in this way, players know less. Consequently, the set of EFCE outcomes is *larger* than the set of strategic-form correlated equilibrium outcomes.

The EFCE is a natural definition of correlated equilibrium for extensive games with perfect recall. It applies to any extensive game with information sets as defined by Kuhn (1953), including games without a well-defined time when each player moves. Earlier extensions of Aumann’s concept applied only to multi-stage games (including Bayesian games and stochastic games) that have a special time and information structure. These known approaches are discussed in Section 2.4 below.

For games in strategic form, correlated equilibria are *easier to compute* than Nash equilibria. The incentive constraints that define the set of correlated equilibria of a game are linear inequalities in terms of the joint probabilities over strategy profiles. An incentive constraint compares any two strategies of a player, so the number of these constraints is polynomial in the size of the strategic form. Consequently, finding a correlated equilibrium with, say, maximum payoff sum amounts to solving a linear program, so this can be done in polynomial time. In contrast, finding a Nash equilibrium with maximum payoff sum defines an NP-hard optimization problem (Gilboa and Zemel (1989), Conitzer and Sandholm (2003); see Garey and Johnson (1979) or Papadimitriou (1994) for notions of computational complexity).

Linear programming can also be applied to computing Nash equilibria when the strategic-form game has two players and zero-sum payoffs. When the game is given in extensive form, its strategic form has an exponentially larger amount of data than the game tree, so

the game cannot be solved in polynomial time using the strategic form. The strategic form is therefore computationally intractable for larger extensive games.

However, extensive two-player zero-sum games with perfect recall can still be solved in polynomial time in the tree size, as shown by Romanovskii (1962), Koller and Megiddo (1992), and von Stengel (1996). These methods use the *sequence form* of an extensive game which represents a behavior strategy by its realization probabilities for sequences of moves along a path in the game tree. These realization probabilities can be characterized by linear equations, one for each information set. Thereby, the sequence form provides a strategic description that has the same size as the game tree, unlike the exponentially large strategic form.

Given an extensive game with perfect recall, is there a “sequence form” to compute a *strategic-form* correlated equilibrium with, say, maximum payoff sum in polynomial time? The answer is *negative* because this problem is NP-hard (Chu and Halpern (2001); see also Theorem 3.10 below). The set of correlated equilibria can therefore not be characterized by a polynomial number of inequalities in the size of the game tree, assuming that $P \neq NP$. Chu and Halpern’s construction applies to extensive games of two players with an initial chance move. In Theorem 3.11, we show that even for an extensive two-player game *without* chance moves, a strategic-form correlated equilibrium with maximum payoff sum is hard to compute.

In contrast, *extensive form* correlated equilibria are computationally tractable for two-player games without chance moves. We show that a polynomial number of linear constraints suffice to characterize EFCE for these games. These constraints extend the sequence form constraints as used for Nash equilibria. They define joint probabilities for correlating moves at any two information sets of the two players by means of suitable *consistency* and *incentive* conditions. From these probabilities, the recommended moves are generated for one information set at a time, taking earlier recommendations into account. This specifies the correlation device compactly, without explicitly using probabilities for strategy profiles.

The polynomial-time computability of EFCE for two-player games without chance moves is *not* straightforward. The consistency constraints on the marginal probabilities of moves that are correlated across information sets are in general only necessary conditions. In extensive games with two players and without chance moves, these constraints are also sufficient to describe the set of EFCE. For such games, the condition of perfect recall imposes strong restrictions on the players’ information sets (see Section 3.4), for example a unique partial “time order” among them. These properties may be of some interest by themselves. The most important consequence is that the generation of a recommended move at a player’s information set can be based on a unique earlier *sequence* of moves by the other player (see Lemma 3.3(c) and (9) in the proof of Theorem 3.8). Because the number of sequences, as opposed to strategies, is polynomial, this reduces the computational complexity.

The EFCE concept can be regarded as the correlated analog of a Nash equilibrium in behavior strategies. It is closer in spirit to the dynamic description of the game by a tree than the strategic-form correlated equilibrium. At the same time, the correlation de-

vice does not have additional power in the sense of observing the game state, because it generates signals at the beginning of the game. In addition, the EFCE concept is computationally tractable, at least for two-player games without chance moves. The EFCE also seems to be the first case of a game-theoretic concept where the introduction of chance moves marks the transition from polynomial-time solvable to NP-hard problems.

Interestingly, Papadimitriou (2005) describes how to find a correlated equilibrium for a compactly represented multi-player game in polynomial time even though finding a correlated equilibrium with maximum payoff sum is NP-hard. That is, maximizing the payoff sum is a stronger computational requirement than merely finding one correlated equilibrium. The compactly represented games considered by Papadimitriou encompass a wide class of games, but not games in extensive form. It is open whether his approach can be applied to the EFCE concept for extensive games with chance moves or with more than two players.

The two main parts of this paper treat the conceptual and the computational aspects of EFCE. In the first part, we first define the EFCE concept. Then we observe that the EFCE can be defined in “canonical form” and by generating only reduced strategy profiles. A signaling game shows that an EFCE can be “type-revealing” when this is not possible with a Nash or strategic-form correlated equilibrium. We then compare the EFCE concepts with related other notions of correlated equilibria. In conclusion, we mention open conceptual problems. Most of our mathematical observations are straightforward, so that we do not state them as theorems with proofs.

The second part of this paper is concerned with computational complexity, and much more technical. Our positive result, a compact description of the set of EFCE, holds for two-player extensive games with perfect recall and without chance moves. We have to prove carefully how the “consistency constraints” describe probabilities for generating moves, because the games we consider have no clear “stages” (see, for example, Figure 6). On the other hand, the consistency constraints provide sufficient structure for generating move recommendations in an unambiguous way. The computational difficulties of more general games with chance moves, and of strategic-form correlated equilibria, are considered at the end of this paper.

2 The EFCE concept

The EFCE concept is defined in Section 2.1. As shown in Section 2.2, an EFCE can be defined with a correlation device that generates reduced strategy profiles. A game with costless signals illustrates the use of EFCE, as explained in Section 2.3. In Section 2.4, we compare the EFCE with other concepts of correlated equilibria that have been defined for games with special time or information structures. Open problems that arise from a conceptual viewpoint are discussed in Section 2.5.

2.1 Definition of EFCE

We use the following standard terminology for extensive games. Let N be the finite set of players. The *game tree* is a finite directed tree, that is, a directed graph with a distinguished node, the *root*, from which there is a unique path to any other node. The non-terminal *decision* nodes of the game tree are partitioned into *information sets*. Each information set belongs to exactly one player i . The set of all information sets of player i is denoted H_i . The set of choices or *moves* at an information set h is denoted C_h . Each node in h has $|C_h|$ outgoing edges, which are labeled with the moves in C_h .

We assume each player has *perfect recall*, defined as follows. Without loss of generality, choice sets C_h and C_k for $h \neq k$ are considered disjoint. A *sequence* of moves of a particular player is a sequence of his moves (ignoring the moves of the other players) along the path from the root to some node in the game tree. By definition, player i has perfect recall if all nodes in an information set h in H_i define the same sequence σ_h of moves for player i .

The set of *pure strategies* of player i is

$$\Sigma_i = \prod_{h \in H_i} C_h. \quad (1)$$

The set of all *strategy profiles* is

$$\Sigma = \prod_{i \in N} \Sigma_i. \quad (2)$$

Definition 2.1. A (canonical) *correlation device* is a probability distribution μ on Σ .

A correlation device μ makes recommendations to the players by picking a strategy profile π according to the distribution μ , and privately recommending the component π_i of π to each player i for play. It defines a *strategic-form correlated equilibrium* if no player can gain by unilaterally deviating from the recommended strategy, given his posterior on the recommendations to the other players (see Aumann (1974)). We define an extensive form correlated equilibrium also by means of a correlation device, but with a different way of giving recommendations to the players.

Definition 2.2. Given a correlation device μ as in Definition 2.1, consider the extended game in which a chance move first selects a strategy profile π according to μ . Then, whenever a player i reaches an information set h in H_i , he receives the move c at h specified in π as a signal, interpreted as a recommendation to play c . An *extensive form correlated equilibrium (EFCE)* is a Nash equilibrium of such an extended game in which the players follow their recommendations.

In an EFCE, the strategy profile selected according to the device defines a move c for each information set h of each player i , which is revealed to player i only when he reaches h . It is optimal for the player to follow the recommended move, assuming that all future moves, including of the player himself, are followed as recommended. However, when a player considers a deviation from a recommended move, he is not obliged to

follow subsequent recommendations. This distinguishes the EFCE from the *agent normal form correlated equilibrium*, which is a correlated equilibrium where at every information set, the move is chosen by a different agent (see also Section 2.4).

The above description of extensive form correlated equilibria is in “canonical form”. That is, the recommendations to players are moves to be made at information sets and not arbitrary signals. In the same way as for strategic-form correlated equilibria, this can be assumed without loss of generality (see Forges (1986a)).

2.2 Reduced strategies suffice

In the *reduced strategic form* of an extensive game, strategies of a player that differ in moves at information sets which are unreachable due to an own earlier move are identified.¹ A reduced strategy can still be considered as a tuple of moves, except that the unspecified move at any unreachable information set is denoted by a new symbol, for example a star “*”, which does not belong to any set of moves C_h .

We denote the set of all reduced strategies of player i by Σ_i^* , and the set of all reduced strategy profiles by

$$\Sigma^* = \prod_{i \in N} \Sigma_i^*. \quad (3)$$

By construction, the payoffs for a profile of reduced strategies are uniquely given as in the strategic form. This defines the *reduced strategic form* of the extensive game.

In Definition 2.1, a correlation device is defined on Σ , that is, using the unreduced strategic form. We now re-define a correlation device to be a probability distribution on Σ^* . Any correlated equilibrium that is specified using the unreduced strategic form can be considered as a correlated equilibrium for the reduced strategic form. This is achieved by defining the probability for a profile π^* of reduced strategies as the sum of the probabilities of the unreduced strategy profiles π that *agree* with π^* (in the sense that whenever π^* specifies a move other than “*” at an information set, then π specifies the same move). Because the incentive constraints hold for the unreduced strategies, and payoffs are identical, appropriate sums of these give rise to the incentive constraints for the reduced strategies, which therefore hold as well.

