

Extensive-Form Correlated Equilibrium: Definition and Computational Complexity

Bernhard von Stengel

Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom,
stengel@maths.lse.ac.uk, <http://www.maths.lse.ac.uk/personal/stengel>

Françoise Forges

CEREMADE, University of Paris–Dauphine, 75775 Paris cedex 16, France,
francoise.forges@dauphine.fr, <http://www.dauphine.fr/edocif/forges/francoise.forges.htm>

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 (CE). Before the game starts, a correlation device generates a move for each information set. This move is recommended to the player only when the player reaches the information set. In two-player perfect-recall extensive games without chance moves, the set of EFCE can be described by a polynomial number of consistency and incentive constraints. Assuming P is not equal to NP, this is not possible for the set of CE, or if the game has chance moves.

Key words: correlated equilibrium; extensive game; polynomial-time computable

MSC2000 subject classification: Primary: 91A18; secondary: 91A05, 91A28, 68Q17

OR/MS subject classification: Primary: noncooperative games; secondary: computational complexity

History: Received March 22, 2006; revised September 19, 2007, and March 20, 2008.

1. Introduction. Aumann [1] defined the concept of correlated equilibrium (abbreviated as CE, also for the plural equilibria) 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 representation* of a CE, these signals are strategies that players are recommended to play.

This paper proposes a new concept of correlated equilibrium for extensive games, called extensive-form correlated equilibrium or EFCE. Like in a CE (which is defined in terms of the strategic form), the recommendations to the players are moves that are generated before the game starts. However, each recommended move is assumed to be in a “sealed envelope” and is only revealed to a player when he reaches the information set where he can make that move.

As recommendations become local in this way, players know less. Consequently, the set of EFCE outcomes is larger than the set of CE outcomes. However, an EFCE is more restrictive than an agent-form correlated equilibrium (AFCE). In the agent form of the game, moves are chosen by a separate agent for each information set of the player. In an EFCE, players remain in control of their future actions, which is important when they consider deviating from their recommended moves.

The EFCE is a natural definition of correlated equilibrium for extensive games with perfect recall as defined by Kuhn [25]. Earlier extensions of Aumann’s concept applied only to multistage games, including Bayesian games and stochastic games, which have a special time and information structure. These earlier approaches are discussed in §2.4.

The main motivation for the EFCE concept is computational. The algorithmic input is some description of the extensive game with its game tree, information sets, moves, chance probabilities and payoffs. Polynomial (or linear or exponential) size and time always refer to the size of this description. The strategic form of the extensive game has typically exponential size. Hence, there are also exponentially many linear constraints that define the set of strategic-form correlated equilibria. In this paper, we are interested in the set of all EFCE of the game, and prove the following result.

THEOREM 1.1. *For a two-player, perfect-recall extensive game without chance moves, the set of EFCE can be described by a system of linear equations and inequalities of polynomial size. For any solution to that system (which defines an EFCE), a pair of pure strategies containing the recommended moves can be sampled in polynomial time.*

This theorem is analogous to the description of the set of CE for a game in strategic form by incentive constraints. The incentive constraints compare any two strategies of a player, so their number is polynomial in the size of the strategic form. Consequently, for games given in strategic form, one can find in polynomial time a CE that maximizes the sum of payoffs to all players, which we call the problem MAXPAY-CE (which

we consider for various descriptions of games as input). In contrast, the problem MAXPAY-NE (finding a Nash equilibrium with maximum payoff sum) for games in strategic form is NP-hard (Gilboa and Zemel [17], Conitzer and Sandholm [6]; see Garey and Johnson [14] or Papadimitriou [30] for notions of computational complexity).

While CE are computationally easier than Nash equilibria for games in strategic form, this is not clear for games in extensive form, because their strategic form may be exponentially large. The following negative result confirms that, unless $P = NP$, the set of (strategic-form) CE does not have a polynomial-sized description.

THEOREM 1.2. *For two-player, perfect-recall extensive games without chance moves, the problem MAXPAY-CE is NP-hard.*

Theorem 1.1 implies that the problem MAXPAY-EFCE (finding an EFCE with maximum payoff sum) can be solved in polynomial time for two-player, perfect-recall games without chance moves. Interestingly, that problem becomes NP-hard when chance moves are allowed, as stated in the following theorem; a closely related result has been shown earlier by Chu and Halpern [5].

THEOREM 1.3. *For two-player, perfect-recall extensive games with chance moves, the problems MAXPAY-NE, MAXPAY-CE, MAXPAY-AFCE, and MAXPAY-EFCE are NP-hard.*

For zero-sum, two-player extensive games with perfect recall, a Nash equilibrium can be found in polynomial time, as shown by Romanovskiĭ [34], Koller and Megiddo [23], and von Stengel [40]. These methods (most explicitly von Stengel [40]) use the sequence form of an extensive game where mixed strategies are replaced by behavior strategies, by Kuhn's theorem [25]. A behavior strategy is represented 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. The sequence form applies also to games with chance moves.

Recently, Hansen et al. [20] have found another case where the introduction of chance moves marks the transition from polynomial-time solvable to NP-hard problems. They give a linear-time algorithm that decides if a two-player zero-sum extensive game with perfect recall and no chance moves has a pure-strategy equilibrium. Blair et al. [3] have shown that this problem is NP-hard if chance moves are allowed. (For games with imperfect recall, even if they are zero-sum and have no chance moves, it is NP-hard to find the unique Nash or correlated equilibrium payoff; see Koller and Megiddo [23, p. 534].)