Conversely, any correlated equilibrium for the reduced strategic form can be applied to the unreduced strategic form by arbitrarily defining a move for every unreachable information set (which is “*”, that is, undefined, in the reduced strategy profile), thereby defining a particular unreduced strategy to be selected by the correlation device.

In the same manner, an EFCE can be defined by assigning probabilities only to reduced strategy profiles. This defines an EFCE for unreduced strategy profiles by recommending an arbitrary move at each unreachable information set. Conversely, consider an EFCE defined using unreduced strategy profiles as in Definition 2.2. Then, just as in the strategic form, this gives rise to an EFCE for reduced profiles, as follows. In the strategy

¹We define the reduced strategic form in this way because it only depends on the game tree structure and not on the payoffs.

profile π generated by the correlation device, any recommendation at an unreachable information set is replaced by “*”. Suppose a player deviates from his recommended move at some information set, and gets a higher payoff by subsequently using moves at previously unreachable information sets where he only gets the recommendation “*”. Then the player could profitably deviate in the same way when getting recommendations of moves for these information sets as in π , which he ignores. This contradicts the assumed equilibrium property.

2.3 Example: A signaling game

Figure 1 shows an example of an extensive game. This is a signaling game as discussed by Spence (1973), Cho and Kreps (1987), and Gibbons (1992, Section 4.2), but with costless signals. Player 1, a student, can be with equal probability of a good (G) or bad (B) type. He applies for a summer research job with a professor, player 2. Player 1 sends a costless signal X or Y . The professor can distinguish the signals but not the type of player 1, as shown by her two information sets. She can either let the student work with her (l) or refuse to do so (r). Move r always gives the pair of payoffs $(0, 3)$ to players 1 and 2, but l results in $(2, 5)$ for G versus $(3, 0)$ for B .

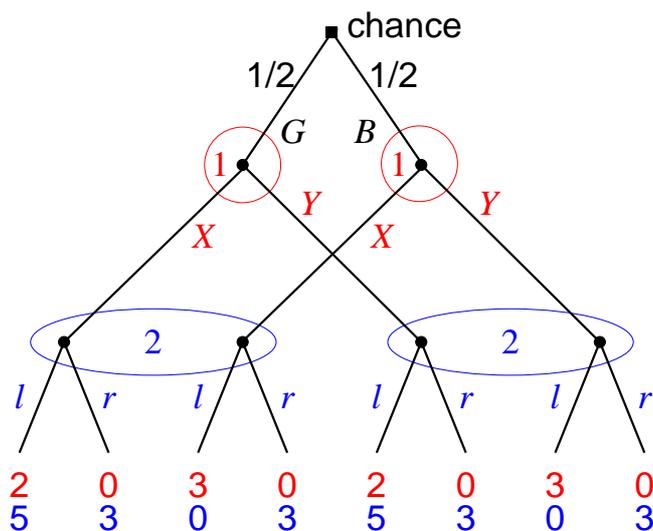


FIGURE 1

In any Nash equilibrium of this game, player 2 refuses to work with the student, choosing r at both information sets. Otherwise, any positive probability for l would induce both types of player 1 to send a (not necessarily unique) signal where that probability of acceptance is the highest; hence, for at least one signal, l is not optimal for player 2 because the bad type is at least as likely as the good type. With player 2 always playing r , the signal sent by player 1 does not matter (he gets payoff 0 anyhow), as long as in no information set of player 2, the conditional probability for G versus B is so high as to make her switch to l .

Similarly, in every strategic-form correlated equilibrium, player 2 is never recommended to choose move l , because otherwise the bad type would “imitate” the good type by sending the signal where player 2 would choose l . So in this game, the sets of Nash and correlated equilibrium outcomes coincide.

However, there is an EFCE with better payoff to both players compared to the outcome with payoff pair $(0, 3)$: A signal X or Y is chosen with equal probability for type G , and player 2 is told to accept (move l) when receiving the chosen signal and to refuse (move r) when receiving the other signal. The bad type B is given an arbitrary recommendation which is independent of the recommendation to type G . Because the move recommended to G is *unknown* to B , the bad type cannot distinguish the two signals and, no matter what he does, will match the signal of G with probability $1/2$. When player 2 receives the signal chosen for G , it is therefore twice as likely to come from G rather than from B , so that her expected payoff $10/3$ for choosing l is higher than 3 when she chose r . When she receives the wrong signal, it comes from B with certainty, and then the best reply is certainly r with payoff 3. The expected payoffs to the two players in this EFCE are 1.75 to player 1 and 3.25 to player 2. In a more elaborate game with M signals instead of just two signals, where the bad type can only guess the correct signal with probability $1/M$, the pair of expected payoffs is $(1 + 1.5/M, 4 - 1.5/M)$.

In the terminology of signaling games, any Nash or correlated equilibrium is the described “pooling equilibrium” with payoff pair $(0, 3)$. This is due to the fact that signals are costless and therefore uninformative. In contrast, the EFCE concept allows for a “partially revealing” equilibrium, where signals can distinguish the types, which has better payoffs for both players.

2.4 Relationship to other solution concepts

Our definition of an EFCE generalizes the Nash equilibrium in behavior strategies and applies to any game in extensive form (with perfect recall). Other extensions of Aumann’s strategic-form correlated equilibrium have been proposed in order to take account of the dynamic structure of *specific* classes of games, namely Bayesian games and multi-stage games.

In a Bayesian game, every player has a type which can be represented by an information set. Players move only once and simultaneously. For Bayesian games, the *agent normal form correlated* equilibrium studied in, e.g., Forges (1986b), Samuelson and Zhang (1989), Cotter (1991), and Forges (1993), exactly coincides with the EFCE.

However, in general extensive form games, the set of agent normal form correlated equilibrium outcomes can be larger than the set of EFCE outcomes. An easy example is a one-player game where the player moves twice, first choosing either “Out” and receiving zero, or “In” and then choosing again between “Out” with payoff zero or “In” with payoff one. If the two agents at the two decision points both choose “Out”, this defines an agent normal form correlated equilibrium, but not an EFCE.

In multi-stage games, the best known extension of the strategic-form correlated equilibrium is the *communication equilibrium* introduced by Myerson (1986) and Forges

(1986a). The underlying canonical scenario is that at every stage of the game, every player is invited to report his new information to a communication device (with perfect memory), which in turn makes private recommendations to the players. In a canonical equilibrium, players are truthful and obedient at every stage. This solution concept differs from the EFCE in two respects: First, the players can send inputs to the device. Second, the outputs can depend on the players' inputs but not on their true information, which is unknown to the device. The device is only informed about the stage (and remembers what players have told it) but cannot distinguish between different information sets of a player at the same stage when generating its outputs. In contrast, the recommendations given in an EFCE are local even at the same stage, as the example in Figure 1 shows. In that example, any communication equilibrium gives only the payoff pair $(0, 3)$. That is, the set of communication equilibrium outcomes in that example is strictly included in the set of EFCE outcomes. The reverse inclusion can also hold. For example, consider a game (described by Forges (1986a), p. 1383) where player 1 learns a move of nature but has only one move himself (like only one signal, say X , in Figure 1). Player 2 does not know the move of nature but would profit from doing so, along with player 1 who has the same payoffs. In a communication equilibrium, player 1 could inform player 2 to their joint benefit, but not in an EFCE.

Like the communication equilibrium, the *autonomous correlated equilibrium* (Forges (1986a)) applies to multistage games; the players receive outputs at every stage, but cannot make any inputs to the device. In the canonical version of the solution concept, the output to every player at every stage is a mapping telling him which move to choose at that stage as a function of his information (i.e., the relevant part of his strategy for the given stage). However, unlike in an EFCE, the respective signal is known to the player for the entire stage and not only locally for each information set.² Obviously, every strategic-form correlated equilibrium outcome is an autonomous correlated equilibrium outcome, but the converse is not true in general, like in a battle of sexes game preceded by a suitable outside option (see Myerson (1986, Fig. 2)). Similarly, the set of autonomous correlated equilibrium outcomes is included in the set of EFCE outcomes, and the inclusion may be strict, as shown in the example of the previous section. The same holds in the class of games we consider later, namely two-player games without chance moves (see Section 3.3).

Solan (2001) defines a concept of communication equilibrium for *stochastic games* where the device knows the game state and all past moves, which are also known to all players. He proves that this concept is outcome equivalent to the autonomous correlated equilibrium. Therefore, in stochastic games, or other games where the players have symmetric information, the equilibrium outcomes for these concepts and the EFCE coincide.

Kamien, Tauman, and Zamir (1990) and Zamir, Kamien, and Tauman (1990) study extensive games with a single initial chance move. The game is modified by introducing a disinterested additional player (the “maven”) who can reveal any partial information about the chance move to each player. In some games, the resulting set of payoffs has

²In Forges (1986a, p. 1378), a correlated equilibrium based on an autonomous device is called “extensive form correlated equilibrium”, but this is now typically referred to as “autonomous correlated equilibrium”. We suggest now to use “EFCE” in our sense.

some similarity with that obtainable in an EFCE. However, the correlation device used in an EFCE is weaker than such a maven, for the following reasons: Recommendations are generated at the beginning of the game. The device does not observe play, and “knows” the game state only implicitly under the assumption that players observe their recommended moves. The device cannot make recommendations conditional on game states that have been determined by a chance move.

Moulin and Vial (1978) proposed a “simple extension” of Aumann’s (1974) correlated equilibrium that is completely different from the ones reviewed above. Like the strategic-form correlated equilibrium, their solution concept, which is sometimes referred to as *coarse correlated equilibrium* (Young (2004)), is described by a probability distribution μ on pure strategy profiles and applies to the strategic form of the game. However, the players do not receive any recommendation on how to play the game: each of them can just choose to either adhere to μ and get the corresponding correlated expected payoff or to deviate ex ante, by picking some strategy. The coarse correlated equilibrium conditions express that no player can gain by unilaterally deviating ex ante. Moulin and Vial’s solution concept assumes in effect some limited commitment from the players, who let the correlation device play for them at equilibrium.

Every EFCE defines a coarse correlated equilibrium: Namely, given an EFCE, it is clear that no player can benefit by ignoring the recommendations of the device at his information sets and deviating unilaterally before the beginning of the extensive form game.

2.5 Discussion and open problems

As observed in the comparison with other solution concepts, the EFCE can give rise to a larger set of outcomes than a communication equilibrium. A device which can give every player a recommendation that depends on a player’s information set, which in a Bayesian game represents the player’s type, may be considered rather powerful. That is, the extended game that defines the EFCE as in Definition 2.2 can be viewed as changing the game quite substantially. However, we think this is a natural approach when information sets define the rules of the game, rather than defining types which may be not easily “verifiable”. In other words, our concept applies to games with imperfect information rather than games with incomplete information. The standard equivalence between these games (Harsanyi (1967)), which is undisputed for Nash equilibria, may become controversial for correlated equilibria.

In distinction to communication equilibria where players can send signals to the device, an EFCE does not change the game by giving additional moves to the players. The recommended move is associated with the information set of the player, but this does not assume an omniscient device that knows the game state.

An interpretation of the moves that are generated in an EFCE would be recommendations that are put into “sealed envelopes” which a player can only open when reaching the respective information set. We assume that the players cannot obtain the information earlier, in the same way as we assume that the information sets describe the rules of the

game which the players must obey. An interesting open question is how to implement “sealed envelopes” in this context by cryptographic techniques. A starting point may be Dodis, Halevi, and Rabin (2000) and Urbano and Vila (2002) who use cryptography to replace the mediator in a correlated equilibrium.

Does the EFCE concept reflect “common knowledge of rationality for extensive games with Bayesian players”, in analogy to Aumann’s (1987) interpretation for strategic-form games? This should be confined to a static description of the game. In a dynamic description, rationality would also mean sequential rationality. Such a concept would lead to refinements such as subgame perfect, or sequential equilibrium. This is not the case for EFCE which include all Nash equilibria, including those that are not sequential. The EFCE concept seems well suited to address refinements such as perfection; see Dhillon and Mertens (1996) or Gerardi (2004, p. 117).

In Forges (1993), Aumann’s (1987) approach is extended to Bayesian games. The corresponding concept of a “belief-invariant Bayesian solution” is described by a probability distribution over types and actions such that the marginal distribution on types is that of the original game. In addition, the action of one player, given his own type, is conditionally independent of the other players’ types. The incentive constraints express that a player should choose the action recommended by an omniscient mediator who uses this distribution. Clearly, any agent normal form correlated equilibrium (which coincides with the EFCE in a Bayesian game), induces a belief-invariant Bayesian solution. However, contrary to the claim in Forges (1993, Proposition 3), there may be other belief-invariant Bayesian solutions; see Forges (2006).

3 Computational complexity

So far, we have argued that the EFCE is a “natural” concept for games in extensive form. In this second part of the paper, we show that the EFCE is also attractive from a computational point of view. We will show that the set of EFCE has a *compact description* given by a polynomial number of inequalities in the size of the game tree, provided the game has only *two players and no chance moves*. The last Section 3.7 gives hardness results showing that, given an extensive game, such a compact description cannot be expected for the set of strategic-form correlated equilibria, and also not for the set of EFCE if the game has chance moves (or a third player).

In an EFCE, the device recommends moves rather than strategies to the players. One motivation for this is a potential reduction in computational complexity, because the corresponding incentive constraints compare any two moves rather than any two pure strategies of a player. In addition to incentive constraints, we need *consistency constraints* that express how the moves at any two information sets are correlated.

First, we review in Section 3.1 the sequence form. This is a compact description of “realization plans” that specify the probabilities for playing sequences of moves, which can be translated to behavior strategy probabilities. Section 3.2 describes how to extend the constraints for realization plans to constraints for joint probabilities for *pairs* of sequences

(we always consider only two players), which we call “correlation plans”. Section 3.3 gives an example that illustrates the use of these constraints.

In general, the consistency constraints apply only to mutually “relevant” information sets that share a path in the game tree, as explained in Section 3.4. That section also describes implications of perfect recall for the information sets in games with two players and without chance moves, and defines the concept of a “reference sequence”, which is used to generate move recommendations. Based on these technical preliminaries, Section 3.5 shows how to use the consistency constraints as a compact description of a correlation device as used in an EFCE. The incentive constraints are described in Section 3.6. Computational difficulties that arise in games with chance moves are discussed in the final Section 3.7.

3.1 Review of the sequence form

The sequence form of an extensive game is similar to the reduced strategic form, but uses sequences of moves of a player instead of reduced strategies. Since player i has perfect recall, all nodes in an information set h in H_i define the same sequence σ_h of moves for player i (see Section 2.1). The sequence σ_h leading to h can be extended by an arbitrary move c in C_h . Hence, any move c at h is the last move of a unique sequence $\sigma_h c$. This defines all possible sequences of a player except for the empty sequence \emptyset . The set of sequences of player i is denoted S_i , so

$$S_i = \{\emptyset\} \cup \{\sigma_h c \mid h \in H_i, c \in C_h\}.$$

We will use the sequence form for characterizing EFCE of two-player games (without chance moves). Then we denote sequences of player 1 by σ and sequences of player 2 by τ , and for readability the sequence leading to an information set k of player 2 by τ_k .

The sequence form is applied to Nash equilibria as follows (see also von Stengel (1996), Koller, Megiddo, and von Stengel (1996), or von Stengel, van den Elzen, and Talman (2002)). Sequences are played randomly according to *realization plans*. A realization plan x for player 1 is given by nonnegative real numbers $x(\sigma)$ for $\sigma \in S_1$, and a realization plan y for player 2 by nonnegative numbers $y(\tau)$ for $\tau \in S_2$. They denote the realization probabilities for the sequences σ and τ when the players use mixed strategies. Realization plans are characterized by the equations

$$\begin{aligned} x(\emptyset) &= 1, & \sum_{c \in C_h} x(\sigma_h c) &= x(\sigma_h) & (h \in H_1), \\ y(\emptyset) &= 1, & \sum_{d \in C_k} y(\tau_k d) &= y(\tau_k) & (k \in H_2). \end{aligned} \tag{4}$$

The reason is that equations (4) hold when a player uses a behavior strategy, in particular a pure strategy, and hence also for a mixed strategy which is a convex combination of pure strategies. A realization plan x (and analogously, y) fulfilling (4) results from a behavior strategy of player 1 (respectively, player 2) that chooses move c at an information set

$h \in H_1$ with probability $x(\sigma_{hc})/x(\sigma_h)$ if $x(\sigma_h) > 0$ and arbitrarily if $x(\sigma_h) = 0$. This yields a canonical proof of the theorem of Kuhn (1953) that asserts that a player with perfect recall can replace any mixed strategy by an equivalent behavior strategy. The behavior at h is unspecified if $x(\sigma_h) = 0$, which means that h is unreachable due to an earlier own move. Not specifying the behavior at such information sets is exactly what is done in the reduced strategic form.

Sequence form *payoffs* are defined for profiles of sequences whenever these lead to a *leaf* (terminal node) of the game tree, multiplied by the probabilities of chance moves on the path to the leaf. Here, we consider the special case of two players and no chance moves, and extend the sequence form to a compact description of the set of EFCE.

The sequence form is much smaller than the reduced strategic form, because a realization plan is described by probabilities for the sequences of the player, whose number is the number of his moves. In contrast, a mixed strategy is described by probabilities for all pure strategies of the player, whose number is generally exponential in the size of the game tree.³ A polynomial number of constraints, namely one equation (4) for each information set (and nonnegativity), characterizes realization plans. These constraints can be used to describe Nash equilibria, as explained in the papers on the sequence form cited above.

3.2 Correlation plans and marginal probabilities

In the following sections, we consider an extensive two-player game with perfect recall and without chance moves. Then any leaf of the game tree defines a unique pair (σ, τ) of sequences of the two players. Let $a(\sigma, \tau)$ and $b(\sigma, \tau)$ denote the respective payoffs to the players at that leaf. Then if the two players use the realization plans x and y , their expected payoffs are given by the expressions, bilinear in x and y ,

$$\sum_{\sigma, \tau} x(\sigma)y(\tau)a(\sigma, \tau), \quad \sum_{\sigma, \tau} x(\sigma)y(\tau)b(\sigma, \tau), \quad (5)$$

respectively. The expressions in (5) represent the sums over all leaves of the payoffs multiplied by the probabilities of reaching the leaves. The sums in (5) may be taken over all $\sigma \in S_1$ and $\tau \in S_2$ by assuming that $a(\sigma, \tau) = b(\sigma, \tau) = 0$ whenever the sequence pair (σ, τ) does not lead to a leaf. This is useful when using matrix notation, where the payoffs in the sequence form are entries $a(\sigma, \tau)$ and $b(\sigma, \tau)$ of sparse $|S_1| \times |S_2|$ payoff matrices and x and y are regarded as vectors.

In order to describe an EFCE, the product $x(\sigma)y(\tau)$ in (5) of the realization probabilities for σ in S_1 and τ in S_2 will be replaced by a more general *joint* realization probability $z(\sigma, \tau)$ that the pair of sequences (σ, τ) is recommended to the two players, for a suitable correlation device μ , as far as this probability is relevant. These probabilities $z(\sigma, \tau)$ define what we call a *correlation plan* for the game.

³A class of games with exponentially large reduced strategic form is described by von Stengel, van den Elzen, and Talman (2002).

As a tentative definition, given in full in Definition 3.7 below, a correlation plan is a function $z: S_1 \times S_2 \rightarrow \mathbb{R}$ for which there is a probability distribution μ on the set of reduced strategy profiles Σ^* such that for each sequence pair (σ, τ) ,

$$z(\sigma, \tau) = \sum_{\substack{(p_1, p_2) \in \Sigma^* \\ (p_1, p_2) \text{ agrees with } (\sigma, \tau)}} \mu(p_1, p_2). \quad (6)$$

Here, the reduced pure strategy pair (p_1, p_2) agrees with (σ, τ) if p_1 chooses all the moves in σ and p_2 chooses all the moves in τ .