For two-player perfect-recall games without chance moves, the set of EFCE has a polynomial-sized description for the following reason. An EFCE describes correlations of moves between information sets of the two players, rather than correlations of entire strategies as in a CE. This is similar to using behavior strategies rather than mixed strategies in a Nash equilibrium. The recommended moves at an information set depend on what has been recommended to the other player (this is stated as “sampling a pure-strategy pair” in Theorem 1.1 and is proved in Theorem 3.9). Consider some information set, say k of player 2, where a move is to be recommended. Perfect recall and the absence of chance moves imply that previous recommendations to the other player must define a sequence of moves, of which there is only a linear number (see Figure 7). Hence, there are only a few conditional distributions for generating the move at k . In contrast, a chance move, when learned by a player, may give rise to parallel information sets (which are preceded by the same own earlier moves of the player; see von Stengel [40, Definition 4.3]). The number of move combinations at parallel information sets may grow exponentially, and each of them may produce a different conditional distribution for the recommended move. This applies in general for CE, with possibly exponentially many recommended strategies and corresponding conditional distributions.

The polynomially many constraints that describe the set of EFCE according to Theorem 1.1 extend, in a relatively natural way, 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. These constraints are valid even when the game has chance moves or more than two players, but they do not characterize the set of EFCE in those cases (otherwise, Theorem 1.3 would imply $P = NP$). The constraints do suffice for two-player games without chance moves, which needs careful reasoning because many subtleties arise; for example, there may be cycles (of length four or more) in the possible temporal order of information sets, as Figure 6 demonstrates.

Papadimitriou and Roughgarden [32] study the computation of CE for various compactly represented games such as certain graphical games, congestion games, and others. For anonymous games, they give an explicit, polynomial-sized description of the set of CE, and (in Papadimitriou and Roughgarden [33]) a way to sample a pure strategy profile from a CE described in that way, analogous to our Theorem 1.1. (The players in an

anonymous game have equal strategy sets, and a player's payoff depends only on how many, but not which, other players choose a particular strategy.)

We consider the problems MAXPAY-CE and MAXPAY-EFCE to see whether the set of CE or EFCE can be described by a polynomial number of linear constraints. Similar to our Theorem 1.3, Papadimitriou [31] and Papadimitriou and Roughgarden [33] prove that for many compactly represented games, the problem MAXPAY-CE is NP-hard. However, their main result (Theorem 3.12 below) states that one CE can often be found in polynomial time, which shows that finding a CE is usually computationally simpler than payoff maximization. In §3.8 we confirm this observation by explicitly constructing Nash equilibria for the games used in the NP-hardness proofs of Theorems 1.2 and 1.3. Moreover, as a corollary to the result of Papadimitriou [31], our Proposition 3.13 states that for any extensive game, an AFCE can be found in polynomial time. This holds because the agent form, unlike the strategic form, has few strategies per player. The computational complexity of finding one CE or EFCE for a general extensive game is an open problem.

Sections 2 and 3 of this paper treat the conceptual and computational aspects of EFCE, respectively; an overview is given at the beginning of each section. Section 4 discusses open problems.

2. The EFCE concept. This section presents the basic properties of the EFCE. In §2.1, we define the solution concept in canonical form. As we explain, this can be done without loss of generality. In §2.2, we show that an EFCE can always be defined with a correlation device that generates profiles of reduced strategies. Section 2.3 discusses a signaling game with costless signals, in which an EFCE is “type-revealing” while all CE are nonrevealing. In §2.4, we compare the EFCE with other extensions of the CE, which have been defined for games with special time or information structures.

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 nonterminal 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 by H_i . The set of choices or moves at an information set h is denoted by 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 CE if no player can gain by unilaterally deviating from the recommended strategy, given his posterior on the recommendations to the other players (see Aumann [1]). We define an EFCE 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 other players follow their recommendations. When a player considers a deviation from a recommended move, he may choose any moves at his subsequent information sets. This distinguishes the EFCE from the AFCE, where each move is optimal assuming that the behavior at all other information sets is fixed (see §2.4).

The above description of an EFCE 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 the CE, this can be assumed without loss of generality (see Forges [10]).

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. (Defined in this way, the reduced strategic form only depends on the game tree structure and not on the payoffs.) In our characterization of EFCE in Theorem 1.1, it is not possible to specify move recommendations for unreachable information sets, so the device can only generate reduced strategy pairs. As shown in this section, this causes no loss of generality.

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 redefine a correlation device to be a probability distribution on Σ^* . Any CE that is specified using the unreduced strategic form can be considered as a CE 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 CE 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 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 [38], Cho and Kreps [4], and Gibbons [16, §4.2], but with costless signals (such games are often referred to as “sender-receiver” games). Player 1, a student, is with equal probability of a good type G or bad type B . He applies for a summer research job with a professor, player 2. Player 1 sends a costless signal X or Y (denoted as move X_G or Y_G for the good type, and as X_B and Y_B for the bad type). The professor can distinguish the signals X and Y but not the type of player 1, as shown by her two information sets. She can either let the student work with her (l_X or l_Y), which gives the payoff pair (4, 10) for G , and (6, 0) for B , or refuse to work with the student (r_X or r_Y), which for either type gives the payoff pair (0, 6).

The CE of this game are found as follows. Figure 2 shows the strategic form and the possible CE probabilities a, a', a'', b, b', \dots , where player 2’s strategy $l_X l_Y$ is strictly dominated by $r_X r_Y$ and never played. The incentive constraints for player 1 imply that $a \geq a'$ (by comparing $X_G X_B$ with any other row), and similarly $d' \geq d$. Comparing $X_G Y_B$ with $X_G X_B$ (respectively, $Y_G Y_B$) implies $b' \geq b$ ($b \geq b'$), so $b = b'$; similarly, $c = c'$. Intuitively, this means that player 2 must not give preference to either signal because otherwise the bad type would switch to that signal. Then the incentive constraints where player 2’s strategies $l_X r_Y$ and $r_X l_Y$ are compared with $r_X r_Y$ state

$$5a + 8b + 3c \geq 6a + 6b + 6c, \quad 3b + 8c + 5d' \geq 6b + 6c + 6d',$$




FIGURE 1. Signaling game with costless signals (X or Y) for player 1.

which when added give $5a + 11b + 11c + 5d' \geq 6a + 12b + 12c + 6d'$ or $0 \geq a + b + c + d'$ and thus $a = b = c = d' = 0 = a' = d$. Any CE is therefore a Nash equilibrium where player 1 plays the mixed strategy (a'', b'', c'', d'') and player 2 plays $r_X r_Y$. The remaining incentive constraints for a'', b'', c'', d'' mean that player 1 must not give player 2 any incentive to accept him (l_X or l_Y) by making the conditional probability for G too high relative to B when she receives signal X or Y .