In an EFCE, a player gets a move recommendation when reaching an information set. The move corresponds uniquely to a sequence ending in that move. For player 1, say, the sequence denotes a row of the $|S_1| \times |S_2|$ correlation plan matrix. From this row, player 1 should have a posterior distribution on the recommendations to player 2. This behavior of player 2 must be specified not only when player 1 follows a recommendation, but also when player 1 deviates, so that player 1 can decide if the own recommendation is optimal; see also the example in Section 3.3. The recommendations to player 2 off the equilibrium path are therefore important, so the collection of recommended moves to player 2 has to define a reduced strategy. Otherwise, one could simply choose a distribution on the leaves of the tree (with a correlation plan that is a sparse matrix like the payoff matrix), and merely recommend to the players the pair of sequences corresponding to the selected leaf.

Our first approach is therefore to define a correlation plan z as a full matrix. Except for a scalar factor, a column of this matrix should be a realization plan of player 1, and a row should be a realization plan of player 2. According to (4) (except for the equations $x_\emptyset = 1$ and $y_\emptyset = 1$ that define the scalar factor), this means that for all $\tau \in S_2$, $h \in H_1$, $\sigma \in S_1$, and $k \in H_2$,

$$\sum_{c \in C_h} z(\sigma_h c, \tau) = z(\sigma_h, \tau), \quad \sum_{d \in C_k} z(\sigma, \tau_k d) = z(\sigma, \tau_k). \quad (7)$$

Furthermore, the pair (\emptyset, \emptyset) of empty sequences is selected with certainty, and the probabilities are nonnegative, which gives the trivial consistency constraints

$$z(\emptyset, \emptyset) = 1, \quad z(\sigma, \tau) \geq 0 \quad (\sigma \in S_1, \tau \in S_2). \quad (8)$$

Clearly, the constraints (7) and (8) hold for the special case $z(\sigma, \tau) = x(\sigma)y(\tau)$ where x and y are realization plans. With properly defined incentive constraints that make it an EFCE, such a correlation plan of rank one should define a Nash equilibrium. In particular, if x and y stand for reduced pure strategies, where each sequence σ or τ is chosen with probability zero or one, then the probabilities $z(\sigma, \tau) = x(\sigma)y(\tau)$ are also zero or one, and equations (7) and (8) hold. For any *convex combination* of pure strategy pairs, as in an EFCE, (7) and (8) therefore hold as well, so these are *necessary* conditions for a correlation plan.

Figure 2 shows a correlation plan defined in this manner for the game in Figure 1. In order to have distinct move names at different information sets, moves X and Y at the information set of the good type are called X_G and Y_G , those of of the bad type X_B and Y_B ,

	\emptyset	l_X	r_X	l_Y	r_Y
\emptyset	1	1	0	1	0
X_G	1	1	0	1	0
Y_G	0	0	0	0	0
X_B	0	0	0	0	0
Y_B	1	1	0	1	0

FIGURE 2

	\emptyset	l_X	r_X	l_Y	r_Y
\emptyset	1	1/2	1/2	1/2	1/2
X_G	1/2	1/2	0	1/2	0
Y_G	1/2	0	1/2	0	1/2
X_B	1/2	0	1/2	1/2	0
Y_B	1/2	1/2	0	0	1/2

FIGURE 3

and the moves of player 2 are l_X and r_X when she receives signal X and l_Y and r_Y when she receives signal Y . Since both players move only once, every non-empty sequence is just a move. The correlation plan in Figure 2 arises from the pure strategy pair $(X_G Y_B, l_X l_Y)$.

Figure 3 shows a possible assignment of probabilities $z(\sigma, \tau)$ that fulfills (7) and (8). These probabilities are “locally consistent” in the sense that the marginal probability of each move is 1/2. However, they *cannot* be obtained as a convex combination of pure strategy pairs like the pure strategy pair in Figure 2. Otherwise, one such pair would have to recommend move X_G to player 1 and move l_X to player 2 to account for the respective entry 1/2. In that pure strategy pair, given that player 2 is recommended move l_X , the recommendation to player 1 at the other information set must be Y_B because the move combination (X_B, l_X) has probability zero. Similarly, move X_G requires that move l_Y is recommended to player 2. This pure strategy pair is thus $(X_G Y_B, l_X l_Y)$ as in Figure 2, but that pair also selects (Y_B, l_Y) , contradicting Figure 3. This shows that (7) and (8) do not suffice to characterize the convex hull of pure strategy profiles. For games with chance moves, Theorem 3.10 below shows that this convex set cannot be characterized by a polynomial number of linear inequalities (unless $P = NP$).

However, we will show that the constraints (7) and (8) suffice to characterize correlation plans when the game has only two players and no chance moves.

3.3 Example of generating move recommendations

Figure 4 is a game very similar to Figure 1, except that the initial chance move is replaced by a move by player 1, as if that player “chose his own type”. A similar analysis as in Section 2.3 shows that there is only one outcome in a strategic-form or autonomous correlated equilibrium, or communication equilibrium, which is non-revealing.

Figure 5 gives an example of probabilities $z(\sigma, \tau)$ that fulfill (7) and (8). We demonstrate how to generate a pair of reduced strategies using z , described in general in Section 3.5 below. We consider only the generation of moves, and not any incentive constraints (treated in Section 3.6), which are in fact violated in Figure 5.

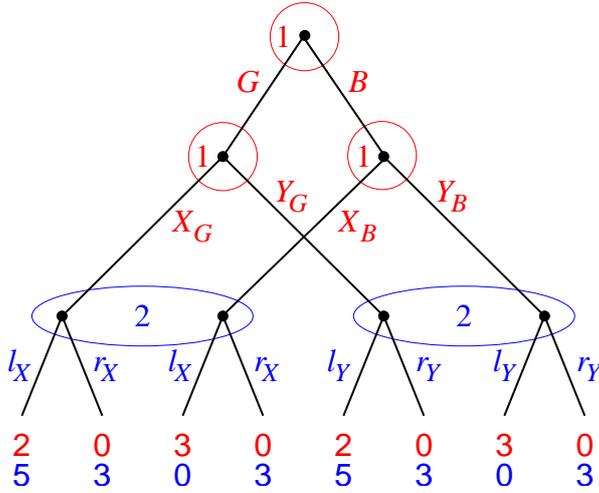


FIGURE 4

	\emptyset	l_X	r_X	l_Y	r_Y
\emptyset	1	1/2	1/2	1/2	1/2
G	1/2	1/4	1/4	1/4	1/4
B	1/2	1/4	1/4	1/4	1/4
GX_G	1/4	1/4	0	1/4	0
GY_G	1/4	0	1/4	0	1/4
BX_B	1/4	0	1/4	1/4	0
BY_B	1/4	1/4	0	0	1/4

FIGURE 5

The generation of moves starts at the root of the game tree. The information set containing the root belongs to player 1 and has the two moves G and B . We consider a “reference sequence” of the other player, which is here $\tau = \emptyset$ of player 2 because that is the sequence of player 2 leading to the root. This reference sequence τ determines a column of z describing the probabilities for making a move G or B . In Figure 5, $z(G, \tau) = z(B, \tau) = 1/2$. Suppose that move G is chosen. The next information set belongs again to player 1 with moves X_G and Y_G . The reference sequence is still $\tau = \emptyset$. The moves of player 1 correspond to the sequences GX_G and GY_G , which have probabilities $z(GX_G, \tau) = z(GY_G, \tau) = 1/4$ in Figure 5. These probabilities have to be divided by $z(G, \tau)$ to obtain the conditional probabilities for generating the moves, which are here both $1/2$; the respective general equation is (10) below. Suppose that move X_G is chosen.

The next information set to be considered (because it still precedes any information set of player 2) is the information set of player 1 with moves X_B and Y_B . However, this information set is unreachable due to player 1’s earlier move G . Because it suffices to generate only a reduced strategy of player 1 as explained in Section 2.2, no move is recommended at this information set. All information sets of player 1 have been considered, so the generated reduced strategy is $(G, X_G, *)$; recall that the moves in that strategy are recommended to player 1 when he reaches his respective information sets.

The remaining information sets belong to player 2. For the information set with moves l_X and r_X , the reference sequence is $\sigma = GX_G$ because these moves have been generated for player 1 and reach player 2’s information set. This reference sequence σ determines a row in Figure 5 where $z(\sigma, l_X) = 1/4$ and $z(\sigma, r_X) = 0$. Normalized by dividing by the probability $z(\sigma, \emptyset) = 1/4$ for the incoming sequence \emptyset of player 2, this means l_X is chosen with certainty.

The information set, say k , of player 2 with moves l_Y and r_Y is interesting because it will not be reached when player 1 plays his recommended moves G and X_G . Nevertheless, a move at k must be recommended to player 2 because player 1 must be able to decide if choosing his recommended move X_G is optimal, or if Y_G is better. Player 1 can only decide this if he has a posterior over the moves l_Y or r_Y of player 2. The reference sequence for player 2's selection is again $\sigma = GX_G$ because its last move X_G is made at the unique information set of player 1 that still allows to reach k , described in generality in Section 3.5. According to Figure 5, $z(\sigma, l_Y) = 1/4$ and $z(\sigma, r_Y) = 0$, so l_Y is also chosen with certainty. The reduced strategy whose moves are recommended to player 2 is therefore (l_X, l_Y) .

The four squares at the bottom right of Figure 5 describe a correlation between the moves at pairs of information sets of player 1 and player 2, with nonzero entries like in Figure 3. However, unlike in Figure 3, these numbers are not only “locally” but also “globally” consistent in the sense that they can arise from a distribution μ on reduced strategy profiles. The reason is that, for example, the moves l_Y and r_Y of player 2 are correlated with *either* X_G and Y_G or X_B and Y_B of player 1, depending on the first move G or B of player 1, but not with both move pairs. In contrast, the conflict in Figure 3 arises because G or B is chosen by a chance move.

3.4 Information structure of two-player games without chance moves

In the following sections, we consider only two-player games without chance moves. Using the condition of perfect recall, we describe structural properties of information sets in such games. We then define the concepts of *relevant* sequence pairs and *reference sequences*, which we use later in Theorem 3.8.

Definition 3.1. In an extensive game, call any two information sets h and k (possibly of the same player) *connected* if there is a path from the root to a leaf containing a node of h and a node of k . If the node in h comes earlier on the path, then h is said to *precede* k .

The following lemma states that two-player games without chance moves have a weak “time structure”.

Lemma 3.2. *Consider a two-player extensive game without chance moves and with perfect recall. Then for any two information sets h and k , if h precedes k , then k does not precede h .*

Proof. Let h and k be two information sets so that h precedes k , let u be a node in h and let v be a node in k so that there is a path from u to v in the tree.

Suppose that, contrary to the claim, k also precedes h , with $v' \in k$ and $u' \in h$ so that there is a path from v' to u' . If h and k belong to the same player, then v is preceded by a move at h (the move made at u), and so is v' by perfect recall, so there is some other node in h from which there is a path via v' to u' ; however, it is easy to see that with perfect recall, no two nodes in an information set share a path. So h and k belong to different players.

Consider the last common node w on the two paths from the root to u and v' , respectively. If $w \in h$, then there is a path from w via v' to $u' \in h$, which is not possible. The same reasoning shows that $w \notin k$, because otherwise there is a path from w via u to $v \in k$. So w belongs to an information set other than h or k , with a move c leading to u and a different move c' leading to v' . Then the player to move at w does not have perfect recall, because in his later information set h or k , there are two nodes that are preceded by different own moves c and c' at w . \square

If σ and σ' are sequences of moves of a player, then the sequence σ is called a *prefix* of σ' if $\sigma = \sigma'$ or if σ' is obtained from σ by appending some moves; it is called a *proper prefix* if $\sigma \neq \sigma'$.

The following simple observations will be used repeatedly.

Lemma 3.3. *Consider a two-player perfect-recall extensive game without chance moves, and let $h, h' \in H_1$ and $k, k' \in H_2$ so that h precedes k . Then the following hold (as well as the symmetric statements with the players exchanged):*

- (a) *if h' precedes h then h' precedes k ;*
- (b) *if k' precedes k then k' and h are connected;*
- (c) *if h' precedes k and h and h' are not connected, then there is an information set h'' in H_1 that precedes both h and h' with different moves $c, c' \in C_{h''}$ leading to h and h' , respectively, that is, σ_h has a prefix of the form $\sigma_{h''}c$ and $\sigma_{h'}$ has a prefix of the form $\sigma_{h''}c'$.*

Proof. Because h precedes k , there is a path from the root to some node v in k that has a node u in h . Then (a) holds because some node of h is preceded by a move at h' , and thus by perfect recall node u is also preceded by that move at h' . Similarly, (b) holds because v is preceded by some node in k' which is therefore also on the path from the root to v , which contains u .

To prove (c), consider two paths from the root to k that intersect h and h' , respectively. These paths split at some point because h and h' are not connected. Consider the last common node u'' on these two paths. That is, from u'' onwards, the paths follow along different moves c and c' to h and h' , respectively, and subsequently reach k . Then u'' belongs to an information set h'' of player 1, because otherwise player 2 would not have perfect recall. That is, $c, c' \in C_{h''}$ so that $c \neq c'$ and h'' precedes h and h' , as claimed. \square

As considered so far in (6), a correlation plan z describes how to correlate moves at any two information sets of player 1 and player 2. However, it suffices to specify only correlations of moves at connected information sets where decisions can affect each other during play. We will specify $z(\sigma, \tau)$ only for “relevant” sequence pairs (σ, τ) . (A motivating example is Figure 6, discussed below.)

Definition 3.4. Consider a two-player extensive game with perfect recall. The pair (σ, τ) in $S_1 \times S_2$ is called *relevant* if σ or τ is the empty sequence, or if $\sigma = \sigma_h c$ and $\tau = \tau_k d$ for connected information sets h and k , where $h \in H_1$, $c \in C_h$, $k \in H_2$, $d \in C_k$. Otherwise, (σ, τ) is called *irrelevant*.

Note that in Definition 3.4, the *information sets* are connected where the respective last move in σ and τ is made. It is not necessary that the sequences themselves share a path. We specify correlations of moves at connected information sets, not just of moves that share a path, because a player may consider deviations from the recommended moves. The following lemma shows that it makes sense to restrict the equations (7) to relevant sequence pairs.

Lemma 3.5. *Consider a two-player extensive game without chance moves and with perfect recall. Assume that the pair (σ, τ) of sequences is relevant, and that σ' is a prefix of σ and that τ' is a prefix of τ . Then (σ', τ') is relevant.*

Proof. If σ or τ is the empty sequence, then so is σ' or τ' , respectively, and (σ', τ') is relevant by definition.

Let $\sigma = \sigma_h c$ and $\tau = \tau_k d$, where h and k are information sets of player 1 and 2, respectively. Since h and k are connected, assume that h precedes k ; the case that k precedes h is symmetric. If σ' or τ' is empty, the claim is trivial, otherwise let $\sigma' = \sigma_{h'} c'$ and $\tau' = \tau_{k'} d'$ for $h' \in H_1$ and $k' \in H_2$.

We first show that (σ', τ) is relevant, so let $h \neq h'$. Then h' precedes h , and h' precedes k by Lemma 3.3(a).

Similarly, (σ', τ') is relevant, which only needs to be shown for $k' \neq k$: Then k' and h' precede k , and k' and h' are connected by Lemma 3.3(b). \square

For an inductive generation of recommended moves, we restrict the concept of relevant sequence pairs further. The concept of a “reference sequence” was mentioned in the example in Section 3.3. A reference sequence τ of player 2, for example, defines a “column” of z (like in Figure 5) to select a move c at some information set h of player 1; then τ is called the reference sequence for $\sigma_h c$. We give the formal definition for both players.

Definition 3.6. Consider a two-player extensive game without chance moves and with perfect recall, and let $(\sigma, \tau) \in S_1 \times S_2$. Then τ is called a *reference sequence* for σ if $\sigma = \sigma_h c$ and

- (a1) $\tau = \emptyset$, or $\tau = \tau_k d$ and k precedes h , and
- (a2) there is no k' in H_2 with $\tau_{k'} = \tau$ that precedes h .

Correspondingly, σ is called a *reference sequence* for τ if $\tau = \tau_k d$ and

- (b1) $\sigma = \emptyset$, or $\sigma = \sigma_h c$ and h precedes k , and
- (b2) there is no h' in H_1 with $\sigma_{h'} = \sigma$ that precedes k .

If τ is a reference sequence for $\sigma_h c$, then all information sets where player 2 has made the moves in τ precede h , according to Definition 3.6(a1), and by (a2), τ cannot be extended to a longer sequence with that property (because the next move in such a longer sequence would be at an additional information set k' with $\tau_{k'} = \tau$ that precedes h). Note, however, that if $\tau = \tau_k d$, the information set h may not be reachable after the move d of player 2; it is only required that the information set k precedes h .

3.5 Using the consistency constraints

In this section, we first restrict the definition (6) of correlation plan probabilities $z(\sigma, \tau)$ to pairs of relevant sequences (σ, τ) . We then show the central result that the constraints (8) and (7), restricted to relevant sequence pairs, characterize a correlation plan. For that purpose, any solution z to these constraints is used to generate, as a random variable, a pair of reduced pure strategies to be recommended to the two players. The moves in that reduced strategy pair are generated inductively, assuming moves at preceding information sets have already been generated; these moves define each time a suitable reference sequence for the next generated move.

Definition 3.7. Consider a two-player extensive game without chance moves and with perfect recall. A *correlation plan* is a partial function $z: S_1 \times S_2 \rightarrow \mathbb{R}$ so that there is a probability distribution μ on the set of reduced strategy profiles Σ^* so that for each relevant sequence pair (σ, τ) , the term $z(\sigma, \tau)$ is defined and fulfills (6).

Theorem 3.8. *In a two-player, perfect-recall extensive game without chance moves, z is a correlation plan if and only if it fulfills (8), and (7) whenever (σ_{hc}, τ) and (σ, τ_{kd}) are relevant, for any $c \in C_h$ and $d \in C_k$. A corresponding probability distribution μ on Σ^* in Definition 3.7 is obtained from z by generating the moves in a reduced pure strategy pair inductively by an iteration over all information sets.*

Proof. As already mentioned, (7) and (8) are necessary conditions for a correlation plan, because they hold for reduced pure strategy profiles and therefore for any convex combination of them, as given by a distribution μ on Σ^* .

Consider now a function z defined on $S_1 \times S_2$ that fulfills (8), and (7) for relevant sequence pairs. Using z , a pair (p_1, p_2) of reduced pure strategies is generated as a random variable. We will show that the resulting distribution μ on Σ^* has the correlation plan z .

The moves in (p_1, p_2) are generated one move at a time, taking the already generated moves into account. For that purpose, we generalize reduced strategies as follows. Define a *partial strategy* of player i as an element of

$$\prod_{h \in H_i} (C_h \cup \{*\}).$$

Let the components of a partial strategy p_i of player i be denoted by $p_i(h)$ for $h \in H_i$. When $p_i(h) = *$, then $p_i(h)$ is undefined for the information set h , otherwise $p_i(h)$ defines a move at h , that is, $p_i(h) \in C_h$.

If σ is a sequence of player i and p_i is a partial strategy of player i , then p_i *agrees* with σ if p_i prescribes all the moves in σ , that is, $p_i(h) = c$ for any move c in σ , where $c \in C_h$. The information set h is *reachable when playing p_i* if p_i agrees with σ_h . It is easy to see that a reduced strategy of player i is a partial strategy p_i so that for all h in H_i , the move $p_i(h)$ is defined if and only if p_i agrees with σ_h .

Initially, p_1 and p_2 are partial strategies that are everywhere undefined, and eventually both are reduced strategies. In an iteration step, an information set h of player i is considered where all information sets (of either player) that precede h have already been treated

in a previous step. For h , a move c in C_h is generated randomly, according to z as described below, provided h is reachable when playing p_i . If this is not the case, that is, if p_i does not agree with σ_h , then $p_i(h)$ remains undefined. In that sense, the partial strategies p_i will always be *reduced* partial strategies. The iteration proceeds “top down” (in the direction of play), starting from the root. It cannot “get stuck” because of Lemma 3.2.

To define the iteration step, consider the pair (p_1, p_2) of reduced partial strategies generated so far, which is not yet a pair of reduced strategies. Let h be an information set, say of player 1 (the case for player 2 is analogous), so that for all information sets k preceding h , where k may belong to either player i , the move $p_i(k)$ is defined, or undefined because k is unreachable when playing p_i . Initially, when p_1 and p_2 are everywhere undefined, h is the information set containing the root of the game tree. If h is unreachable when playing p_1 , the move $p_1(h)$ stays undefined. Otherwise, p_1 agrees with σ_h .