However, there is an EFCE with better payoff to both players compared to the outcome with payoff pair $(0, 6)$: A signal X_G or Y_G is chosen with equal probability for type G , and player 2 is told to accept when receiving the chosen signal and to refuse when receiving the other signal (so X_G and $l_X r_Y$ are perfectly correlated, as well as Y_G and $r_X l_Y$). 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 of $20/3$ for choosing l is higher than 6 when she chooses r . When she receives the wrong signal, it comes from B with certainty, and then the best reply is certainly r with payoff 6 . The expected payoffs in this EFCE are 3.5 to player 1 and 6.5 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 $(2 + 3/M, 8 - 3/M)$.

In the terminology of signaling games, any Nash or correlated equilibrium is the described “pooling equilibrium” with payoff pair $(0, 6)$. 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 CE have been proposed to take into account the dynamic structure of specific classes of games, namely Bayesian games and multistage games.

	2				
1		$l_X l_Y$	$l_X r_Y$	$r_X l_Y$	$r_X r_Y$
$X_G X_B$		5	5	6	6
$X_G Y_B$		5	8	3	6
$Y_G X_B$		5	3	8	6
$Y_G Y_B$		5	6	5	6
		5	0	5	0

	2				
1		$l_X l_Y$	$l_X r_Y$	$r_X l_Y$	$r_X r_Y$
$X_G X_B$		0	a	a'	a''
$X_G Y_B$		0	b	b'	b''
$Y_G X_B$		0	c	c'	c''
$Y_G Y_B$		0	d	d'	d''

FIGURE 2. Left: Strategic form of the game in Figure 1. Right: Correlated equilibrium probabilities.

In a Bayesian game, every player has a type which can be represented by an information set. Players move simultaneously and only once, so that an AFCE is the same as an EFCE. For Bayesian games, AFCE have been studied by Forges [11, 12, 13], Samuelson and Zhang [35], and Cotter [7].

In general extensive-form games, any EFCE is an AFCE, by giving arbitrary recommendations at unreachable information sets that in an EFCE are left unspecified (see §2.2). However, the set of AFCE 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 AFCE but not an EFCE.

In multistage games, the best known extension of the CE is the communication equilibrium introduced by Myerson [28] and Forges [10]. This solution concept differs from the EFCE, because the players can send inputs to the device, which they cannot do in an EFCE.

Like the communication equilibrium, the autonomous correlated equilibrium (Forges [10]) applies to multi-stage games. However, the players cannot make any inputs to the device. They still receive outputs at every stage. 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.¹ 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 in §2.3, where an autonomous correlated equilibrium is the same as a CE. The inclusion may also be strict for two-player games without chance moves, which we consider later (see the example in §3.3).

Solan [37] defines the concept of general communication equilibrium for stochastic games, where the device knows the game state and all past moves. He proves that this concept is outcome equivalent to the autonomous correlated equilibrium. Because any autonomous correlated equilibrium outcome is an EFCE outcome, which, by definition, is a general communication equilibrium outcome, these concepts coincide for stochastic games.

Kamien et al. [22] and Zamir et al. [44] 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 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 [27] proposed a “simple extension” of Aumann’s [1] correlated equilibrium that is completely different from the ones reviewed above. Like the CE, their solution concept, which is also referred to as a coarse correlated equilibrium (Young [43]), 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 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 state 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.

3. Computational complexity. So far, we have argued that the EFCE is a “natural” concept for games in extensive form. This section deals with computational aspects of the EFCE. The main technical work is to prove Theorem 1.1, which concerns two-player games without chance moves. In §3.1, we review 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 consistency constraints for joint probabilities of pairs of sequences, which define what we call a correlation plan. Section 3.3 gives an example that illustrates the use of the consistency constraints.

¹ In Forges [10, p. 1378], a correlated equilibrium based on an autonomous device is called an “extensive-form correlated equilibrium,” but this is now typically referred to as an “autonomous correlated equilibrium.” We suggest now using “EFCE” in our sense.

In general, the consistency constraints apply only to mutually relevant information sets that share a path in the game tree, as explained in §3.4. That section also describes the special structure of information sets in two-player perfect-recall games without chance moves, and defines the concept of a reference sequence, which is used to generate move recommendations. Based on these technical preliminaries, §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 §3.6.

In §3.7, we prove the hardness results of Theorems 1.3 and 1.2. These hardness results do not apply to the problem of finding one CE, which is the topic of §3.8.

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. Because 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 §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 the set of 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 [40], Koller et al. [24], or von Stengel et al. [42]). 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

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

The reason is that Equations (4) hold when a player uses a behavior strategy, in particular a pure strategy, and therefore also for any mixed strategy, because the equations are preserved when taking convex combinations.

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 $x(\sigma_h c)/x(\sigma_h)$ if $x(\sigma_h) > 0$ and arbitrarily if $x(\sigma_h) = 0$. The probability of reaching any node of the game tree depends only on the probabilities for the players' move sequences defined by the path to the node. So, via x , every mixed strategy has a realization-equivalent behavior strategy, as stated by Kuhn [25]. This canonical proof of Kuhn's theorem (essentially due to Selten [36]) works for any number of players. 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.² 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.

² A class of games with exponentially large reduced strategic form is described by von Stengel et al. [42].

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 ,

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

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.

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, as far as this probability is relevant. These probabilities $z(\sigma, \tau)$ define what we call a correlation plan for the game.

As a tentative definition, given in full in Definition 3.8 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 recommendation given to him is optimal; see the example in §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 $k \in H_2$,

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

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 3 shows on the left a correlation plan defined in this manner for the game in Figure 1. Because both players move only once, every nonempty sequence is just a move. The correlation plan on the left in Figure 3 arises from the pure strategy pair $(X_G Y_B, l_X l_Y)$.