The move $c = p_1(h)$ will be generated based on a reference sequence for $\sigma_h c$. This sequence consists of the moves that player 2 makes at the information sets that precede h when player 2 plays as in p_2 . These moves form a sequence because of Lemma 3.3(c): Let

$$K = \{k \in H_2 \mid k \text{ precedes } h \text{ and } \tau_k \text{ agrees with } p_2\}. \quad (9)$$

We claim that for any two information sets k and k' in K , one precedes the other or vice versa. Otherwise, if there are k and k' in K that are not connected, we obtain a contradiction as follows: Lemma 3.3(c) (with the players exchanged) shows that k and k' are preceded by distinct moves d and d' at an information set k'' of player 2 that precedes h . Because k and k' were reachable when playing p_2 , so is k'' , so that $p_2(k'')$ is defined. However, of the two moves d and d' , at most one can be chosen by p_2 , and so that p_2 cannot agree with both τ_k and $\tau_{k'}$, that is, k and k' cannot both belong to K . This proves our claim.

If K in (9) is empty, let $\tau = \emptyset$. Otherwise, let k be the unique last information set in K not preceding any other, and let $\tau = \tau_k d$, where $d = p_2(k)$. Then τ is a reference sequence for $\sigma_h c$ for any move c at h by construction of K .

The pair of partial strategies (p_1, p_2) generated so far agrees with (σ_h, τ) . Consequently, all moves in (σ_h, τ) have been generated, and this event has positive probability. We will show shortly by induction that this probability is $z(\sigma_h, \tau)$. For the base case of the induction where $(\sigma_h, \tau) = (\emptyset, \emptyset)$, this is true because $z(\emptyset, \emptyset) = 1$ by (8).

Given the described reference sequence τ , the move c at h is generated randomly according to the probability

$$\beta(c, \tau) = \frac{z(\sigma_h c, \tau)}{z(\sigma_h, \tau)} \quad (c \in C_h), \quad (10)$$

where by inductive assumption $z(\sigma_h, \tau) > 0$. The probability $\beta(c, \tau)$ is well defined when considering h in the induction, because it only depends on having generated the moves in σ_h (as part of p_1) and in τ (as part of p_2); any other moves in p_2 do not matter because they are not at information sets that precede h , by the definition of K in (9). By construction of τ , the sequence pairs (σ_h, τ) and $(\sigma_h c, \tau)$ in (10) are relevant. Moreover, (10) defines a probability distribution on C_h by (7) and (8).

When all information sets have been considered, (p_1, p_2) is a pair of reduced strategies. The described process of generating moves defines a distribution μ on Σ^* .

For any relevant pair of sequences (σ, τ) , let

$$\mu(\sigma, \tau) = \sum_{\substack{(p_1, p_2) \in \Sigma^* \\ (p_1, p_2) \text{ agrees with } (\sigma, \tau)}} \mu(p_1, p_2).$$

In the process described above, a move is generated once for each reachable information set, so $\mu(\sigma, \tau)$ is the probability that all moves in (σ, τ) are generated. We want to show (6), that is,

$$\mu(\sigma, \tau) = z(\sigma, \tau), \quad (11)$$

for all relevant sequence pairs (σ, τ) . If σ or τ is the empty sequence, this imposes no constraint on the moves of the respective player. Thus, if $(\sigma, \tau) = (\emptyset, \emptyset)$, then (11) holds because $z(\emptyset, \emptyset) = 1$ by (8). If at least one of the sequences σ or τ is not empty, then according to Definition 3.4 one of the following cases applies:

- (a) $(\sigma, \tau) = (\sigma_h c, \emptyset)$, or $(\sigma, \tau) = (\sigma_h c, \tau_k d)$ and k precedes h ; or, symmetrically,
- (b) $(\sigma, \tau) = (\emptyset, \tau_k d)$, or $(\sigma, \tau) = (\sigma_h c, \tau_k d)$ and h precedes k .

Using Definition 3.6 and Lemma 3.5, it is easy to see that (a) and (b) are, respectively, equivalent to the statements

- (a') τ is the prefix of a reference sequence for $\sigma = \sigma_h c$,
- (b') σ is the prefix of a reference sequence for $\tau = \tau_k d$.

We prove (11) for case (a') with a two-part induction; the same reasoning applies to (b') by symmetry. The ‘‘outer’’ inductive assumption is that (11) holds for $(\sigma, \tau) = (\emptyset, \emptyset)$, and for case (a') with h' instead of h for any information set h' that precedes h , and for case (b') for any k that precedes h .

We prove (11) with a second ‘‘inner’’ induction over the prefixes τ of reference sequences for σ_h as in (a'), where we consider the longest prefixes first. We say that the prefix τ of a reference sequence for $\sigma_h c$ has *distance* n if n is the largest number of moves d_1, d_2, \dots, d_n of player 2 so that $\tau d_1 d_2 \dots d_n$ is a reference sequence for $\sigma_h c$. We will prove by induction on n : If τ is the prefix of a reference sequence for $\sigma_h c$ of distance n , then $\mu(\sigma_h c, \tau) = z(\sigma_h c, \tau)$. Then this shows (11) for case (a').

If $n = 0$, the sequence τ is itself a reference sequence for $\sigma_h c$. That is, move c is generated according to (10) with probability $\beta(c, \tau)$, so that $\mu(\sigma_h c, \tau) = \beta(c, \tau) \cdot \mu(\sigma_h, \tau)$. The moves in σ_h and τ are all made at information sets that precede h , so by the ‘‘outer’’ inductive hypothesis, $\mu(\sigma_h, \tau) = z(\sigma_h, \tau)$. Consequently, $\mu(\sigma_h c, \tau) = \beta(c, \tau) \cdot z(\sigma_h, \tau) = z(\sigma_h c, \tau)$. This proves the base case $n = 0$ for the ‘‘inner’’ induction.

Suppose that $n > 0$ and that τ is the prefix of a reference sequence for $\sigma_h c$ of distance n . As ‘‘inner’’ inductive hypothesis, (11) holds for such sequences for all smaller values of n . Because $n > 0$, there is an information set k in H_2 with $\tau_k = \tau$ so that k precedes h ; similar to the construction of K in (9), this information set k is seen to be unique with the help of Lemma 3.3(c). Then for all $d \in C_k$, the sequences $\tau_k d$ are all prefixes of reference sequences for $\sigma_h c$ of distance less than n , so by the ‘‘inner’’ inductive hypothesis,

$\mu(\sigma_{hc}, \tau_k d) = z(\sigma_{hc}, \tau_k d)$. If all the moves in σ_{hc} and τ_k are generated, then exactly one of the moves in C_k is generated. This implies

$$\mu(\sigma_{hc}, \tau_k) = \sum_{d \in C_k} \mu(\sigma_{hc}, \tau_k d) = \sum_{d \in C_k} z(\sigma_{hc}, \tau_k d) = z(\sigma_{hc}, \tau_k). \quad (12)$$

This completes the “inner” and thereby also the “outer” induction.

This shows (6) for all relevant sequence pairs (σ, τ) , so that z is indeed the correlation plan corresponding to μ . \square

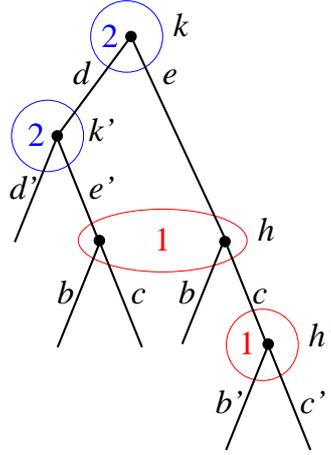


FIGURE 6

The example in Figure 6 demonstrates the two-part induction in the preceding proof. Player 1 has the information sets h and h' , and player 2 has k and k' . Except for h' and k' , any two of these are connected. The sets of sequences of player 1 and 2 are $S_1 = \{\emptyset, b, c, cb', cc'\}$ and $S_2 = \{\emptyset, d, e, dd', de'\}$. Any non-empty sequence of player 2 has the reference sequence \emptyset of player 1, so that player 2's move recommendations are generated first. For the sequences of player 1 that end in a move at h , the possible reference sequences are dd' , de' , or e . For the sequences that end in a move at h' , the reference sequences are d or e . That is, reference sequences can be “non-monotonic”, in the sense that “later” information sets (here h' , preceded by h) can have “shorter” reference sequences (here d , which is a reference sequence for $\sigma_{h'}c'$, which is a proper prefix of the reference sequence dd' or de' for the sequence $\sigma_hc = \sigma_{h'}$). For this reason, one needs the second, “inner” induction step (12) in the preceding proof, which amounts here to proving that $\mu(c, d) = \mu(c, dd') + \mu(c, de')$. In this example, all other cases of (11) involve a reference sequence directly, so that only the base case of the inner induction is required.

Figure 6 also demonstrates the use of reference sequences, and why it is useful to restrict $z(\sigma, \tau)$ to relevant sequence pairs. If one did not do the latter, one could specify probabilities $z(cb', dd')$, $z(cb', de')$, $z(cc', dd')$, and $z(cc', de')$ subject to (7), which correlate the moves at the two information sets h' and k' . However, this correlation would not matter, but only the marginal probabilities $z(cb', d)$ and $z(cc', d)$ when player 2 has to

decide at her information set k whether to follow the recommendation to play d or not. Moreover, such an over-specified correlation plan would be hard to translate into a generation of moves, since it is not clear at what point one should generate the move at h' (which would be even more difficult if there were other information sets following move d' of player 2). Instead, the concept of a reference sequence leads to a well-defined way for generating moves using z , as described in the proof of Theorem 3.8.

3.6 Incentive constraints

In an EFCE, a player gets a move recommendation when reaching an information set. This recommendation induces a posterior distribution on the recommendations given to the other player. For past moves, this induces a certain distribution on where the player is in the information set. For future moves, it expresses the subsequently expected play. Both are represented by the eventual distribution on the leaves of the game tree. The players want to optimize the expected payoffs which they receive at the leaves, assuming the other player follows his or her recommendations.

The *incentive constraints* in an EFCE express that it is optimal to follow any move recommendation, under two assumptions about the player's own behavior: When following the recommended move, the player considers his expected payoff when he follows recommendations in the future. When deviating from the recommended move, the player optimizes his payoff, given the current knowledge about the other player's behavior. Any recommendations given after a deviation are ignored, and are in fact not given, because an EFCE only generates a pair of reduced strategies: When a player deviates, he subsequently only reaches own information sets that would be unreachable when following the original move in the strategy, so these later moves are left unspecified in a reduced strategy.

Assume that a pair of reduced strategies is generated according to a correlation plan z as in Theorem 3.8. Suppose that player 1, say, gets a recommendation for a move c at h which is the last move of the sequence $\sigma = \sigma_h c$. For the sequences τ of player 2, the row entries $z(\sigma, \tau)$ of the correlation plan z define, up to normalization, a realization plan that describes player 2's behavior. This is only given where (σ, τ) is relevant, which suffices for the decision of player 1 whether to accept the recommendation to play move c .

In order to state the incentive constraints, we first introduce auxiliary variables $u(\sigma)$ for any $\sigma \in S_1$ (and, throughout, analogously for player 2). These denote the expected payoff contribution of σ (that is, of all reduced strategies agreeing with σ) when player 1 follows his recommendations. They are given by

$$u(\sigma) = \sum_{\tau} z(\sigma, \tau) a(\sigma, \tau) + \sum_{k \in H_1: \sigma_k = \sigma} \sum_{d \in C_k} u(\sigma_k d). \quad (13)$$

(All incentive constraints will refer to information sets h, k, l and moves c, d of a single player.) In (13), $a(\sigma, \tau)$ is the payoff to player 1 at the leaf that defines the sequence pair (σ, τ) , which is then obviously a relevant pair; if there is no such leaf, $a(\sigma, \tau) = 0$. The first sum in (13) captures the expected payoff contribution where σ and suitable

sequences τ of player 2 are defined by leaves. The second, double sum in (13) concerns the information sets k of player 1 reached by σ . The sum of the payoff contributions $u(\sigma_k d)$ for $d \in C_k$ is the expected payoff when player 1 follows the recommendation to choose d at k , given the new posterior information that he obtains there.

Applying (13) inductively, starting with the longest sequences, gives eventually for the empty sequence $u(\emptyset) = \sum_{\sigma, \tau} z(\sigma, \tau) a(\sigma, \tau)$. This denotes the overall payoff to player 1 under the correlation plan z (and similarly for player 2), which generalizes (5).

When σ is not the empty sequence, the payoff $u(\sigma)$ when player 1 chooses the recommended last move c of $\sigma = \sigma_h c$ must be compared with the possible payoff when player 1 deviates from his recommendation. This is described by an optimization against the behavior of player 2 in row σ of z , by considering the other moves at h , as well as moves at information sets k that player 1 can reach later on. By optimizing in this way, the payoff contribution at an information set k of player 1 is denoted by $v(k, \sigma)$. The parameter σ indicates the given row of the correlation plan z against which player 1 optimizes. The optimal payoff $v(k, \sigma)$ at an information set k of player 1 is the maximum of the payoffs for the possible moves at k , which may either directly give a payoff when they lead to a leaf, or are obtained from subsequent optimal payoffs at later information sets. This is expressed by the following inequalities, for any $k \in H_1$ with $k = h$ or h preceding k (where $\sigma = \sigma_h c$), and all moves d at k :

$$v(k, \sigma) \geq \sum_{\tau} z(\sigma, \tau) a(\sigma_k d, \tau) + \sum_{l \in H_1: \sigma_l = \sigma_k d} v(l, \sigma) \quad (d \in C_k). \quad (14)$$

The first sum in (14) is well defined, because when $(\sigma_k d, \tau)$ leads to a leaf, then (σ, τ) is relevant because $\sigma = \sigma_h c$ and σ_h is a prefix of $\sigma_k d$. These incentive constraints are completed by

$$v(h, \sigma_h c) = u(\sigma_h c) \quad (15)$$

which says that the recommended move c at h is optimal.

As an illustration of the incentive constraints, consider an information set h that precedes no further information sets of player 1. Then (13), (15) and (14) show that

$$u(\sigma_h c) = \sum_{\tau} z(\sigma_h c, \tau) a(\sigma_h c, \tau) \geq \sum_{\tau} z(\sigma_h c, \tau) a(\sigma_h d, \tau), \quad (d \in C_h) \quad (16)$$

which says that player 1 cannot gain by changing his move c at h to d . This is analogous to the incentive constraint in a strategic-form correlated equilibrium that states that player 1, say, cannot gain by changing from the recommended strategy to some other strategy. In both cases, the posterior on player 2's behavior is given by the recommended "row" of the joint distribution, in (16) given by row $\sigma_h c$ of z .

The number of variables $v(k, \sigma)$ is quadratic in the number of sequences of player 1 because they are parameterized by the information sets k and the sequences σ . The latter reflect player 1's current information about the behavior of the other player, which varies in a correlated equilibrium. By comparison, this behavior is fixed in a Nash equilibrium, where $z(\sigma, \tau)$ is replaced by $y(\tau)$ with a constant realization plan y of player 2. Furthermore, the variables $v(k, \sigma)$ are replaced by variables $v(k)$, one for each information set k

of player 1. Then the inequalities (14) are exactly those expressing the Nash equilibrium condition, with certain *dual* variables $v(k)$, normally derived from linear programming duality. These dual variables also express, like here, the computation of the player’s optimal payoff by “dynamic programming”, as described by von Stengel (1996, p. 239).

Together with the consistency constraints, the stated incentive constraints characterize an EFCE. We summarize our main result as follows.

Theorem 3.9. *In a two-player, perfect-recall extensive game without chance moves, a correlation plan z as in Theorem 3.8 that fulfills for both players the incentive constraints (13), (14), and (15) defines an EFCE. The size of these constraints that describe the set of EFCE is polynomial in the size of the game tree, so that an EFCE is polynomial-time computable.*

Proof. The correctness of the incentive constraints (13), (14), and (15) has already been argued above. In particular, suppose that player 1 considers deviating from the recommended move c at h . Then player 1’s posterior on the behavior of player 2 is given by $z(\sigma_h c, \tau)$ for any sequence τ in S_2 . When player 1 deviates and considers another move d at h , or a move d at any subsequent information set k , as in (14), his posterior about player 2 does not change, but remains based on the last own recommendation that player 1 received, as in $z(\sigma_h c, \tau)$, because player 1 will not get any further information. The sequences τ recommended to player 2 have probabilities $z(\sigma_h c, \tau)$ by Theorem 3.8. The recommendations are already generated, so they are not affected by player 1’s deviation, and the equilibrium condition assumes that player 2 follows the recommendations (even if she can conclude that player 1 has deviated, because the EFCE does not involve any kind of “subgame perfection”, see Section 2.5). \square

We mention an interesting case of (14), namely $k = h$ and $c = d$. This is the optimality condition applied to the *recommended* move c , where the player chooses to follow the move recommendation now, but henceforth ignores all future recommendations and the associated Bayesian update about the other player’s behavior. The constraint (14) with $k = h$ and $c = d$ states that such an optimization following move c , given the current knowledge about the other player as represented by the parameter $\sigma_h c$ of the variables $v(l, \sigma_h c)$, will not give higher payoff to the player than when following the recommendation as expressed by $u(\sigma_h c)$ in (13). In fact, this constraint can be omitted because it is implied by the other conditions. Intuitively, this holds because the player cannot gain by ignoring private information. The simple proof of this fact is analogous to the observation that any correlated equilibrium is a coarse correlated equilibrium as defined by Moulin and Vial (1978).

3.7 Hardness results

In the preceding sections, we have characterized the set of EFCE for two-player games without chance moves by means of a polynomial number of inequalities. In this final section, we show that such a compact description cannot be expected when chance moves

are admitted, because one could then maximize in polynomial time a linear function of the payoffs over the set of EFCE, for example the sum of the expected payoffs to the players. However, the following theorem shows that this problem is NP-hard, so that a polynomial number of constraints would imply $P = NP$. (Chu and Halpern (2001) have given a similar NP-hardness proof for finding optimal play in a “possible worlds” model.) We then show the corresponding theorem for strategic-form correlated equilibria of extensive games with two players and without chance moves. Examples and pictures illustrating the proofs can be found in von Stengel (2001).

Theorem 3.10. *For two-player, perfect-recall extensive games with chance moves, it is NP-hard to find a strategic-form correlated equilibrium (or an EFCE, or a Nash equilibrium) with maximum payoff sum.*

Proof. We give a reduction from 3SAT that converts a 3SAT formula ϕ with n clauses to an extensive two-player game with perfect recall, so that the game has a pure strategy pair (and Nash equilibrium) with payoff sum 2 if ϕ is satisfiable, and so that the payoff sum for any pure strategy pair is at most $2 - 2/n$ if ϕ is not satisfiable. This applies also to mixed-strategy Nash equilibria, correlated equilibria, and EFCE, because they are convex combinations of pure strategy pairs. This shows that finding any such equilibrium with maximum payoff sum is NP-hard.

We construct a two-player game from ϕ as follows. Player 2 has n decision nodes in singleton information sets, which correspond to the clauses. Player 1 has $3n$ decision nodes, corresponding to the $3n$ literals (negated or nonnegated variables) in the clauses. The game has $6n$ terminal nodes. If ϕ has m variables, then player 1 has m information sets, where each information set contains the “literal” nodes that have the same variable. An initial chance move at the root chooses with probability $1/n$ one of the n nodes of player 2. Player 2 is informed about the chance move and, for each clause chosen, selects one of the literals in the clause, which are nodes of player 1. Player 1 has two moves at each information set, with a move setting the respective variable to true or false. Both players then receive the same payoff, which is 0 if the literal (chosen by player 2 from the clause) is false and 1 if it is true.