Figure 3 shows on the right 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,

	\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

	\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. Left: Correlation plan representing the pure strategy pair $(X_G Y_B, l_X l_Y)$. Right: Distribution on sequence pairs that is “locally” (in each row and column) consistent, but which is not a convex combination of pure strategy pairs.

they cannot be obtained as a convex combination of pure strategy pairs like the pure strategy pair on the left in Figure 3. 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 be recommended to player 2. This pure strategy pair is thus $(X_G Y_B, l_X l_Y)$ as in the left picture of Figure 3, but that pair also selects (Y_B, l_Y) , which is not possible according to the right picture. This shows that (7) and (8) do not suffice to characterize the convex hull of pure strategy profiles. For games with chance moves, Theorem 1.3 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. The left picture in Figure 4 shows 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 §2.3 shows that there is only one outcome in a strategic-form or autonomous correlated equilibrium, or even communication equilibrium (see §2.4), which is nonrevealing.

Figure 4 shows on the right 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 §3.5 below. We consider only the generation of moves, and not any incentive constraints (treated in §3.6), which are in fact violated for this z .

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 4, $z(G, \tau) = z(B, \tau) = 1/2$. 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 4. 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




FIGURE 4. Left: Game similar to Figure 1 with a move by player 1 instead of chance. Right: A possible distribution on sequence pairs for this game.

earlier move G . Because it suffices to generate only a reduced strategy of player 1 as explained in §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 4 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 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 player 2 to reach k , described in generality in §3.5. According to Figure 4, $z(\sigma, l_Y) = 1/4$ 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 4 describe a correlation between the moves at pairs of information sets of player 1 and player 2, with nonzero entries—as in the right picture of Figure 3. However, unlike in that picture, 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 3.9.

DEFINITION 3.1. Let u and v be two nodes in an extensive game, where u is on the path from the root to v . Then u is called an ancestor of v (or earlier than v), and v is said to be later than u . If h and k are information sets (possibly of the same player) with $u \in h$ and $v \in k$, then h is said to precede k , and h and k are called connected (sharing a path).

LEMMA 3.2. Consider an extensive game with perfect recall.

- (a) If h' , h , k are information sets so that h' and h belong to the same player, h' precedes h , and h precedes k , then h' precedes k .
- (b) Restricted to the information sets of a single player, “precedes” is an irreflexive and transitive relation.

PROOF. For (a), some node u in h is earlier than some node v in k , and some node in h has an earlier node in h' , so by perfect recall all nodes of h have an earlier node in h' , including u , which implies that h' precedes v .

For (b), it is easy to see that no two nodes in an information set share a path in the tree, so “precedes” is irreflexive, and by (a) it is transitive. \square

By Lemma 3.2(b), any set H' of information sets of a single player is partially ordered. We call an information set h in H' maximal if it is not preceded by any other information set in H' .

The following lemma states that for two-player games without chance moves, “precedes” is antisymmetric even for information sets of different players (which is easily seen to be false if there are chance moves or a third player).

LEMMA 3.3. 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' . Let w be the last common ancestor of u and v' . If $w = u$ (or $w = v'$), then two nodes in h (respectively, k) share a path, which is not possible with perfect recall. Otherwise, Figure 5 shows that perfect recall is violated if h and k belong to the same player, or else for the player who moves at w . \square




FIGURE 5. Proof of Lemma 3.3.

For two-player games without chance moves, “precedes” is in general not a transitive relation on all information sets, and it may even have cycles, as shown in Figure 6.

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. The following lemma is illustrated by Figure 7.

LEMMA 3.4. Consider a two-player perfect-recall extensive game without chance moves, and let $h, h' \in H_1$ and $k \in H_2$ so that h and h' both precede k but 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. Consider two paths from the root to k that intersect h and h' , respectively. These paths split at some node u because h and h' are not connected (see Figure 7). 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) .




FIGURE 6. Extensive game of two players with perfect recall where the information sets h, k, h', k' form a cycle with respect to the “precedes” relation.




FIGURE 7. Proof of Lemma 3.4.

DEFINITION 3.5. 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 3.5, 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. In the example in Figure 8, player 1 has information sets h and h' , and player 2 has k and k' . The sets of sequences of players 1 and 2 are $S_1 = \{\emptyset, b, c, cb', cc'\}$ and $S_2 = \{\emptyset, d, e, dd', de'\}$. The two information sets h' and k' are not connected (all others are), so the sequence pairs (cb', dd') , (cb', de') , (cc', dd') , and (cc', de') are irrelevant. We will not specify probabilities $z(\sigma, \tau)$ for such irrelevant sequence pairs (σ, τ) , because correlating the moves at the two information sets h' and k' would not matter. Moreover, such an over-specified correlation plan z would be hard to translate into a generation of moves.

We do 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 Equations (7) to relevant sequence pairs.

LEMMA 3.6. Consider a two-player extensive game without chance moves and with perfect recall. Assume that the pair $(\sigma, \tau) \in S_1 \times S_2$ of sequences is relevant, and that $\sigma' \in S_1$ is a prefix of σ and that $\tau' \in S_2$ 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 players 1 and 2, respectively. Because h and k are connected, assume that h precedes k ; the case where 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 3.2(a).

Similarly, (σ', τ') is relevant, which only needs to be shown for $k' \neq k$: Then k' and h' precede k , with some node v in k having an earlier node u in h' . Because some node in k has an earlier node in k' , node v also has an earlier node in k' , which is therefore on the path from the root to v which also contains u . This shows that k' and h' are connected. \square




FIGURE 8. Example demonstrating relevant sequence pairs, reference sequences, and the proof of Theorem 3.9.

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 §3.3. A reference sequence τ of player 2, for example, defines a “column” of z (as in Figure 4) 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 3.7. 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 $\tau_{k'} = \tau$ so that k' 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 $\sigma_{h'} = \sigma$ so that h' 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 3.7(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 precede h .

In Figure 8, any sequence of player 2 has the reference sequence \emptyset of player 1. 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 .

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; each time, these moves define a suitable reference sequence for the next generated move.

DEFINITION 3.8. 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 3.9. 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 $(\sigma_h c, \tau)$ and $(\sigma, \tau_k d)$ are relevant, for any $c \in C_h$ and $d \in C_k$. A corresponding probability distribution μ on Σ^* in Definition 3.8 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. 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 so that for all information sets k that precede 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 .

We claim that such an information set h always exists. Initially, when p_1 and p_2 are everywhere undefined, h is the information set containing the root of the game tree. In general, let H'_1 be the set of information sets of player 1 where p_1 is not yet defined, and which are not unreachable due to an earlier move in p_1 , and let H'_2 be the analogous set for player 2. The sets H'_1 and H'_2 are partially ordered by “precedes” (see Lemma 3.2(b)). Consider a maximal information set in H'_1 . It may be preceded by a information set in H'_2 (otherwise we are done) that by Lemma 3.2(a) is maximal in H'_2 , which in turn may be preceded by another maximal information set in H'_1 , and so on, but we claim that there cannot be a cycle of such sets. Otherwise, let the cycle begin with h preceding k and k preceding h' , with $h, h' \in H'_1$ and $k \in H'_2$; the cycle must have length at least four by Lemma 3.3, because h' cannot precede h because otherwise it would precede k by Lemma 3.2(a). Let $u \in h$ and $v, v' \in k$ so that u is earlier than v , and v' is earlier than some node in h' , as in the example in Figure 6. Let w be the last common ancestor of u and v' . Then (see Figure 6) w must belong to an information set of player 1 because otherwise player 2 would not have perfect recall. That information set is not h , because otherwise h would precede h' , but h' is maximal in H'_1 , so w belongs to an information set where the move under p_1 is already specified, or not specified because the information set is unreachable under p_1 , but then h is unreachable as well. According to that move at w , at least one of h or h' is unreachable, contradicting the definition of H'_1 . So, as claimed, there is an information set h in $H'_1 \cup H'_2$ not preceded by any other such set. Assume that h belongs to player 1; the case for player 2 is analogous. Because h is not unreachable when playing p_1 (where the move $p_1(h)$ stays undefined), 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 3.4: 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. Otherwise, if there are k and k' in K that are not connected, we obtain a contradiction as follows: Lemma 3.4 (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 , so 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 (minimal) 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 that (6) holds, 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 3.5 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 3.7 and Lemma 3.6, 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 $\sigma_h c$ 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 the “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 3.4. 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, $\mu(\sigma_h c, \tau_k d) = z(\sigma_h c, \tau_k d)$. If all the moves in $\sigma_h c$ and τ_k are generated, then exactly one of the moves in C_k is generated. This implies

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

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

The example in Figure 8 demonstrates the two-part induction in the preceding proof. Recall that the sequences b and c of player 1 have possible reference sequences dd' , de' , or e , whereas the sequences cb' and cc' , which are longer than c , have possibly shorter reference sequences d or e . That is, reference sequences can be “nonmonotonic,” in the sense that later information sets (here h' , preceded by h) can have shorter reference sequences. 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.

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 3.9. 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, eventually gives $u(\emptyset) = \sum_{\sigma, \tau} z(\sigma, \tau) a(\sigma, \tau)$ for the empty sequence. 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 give a payoff directly 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 CE 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 because correlation is allowed. 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, as here, the computation of the player's optimal payoff by dynamic programming, as described by von Stengel [40, p. 239].

In summary, in a two-player, perfect-recall extensive game without chance moves, a correlation plan z as in Theorem 3.9 that fulfills for both players the incentive constraints (13), (14), and (15) defines an EFCE. This proves Theorem 1.1.

We conclude with 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 anything by ignoring private information. The simple proof of this fact is analogous to the observation that any CE is a coarse correlated equilibrium as defined by Moulin and Vial [27].

3.7. Hardness results. Theorem 1.1 shows that the set of EFCE for two-player games without chance moves has a compact description, with a polynomial number of linear equations and inequalities. In this section, we first prove Theorem 1.3. This theorem implies that a compact description of the set of CE, AFCE, or EFCE cannot be expected when chance moves are admitted (which is already known for Nash equilibria). Otherwise, one could maximize in polynomial time the sum (or any linear function) of the expected payoffs to the players over the respective set, which would imply $P = NP$. Chu and Halpern [5, Theorem 3.1] have given a very similar NP-hardness proof for finding optimal play in a “possible worlds” model. They also give a reduction from SAT where player 1 chooses a truth assignment and player 2 a literal to be made true in a clause, but the chance move picks a literal rather than a clause, and player 2 has to match the chosen “possible world.”

PROOF OF THEOREM 1.3. We give a reduction from SAT that applies to all four problems. An instance of SAT is a boolean formula ϕ in conjunctive normal form. If ϕ has n clauses, the reduction gives an extensive two-player game $\Gamma(\phi)$ with perfect recall that has size proportional to the size of ϕ , with identical payoffs to the players. The game has a pure strategy profile (which is a Nash equilibrium) with payoff 1 for each player if ϕ is satisfiable, and payoff at most $1 - 1/n$ if ϕ is not satisfiable; the payoff sum is that payoff times the number of players (which is higher in the agent form). This also applies to mixed-strategy Nash equilibria, CE, AFCE, and EFCE, because they are convex combinations of pure strategy profiles. Finding any such equilibrium with maximum payoff sum is therefore NP-hard.

We construct $\Gamma(\phi)$ as follows; Figure 9 shows an example. When referring to the two players, we mean the original players in the case of the agent form. Player 2 has n decision nodes in singleton information sets, which correspond to the clauses of ϕ . Player 1 has as many decision nodes as there are literals (negated or nonnegated variables) in ϕ , and the game has twice as many terminal nodes as decision nodes of player 1. 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




FIGURE 9. Extensive game $\Gamma(\phi)$ for the SAT instance $\phi = (x) \wedge (\neg x \vee y) \wedge (\neg x \vee \neg y)$, which is not satisfiable. Chance chooses a clause (this part, above the dotted line, is replaced in Figure 10). Player 2 picks a literal within the clause, and player 1 chooses for each variable the literal to be made true. The payoffs 0 and 1 are the same for both players.