This game has a pair of pure strategies for the two players with payoff 1 if and only if ϕ is satisfiable. The 2^m pure strategies of player 1 are the possible truth assignments to the variables in ϕ . A satisfying assignment defines a pure strategy for player 1, and player 2 can pick for each clause a literal that makes the clause true, so that both players get their maximum possible payoff 1. Conversely, if ϕ is not satisfiable, then any truth assignment to the m variables has at least one clause that is false, so that the respective move of player 2, which is chosen with probability $1/n$ by the chance move, leads to a payoff zero. The overall expected payoff to each player is then at most $1 - 1/n$. \square

Theorem 3.10 holds also when instead of chance moves, a third player is allowed in the game. In that case, the chance move is replaced by a move of player 3, who receives the negative of the identical payoffs to players 1 and 2. Player 3 has then an incentive to randomize, and the maximum payoff sum (which is equal to player 1’s payoff) in equilibrium is equal to 1 if and only if the 3SAT formula is satisfiable.

In the game constructed in the proof of Theorem 3.10 for a 3SAT formula with m variables and n clauses, the 2^m strategies of player 1 are truth assignments to the variables, like rows in a truth table. The game has an initial chance move, so that player 2 has 3^n strategies. If the chance move is replaced by a move of player 2, then the number of reduced strategies of player 2 is $3n$, but the strategic form is still exponential. Our second result uses such a construction to show that even for two-player games without chance moves, it is NP-hard to find a strategic-form correlated equilibrium with maximum payoff sum. The constructed game has a first stage given by a zero-sum generalized “rock-scissors-paper” game that induces player 2 to randomize, which replaces the initial chance move.

Theorem 3.11. *For two-player, perfect-recall extensive games without chance moves, it is NP-hard to find a strategic-form correlated equilibrium with maximum payoff sum.*

Proof. For a 3SAT formula ϕ with n clauses, consider the extensive game constructed in the proof of Theorem 3.10, but with the initial chance move replaced by a decision node of player 2 with n moves called c_1, \dots, c_n . These moves lead to n decision nodes of player 1 that all belong to a single information set. At that information set, player 1 has $n + 1$ moves called o_1, \dots, o_n , and “in”. Any move combination c_j, o_i for $1 \leq i, j \leq n$ leads to a separate terminal node with payoff a_{ij} to player 1 and $-a_{ij}$ to player 2, where

$$a_{ij} = 2n \cdot ((j - i) \bmod n - (n - 1)/2). \quad (17)$$

The n edges for move “in” of player 1 lead to the n “clause” nodes of player 2, with the rest of the game defined as before. Player 1 has perfect recall because all his later information sets (for the m variables of ϕ) are preceded by the same move “in”. Player 2 has perfect information and therefore perfect recall.

If player 1 uses only one of his “outside options” o_1, \dots, o_n , then the game is a zero-sum game with payoffs as in (17). It can be shown (for details see von Stengel (2001)) that this game has value zero with the uniform mixed strategy for each player as his or her optimal strategy. Thus, when player 2 chooses each move c_1, \dots, c_n with probability $1/n$, then player 1 will get payoff 0 when choosing any move o_1, \dots, o_n , which is worse than if he chooses “in” and, for example, subsequently randomizes with probability $1/2$ at each of his later information sets. In contrast, suppose player 2 chooses one move, say by symmetry move c_1 , with probability $q < 1/n$. Then it is easy to show that player 1 gets payoff $2 - 2nq$ when responding with the mixed strategy that chooses o_1, \dots, o_n with probabilities $2/n, 0, 1/n, \dots, 1/n$, and “in” with probability zero. If $q < 1/2n$, this payoff is larger than 1 so that player 1 would choose some of his outside options rather than “in”, with corresponding negative payoff to player 2. Clearly, this does not happen in any equilibrium.

Consider a strategic-form correlated equilibrium, and suppose that ϕ is not satisfiable. When player 1 is told his pure strategy by the correlation device, the conditional probability for each move c_1, \dots, c_n of player 2 is at least $1/2n$ since otherwise player 1 chooses an outside option rather than “in”. Since one of the clauses is not true for player 1’s strategy,

each player gets at most $1 - 1/2n$. This holds also for the expected equilibrium payoff, and the sum of the (identical) payoffs to the two players is at most $2 - 1/n$. If ϕ is satisfiable, then there is a correlated equilibrium with payoff 1 to both players, and payoff sum 2, as before. So the maximum payoff sum of a correlated equilibrium would answer whether ϕ is satisfiable. \square

The preceding proof does not apply to the EFCE concept where player 1 is not told his full strategy. Instead, the following defines an EFCE with payoff 1 to each player for any formula ϕ : With probability $1/n$, choose any of the n pure strategy pairs where player 2 chooses c_i for $1 \leq i \leq n$, and any literal in the i th clause of ϕ , and player 1 chooses “in” and a truth assignment that makes this literal true (with arbitrary assignments to the other variables). This is an EFCE because at his first information set, player 1 only receives the recommendation to play “in”, which is an optimal move for player 1 because each of player 2’s moves c_1, \dots, c_n has conditional probability $1/n$.

Papadimitriou and Roughgarden (2005) study correlated equilibria for compactly specified games, for example symmetric games of n players that have two strategies each, which can be specified by $2n$ payoffs rather than $n2^n$ payoffs when the game is given explicitly in strategic form. For such symmetric games, they describe the set of correlated equilibria by a polynomial number of constraints, which is similar in spirit to Theorem 3.9.

Papadimitriou (2005) shows how to compute *some* correlated equilibrium for other types of compactly specified games in polynomial time, even though computing a correlated equilibrium with maximum payoff sum, as considered in Theorems 3.10 and 3.11, is NP-hard. Extensive games can also be regarded as “compact specifications” of games in strategic form. It is an interesting open problem whether the approach by Papadimitriou (2005) can be applied to the computation of some EFCE, or even of some strategic-form correlated equilibrium, of an extensive game.

Acknowledgments

We thank Nimrod Megiddo, Roger Myerson, Andrés Perea, and Shmuel Zamir for stimulating early discussions, and seminar participants at Johns Hopkins, Paris, Stockholm, Stony Brook, Valencia, Waseda (Tokyo), and Yale for their comments.

References

- Aumann, R. J. (1974), Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
- Aumann, R. J. (1987), Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
- Cho, I.-K., and D. M. Kreps (1987), Signaling games and stable equilibria. *The Quarterly Journal of Economics* 102, 179–222.
- Chu, F. and J. Halpern (2001), On the NP-completeness of finding an optimal strategy in games with common payoffs. *International Journal of Game Theory* 30, 99–106.
- Conitzer, V., and Sandholm, T. (2003), Complexity results about Nash equilibria. In: *Proc. 18th International Joint Conference on Artificial Intelligence (IJCAI)*, 765–771.
- Cotter, K. (1991), Correlated equilibrium in games with type-dependent strategies. *Journal of Economic Theory* 54, 48–68.
- Dhillon, A., and J. F. Mertens (1996), Perfect correlated equilibria. *Journal of Economic Theory* 68, 279–302.
- Dodis, Y., S. Halevi, and T. Rabin (2000), A cryptographic solution to a game theoretic problem. In: *Proc. Crypto 2000*, Springer Lecture Notes in Computer Science, Vol. 1880, 112–130.
- Forges, F. (1986a), An approach to communication equilibria. *Econometrica* 54, 1375–1385.
- Forges, F. (1986b), Correlated equilibria in repeated games with lack of information on one side: a model with verifiable types. *International Journal of Game Theory* 15, 65–82.
- Forges, F. (1993), Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision* 35, 277–310.
- Forges, F. (2006), Correlated equilibrium in games with incomplete information revisited. Mimeo, University of Paris–Dauphine.
- Garey, M. R., and D. S. Johnson (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco.
- Gerardi, D. (2004), Unmediated communication in games with complete and incomplete information. *Journal of Economic Theory* 114, 104–131.
- Gibbons, R. (1992), *Game Theory for Applied Economists*. Princeton Univ. Press, Princeton, NJ.
- Gilboa, I. and E. Zemel (1989), Nash and correlated equilibria: some complexity considerations. *Games and Economic Behavior* 1, 80–93.
- Harsanyi, J. C. (1967), Games with incomplete information played by “Bayesian” players. Part I. The basic model. *Management Science* 14, 159–182.
- Koller, D. and N. Megiddo (1992), The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior* 4, 528–552.
- Koller, D., N. Megiddo, and B. von Stengel (1996), Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior* 14, 247–259.

- Kuhn, H. W. (1953), Extensive games and the problem of information. In: *Contributions to the Theory of Games II*, eds. H. W. Kuhn and A. W. Tucker, Annals of Mathematics Studies 28, Princeton Univ. Press, Princeton, 193–216.
- Moulin, H., and J.-P. Vial (1978), Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7, 201–221.
- Myerson, R. B. (1986), Multistage games with communication. *Econometrica* 54, 323–358.
- Papadimitriou, C. H. (1994), *Computational Complexity*. Addison-Wesley, Reading, Mass.
- Papadimitriou, C. H. (2005), Computing correlated equilibria in multiplayer games. In: *Proc. 37rd Annual ACM Symposium on the Theory of Computing (STOC)*, 49–56.
- Papadimitriou, C. H., and T. Roughgarden (2005), Computing equilibria in multi-player games. In: *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 82–91.
- Romanovskii, I. V. (1962), Reduction of a game with complete memory to a matrix game. *Doklady Akademii Nauk SSSR* 144, 62–64 [English translation: *Soviet Mathematics* 3, 678–681].
- Samuelson, L., and J. Zhang (1989), Correlated equilibria and mediated equilibria in games with incomplete information. Mimeo, Penn State University.
- Solan, E. (2001), Characterization of correlated equilibria in stochastic games. *International Journal of Game Theory* 30, 259–277.
- Spence, M. (1973), Job market signaling. *The Quarterly Journal of Economics* 87, 355–374.
- Urbano, A., and J. E. Vila (2002), Computational complexity and communication: coordination in two-player games. *Econometrica* 70, 1893–1927.
- von Stengel, B. (1996), Efficient computation of behavior strategies. *Games and Economic Behavior* 14, 220–246.
- von Stengel, B. (2001), Computational complexity of correlated equilibria for extensive games. Research Report LSE-CDAM-2001-03, London School of Economics.
- von Stengel, B., A. H. van den Elzen, and A. J. J. Talman (2002), Computing normal form perfect equilibria for extensive two-person games. *Econometrica* 70, 693–715.
- Young, H. P. (2004), *Strategic Learning and its Limits*. Oxford University Press, Oxford.