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.

The game $\Gamma(\phi)$ has a pair of pure strategies for the two players (or a strategy profile for the $m + n$ players in the agent form) with payoff 1 for each player 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 1.3 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 then has 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 SAT formula is satisfiable.

In the game $\Gamma(\phi)$ constructed in the proof of Theorem 1.3, both players have an exponential number of pure strategies. If the chance move is replaced by a move of player 2, then the number of reduced strategies of player 2 equals the number of literals of ϕ . However, the strategic form is still exponential.

Our second result, stated as Theorem 1.2 in the introduction, uses such a construction to show that even for two-player games without chance moves, it is NP-hard to find a CE with maximum payoff sum. The constructed game has a first stage given by a high-stakes zero-sum game with a unique mixed equilibrium that induces player 2 to randomize, which replaces the initial chance move. Goldberg and Papadimitriou [18] have introduced this game of generalized matching pennies for a similar purpose. One can also use the “generalized rock-scissors-paper game” by von Stengel [41], with payoffs $(j - i) \bmod n$ (suitably scaled) for a_{ij} in (17) below, but the explicit construction of a Nash equilibrium in Proposition 3.11 below is more complicated for that game.

PROOF OF THEOREM 1.2. The proof is again by reduction from SAT, by constructing a game $\Gamma'(\phi)$ from an SAT formula ϕ ; let ϕ have n clauses. The game is derived from the extensive game $\Gamma(\phi)$ constructed in the proof of Theorem 1.3, but with the initial chance move replaced by a decision node of player 2 with n moves called c_1, \dots, c_n (see Figure 10). 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

$$a_{ij} = \begin{cases} 2 - 2n & \text{if } i = j, \\ 2 & \text{if } i \neq j. \end{cases} \quad (17)$$

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 his outside options o_1, \dots, o_n , then the game is a zero-sum game with payoffs a_{ij} as in (17). Players 2 and 1 are here pursuer and evader in a game of generalized matching pennies (Goldberg and Papadimitriou [18, Definition 1]), scaled so that it has value zero, with the uniform mixed strategy as the optimal




FIGURE 10. Preplay of a zero-sum pursuit-evasion game between players 2 and player 1 that induces player 2 to completely mix between her choices c_1, \dots, c_n . After the dotted line following move in of player 1, the game $\Gamma'(\phi)$ continues as in Figure 9.

strategy for each player. If player 2 chooses move c_j with probability $q < 1/2n$, then player 1 can respond with move o_j and get payoff $2 - 2nq > 1$, which is larger than any payoff following move in . Consequently, in any CE of $\Gamma'(\phi)$ where a pure strategy involving in is recommended to player 1, the conditional probability for each move c_1, \dots, c_n of player 2 is at least $1/2n$ because otherwise player 1 would deviate to one of his outside options.

Consider a CE of $\Gamma'(\phi)$, and suppose that ϕ is not satisfiable. Any recommendation to the players where player 1 is recommended an outside option contributes payoff sum zero. When player 1 is told a pure strategy where he plays in , the payoff sum is at most $2 - 1/n$ because each clause has probability at least $1/2n$ and at least one of the clauses is false. So the expected payoff sum is at most $2 - 1/n$. If ϕ is satisfiable, then player 2 can play c_1, \dots, c_n uniformly, player 1 chooses in , and $\Gamma'(\phi)$ has a CE with payoff 1 to each player and payoff sum 2, as before. So solving MAXPAY-CE for $\Gamma'(\phi)$ 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 of $\Gamma'(\phi)$ 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 $1/n$.

3.8. Finding one correlated equilibrium. The hardness results in Theorems 1.3 and 1.2 apply only to finding a correlated equilibrium with maximum payoff sum. Finding one correlated equilibrium is computationally easier than payoff maximization. To demonstrate this, we first give explicit constructions of Nash equilibria (which are also CE, AFCE, and EFCE) for the games $\Gamma(\phi)$ and $\Gamma'(\phi)$ in the proofs of the above NP-hardness theorems. Then, we apply the central result of Papadimitriou [31] to finding one AFCE.

PROPOSITION 3.10. *Given an SAT formula ϕ , a Nash equilibrium of the game $\Gamma(\phi)$ in the proof of Theorem 1.3 can be found in polynomial time.*

PROOF. The following straightforward algorithm produces a pure-strategy Nash equilibrium of $\Gamma(\phi)$. For initialization, declare all clauses of ϕ as active. Consider the variables of ϕ in some arbitrary order. Each variable x occurs among the active clauses (possibly none) as a positive literal x or negative literal $\neg x$. At the information set of $\Gamma(\phi)$ for x , let player 1 choose the truth value for x that satisfies the majority of active clauses (ties broken arbitrarily), for example false (so that $\neg x$ is true) in Figure 9. Then let player 2 choose that literal $\neg x$ for each clause in that majority of clauses, and let these clauses become inactive.

After all variables of ϕ have been assigned truth values in this way, all remaining active clauses are unsatisfied. Let player 2 choose an arbitrary literal from each of those clauses.

This defines a Nash equilibrium of $\Gamma(\phi)$: Player 2 cannot improve her payoff by changing any move, and neither can player 1 because at each of his information sets (variables of ϕ), at least as many moves of player 2 (from the clauses) lead to the high payoff 1 as to the low payoff 0. \square

PROPOSITION 3.11. *Given an SAT formula ϕ , a Nash equilibrium of the game $\Gamma'(\phi)$ in the proof of Theorem 1.2 can be found in polynomial time.*

PROOF. We extend the Nash equilibrium of $\Gamma(\phi)$ given in the proof of Proposition 3.10 with suitable mixed strategies of the initial zero-sum game of $\Gamma'(\phi)$. Namely, let S be the subset of $\{1, \dots, n\}$ indicating the satisfied clauses, and $s = |S|$. Let player 2 choose c_j with probability $2/(n+s)$ if $j \in S$ and with probability $1/(n+s)$ otherwise. Because the satisfied clauses are thereby given higher weight than the unsatisfied clauses, the moves of player 1 that choose the truth values as in $\Gamma(\phi)$, now following his move in , are still optimal (for this we need that player 1 is the evader and player 2 the pursuer in the initial generalized matching pennies game).

With these probabilities for c_1, \dots, c_n , the expected payoff to player 1, according to (17), is $2 - 4n/(n+s)$ when he chooses o_i and $i \in S$, and higher, namely $2 - 2n/(n+s)$, when he chooses o_i and $i \notin S$, and the same, that is, $2s/(n+s)$, when he chooses in . So player 1 can mix between o_i for the unsatisfied clauses i and in . He chooses o_i with probability $1/(3n-s)$ if $i \notin S$ (and zero if $i \in S$), and in with probability $2n/(3n-s)$. The expected payoff to player 2 following any initial move c_j is then $(-2(n-s) + 2n)/(3n-s)$, that is, $2s/(3n-s)$. (This works for any $0 \leq s \leq n$.) \square

The following result of Papadimitriou [31] is of interest for compactly represented games, where the description of the game is much smaller than its strategic form.

THEOREM 3.12 (PAPADIMITRIOU [31]). *Consider a game G in some description so that the number of players and the number of strategies per player is polynomial in the size of the description, and so that the expected*

payoff for any product distribution (mixed strategy profile) can be computed in polynomial time. Then one can compute in polynomial time a CE of G that is a convex combination of polynomially many product distributions.

An extensive game can be viewed as a compact representation of its strategic form or agent form.

PROPOSITION 3.13. *For any extensive game, an AFCE can be found in polynomial time.*

PROOF. In the agent form of the game, every information set h is represented by a separate player with the moves at h as his strategies, which are polynomial in number in the size of the game. The expected payoffs for a product distribution are found in polynomial time by summing over all leaves of the game tree. Thus, by Theorem 3.12, a CE of the agent form can be found in polynomial time. \square

Theorem 3.12 cannot be applied directly to the problem of finding one CE or EFCE of an extensive game because a player may have exponentially many strategies.

In contrast to finding one CE, the problem MAXPAY-CE is NP-hard for most types of compactly represented games (see Papadimitriou [31], Papadimitriou and Roughgarden [33]). These results are similar to our Theorem 1.3.

4. Discussion and open problems. For an extensive game, a CE generalizes a Nash equilibrium in mixed strategies, by allowing for correlations of strategies. An EFCE recommends moves rather than strategies to the players, so an EFCE can be seen as the correlated counterpart to a Nash equilibrium in behavior strategies. The EFCE is therefore closer in spirit to the dynamic description of the game by a tree than the CE. At the same time, the correlation device does not have additional power in the sense of observing the game state, because it generates its recommendations at the beginning of the game. In addition, the EFCE concept is computationally tractable for two-player games without chance moves.

Does the EFCE concept reflect “common knowledge of rationality for extensive games with Bayesian players,” in analogy to Aumann’s [2] 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, in particular those that are not sequential. The reason is that when, say, player 1 considers deviating from a recommended move, he assumes that player 2 follows her recommendations (which have already been generated), even if player 2 can conclude that player 1 has deviated. On the other hand, the EFCE seems easily extendable to subgames and subgame perfection. In further research, one could investigate refinements of the EFCE in the same way as, for example, Dhillon and Mertens [8] and Gerardi [15, p. 117] did for the CE.

For implementing an EFCE, the recommended moves have to be 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 the information sets themselves describe the rules of the game. “Sealed envelopes” should be implementable by cryptographic techniques. Dodis et al. [9], Urbano and Vila [39], and Lepinski et al. [26] have shown how to use cryptography to replace the mediator in a CE.

By Proposition 3.13, the AFCE seems more suitable than the EFCE if one is only interested in finding one equilibrium. However, the underlying Theorem 3.12 is based on the ellipsoid algorithm (see Grötschel et al. [19]), and is not very practical. The agent form has exponentially many strategy profiles. For a direct comparison of EFCE and AFCE, is there a compact description of the set of AFCE for two-player games without chance moves? How difficult is the problem MAXPAY-AFCE for these games?

A related open problem is the complexity of finding one EFCE or CE for a general extensive game with perfect recall. One possible approach, following Papadimitriou’s proof of Theorem 3.12, would be to extend the existence proof for CE by Hart and Schmeidler [21] or Nau and McCardle [29] to extensive games, where one would need to find a suitable compact encoding of strategies.

Acknowledgments. We thank Nimrod Megiddo, Roger Myerson, Andrés Perea, and Shmuel Zamir for stimulating early discussions. The paper benefited from comments of seminar participants at Dagstuhl (Germany), Johns Hopkins, Paris, Stockholm, Stony Brook, Valencia, Waseda (Tokyo), and Yale. We thank Eilon Solan, an associate editor, and two referees for their helpful suggestions.

References

- [1] Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *J. Math. Econom.* **1** 67–96.
- [2] Aumann, R. J. 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55** 1–18.

- [3] Blair, J. R. S., D. Mutchler, M. van Lent. 1996. Perfect recall and pruning in games with imperfect information. *Comput. Intell.* **12** 131–154.
- [4] Cho, I.-K., D. M. Kreps. 1987. Signaling games and stable equilibria. *Quart. J. Econom.* **102** 179–221.
- [5] Chu, F., J. Halpern. 2001. On the NP-completeness of finding an optimal strategy in games with common payoffs. *Internat. J. Game Theory* **30** 99–106.
- [6] Conitzer, V., T. Sandholm. 2003. Complexity results about Nash equilibria. *Proc. 18th Internat. Joint Conf. Artificial Intelligence (IJCAI)* 765–771.
- [7] Cotter, K. D. 1994. Type correlated equilibria for games with payoff uncertainty. *Econom. Theory* **4** 617–627.
- [8] Dhillon, A., J. F. Mertens. 1996. Perfect correlated equilibria. *J. Econom. Theory* **68** 279–302.
- [9] Dodis, Y., S. Halevi, T. Rabin. 2000. A cryptographic solution to a game theoretic problem. *Advances in cryptography—CRYPTO 2000 (Santa Barbara, CA)*, *Lecture Notes in Computer Science*, Vol. 1880. Springer-Verlag, Berlin, 112–130.
- [10] Forges, F. 1986. An approach to communication equilibria. *Econometrica* **54** 1375–1385.
- [11] Forges, F. 1986. Correlated equilibria in repeated games with lack of information on one side: A model with verifiable types. *Internat. J. Game Theory* **15** 65–82.
- [12] Forges, F. 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory Decision* **35** 277–310.
- [13] Forges, F. 2006. Correlated equilibrium in games with incomplete information revisited. *Theory and Decision* **61** 329–344.
- [14] Garey, M. R., D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, CA.
- [15] Gerardi, D. 2004. Unmediated communication in games with complete and incomplete information. *J. Econom. Theory* **114** 104–131.
- [16] Gibbons, R. 1992. *Game Theory for Applied Economists*. Princeton University Press, Princeton, NJ.
- [17] Gilboa, I., E. Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. *Games Econom. Behav.* **1** 80–93.
- [18] Goldberg, P. W., C. H. Papadimitriou. 2006. Reducibility among equilibrium problems. *Proc. 38th Annual ACM Sympos. Theory Comput. (STOC)* 61–70.
- [19] Grötschel, M., L. Lovász, A. Schrijver. 1993. *Geometric Algorithms and Combinatorial Optimization*, 2nd ed. Springer-Verlag, Berlin.
- [20] Hansen, K. A., P. B. Miltersen, T. B. Sørensen. 2007. Finding equilibria in games of no chance. COCOON 2007. G. Lin, ed. *Lecture Notes in Computer Science*, Vol. 4598. Springer-Verlag, Berlin, 274–284.
- [21] Hart, S., D. Schmeidler. 1989. Existence of correlated equilibria. *Math. Oper. Res.* **14** 18–25.
- [22] Kamien, M. I., Y. Tauman, S. Zamir. 1990. On the value of information in a strategic conflict. *Games Econom. Behav.* **2** 129–153.
- [23] Koller, D., N. Megiddo. 1992. The complexity of two-person zero-sum games in extensive form. *Games Econom. Behav.* **4** 528–552.
- [24] Koller, D., N. Megiddo, B. von Stengel. 1996. Efficient computation of equilibria for extensive two-person games. *Games Econom. Behav.* **14** 247–259.
- [25] Kuhn, H. W. 1953. Extensive games and the problem of information. *Contributions to the Theory of Games*, Vol. 2. *Annals of Mathematics Studies*, No. 28. Princeton University Press, Princeton, NJ, 193–216.
- [26] Lepinski, M., S. Micali, C. Peikert, A. Shelat. 2004. Completely fair SFE and coalition-safe cheap talk. *Proc. 23rd Annual ACM Sympos. Principles of Distributed Comput.* 1–10.
- [27] Moulin, H., J.-P. Vial. 1978. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *Internat. J. Game Theory* **7** 201–221.
- [28] Myerson, R. B. 1986. Multistage games with communication. *Econometrica* **54** 323–358.
- [29] Nau, R. F., K. F. McCardle. 1990. Coherent behavior in noncooperative games. *J. Econom. Theory* **50** 424–444.
- [30] Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley, Reading, MA.
- [31] Papadimitriou, C. H. 2005. Computing correlated equilibria in multi-player games. *Proc. 37th Annual ACM Sympos. Theory Comput. (STOC)* 49–56.
- [32] Papadimitriou, C. H., T. Roughgarden. 2005. Computing equilibria in multi-player games. *Proc. ACM-SIAM Sympos. Discrete Algorithms (SODA)* 82–91.
- [33] Papadimitriou, C. H., T. Roughgarden. 2008. Computing equilibria in multi-player games. *J. ACM* **55**, Article 14.
- [34] Romanovskii, I. V. 1962. Reduction of a game with complete memory to a matrix game. *Dokl. Akad. Nauk SSSR* **144** 62–64, English translation: *Soviet Mathematics* **3** 678–681.
- [35] Samuelson, L., J. Zhang. 1989. Correlated equilibria and mediated equilibria in games with incomplete information. Papers 8-89-5, Department of Economics, Pennsylvania State University, University Park, PA.
- [36] Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *Internat. J. Game Theory* **4** 25–55.
- [37] Solan, E. 2001. Characterization of correlated equilibria in stochastic games. *Internat. J. Game Theory* **30** 259–277.
- [38] Spence, M. 1973. Job market signaling. *Quart. J. Econom.* **87** 355–374.
- [39] Urbano, A., J. E. Vila. 2002. Computational complexity and communication: Coordination in two-player games. *Econometrica* **70** 1893–1927.
- [40] von Stengel, B. 1996. Efficient computation of behavior strategies. *Games Econom. Behav.* **14** 220–246.
- [41] von Stengel, B. 2001. Computational complexity of correlated equilibria for extensive games. Research Report LSE-CDAM-2001-03, London School of Economics, London.
- [42] von Stengel, B., A. van den Elzen, D. Talman. 2002. Computing normal form perfect equilibria for extensive two-person games. *Econometrica* **70** 693–715.
- [43] Young, H. P. 2004. *Strategic Learning and Its Limits*. Oxford University Press, Oxford, UK.
- [44] Zamir, S., M. I. Kamien, Y. Tauman. 1990. Information transmission. T. Ichiishi, A. Neyman, Y. Tauman, eds. *Game Theory and Applications*. Academic Press, San Diego, CA, 273–281.