

Finding Equilibria in Games of No Chance

Kristoffer Arnsfelt Hansen, Peter Bro Miltersen, and Troels Bjerre Sørensen

Department of Computer Science, University of Aarhus, Denmark
{arnsfelt,bromille,trold}@daimi.au.dk

Abstract. We consider finding maximin strategies and equilibria of explicitly given extensive form games with imperfect information but with *no moves of chance*. We show:

1. A maximin *pure* strategy for a two-player extensive form game with perfect recall and no moves of chance can be found in time linear in the size of the game tree. In contrast, it is known that this problem is **NP**-hard for games *with* chance moves.
2. All *pure* Nash equilibrium outcomes of a two-player *general-sum* extensive form game with perfect recall and no moves of chance can be enumerated in time linear in the size of the game tree. In contrast, finding a *behavior* Nash equilibrium is known to be **PPAD**-hard.
3. Finding an optimal *behavior* strategy for a one-player game of no chance in extensive form *without* perfect recall is **NP**-hard. In contrast, finding an optimal *pure* strategy in such a game can be done in linear time.
4. Determining whether an equilibrium in behavior strategies exists in a two-player zero-sum extensive form game of no chance without perfect recall is **NP**-hard.

1 Introduction

In a seminal paper, Koller and Megiddo [3] considered the complexity of finding maximin strategies in *two-player zero-sum imperfect-information extensive form games*. An extensive form game is an explicitly given game tree with *information sets* modelling hidden information (for details, see [3] or any text book on game theory). A main result of Koller and Megiddo was the existence of a polynomial time algorithm for finding an equilibrium in behavior strategies (or equivalently, a pair of maximin behavior strategies) of such a game when the game has *perfect recall*. Informally speaking, a game has perfect recall when a player never forgets what he once knew (for a formal definition, see below). In contrast, for the case of imperfect recall, the problem of finding a maximin strategy was shown to be **NP**-hard.

Pure equilibria (i.e., equilibria avoiding the use of randomization) play an important role in game theory and it is of special interest to know if a game possesses such an equilibrium. For the case of a zero-sum games, one may determine if a game has a pure equilibrium by computing a maximin pure strategy for each of the two players and checking that these strategies are best responses to one another. Unfortunately, Blair *et al* [1] established that the problems of finding a maximin pure strategy of a two-player extensive form game or determining whether a pure equilibrium exists are both **NP**-hard, even for the case of zero-sum games of perfect recall. Their proof is an elegant reduction from the EXACT PARTITION (or BINPACKING) problem and relies heavily on the fact that the extensive form game is allowed to contain *chance* nodes, i.e., random events not controlled by either of the two players.

Extensive form games *without* chance nodes is a very natural special case to consider (natural non-trivial examples include such popular parlor games as variants of *Spoof*). In this paper we consider the equilibrium computation problems considered by Koller and Megiddo and by Blair *et al* for this special case. Our main results are the following:

1. First, we show that *a maximin pure strategy for a two-player extensive form game of no chance with imperfect information but perfect recall can be found in time linear in the size of the game tree*. As stated above, Blair *et al* show that *with* chance moves, the problem is **NP**-hard. Apart from the obvious practical interest, the example is also interesting in light of the recent work of von Stengel and Forges [6]. They introduced the notion of *extensive form correlated equilibria* (EFCEs) of two-player extensive form games. They showed that finding such equilibria in games *without* chance moves can be done in polynomial time while finding them in games *with* chance moves is **NP**-hard. They remark that EFCE seems to be the first example of a game-theoretic solution concept where the introduction of chance moves marks the transition from polynomial-time solvability to **NP**-hardness. Our result combined with the result of Blair *et al* provides a second and much more elementary such example.
2. Second, we extend the above result from maximin pure strategies to pure Nash equilibria. We show that *all pure Nash equilibrium outcomes of a two-player general-sum extensive form game of no chance with imperfect information but perfect recall can be enumerated in time linear in the size of the game tree*. Here, an outcome is a leaf of the tree defining the extensive form. Also, given one such pure Nash equilibrium outcome, we can in linear time construct a pure equilibrium (in the form of a strategy profile) with that particular outcome. In contrast, the recent breakthrough result of Chen and Deng [2] shows that finding a behavior Nash equilibrium for a game of this kind is **PPAD**-hard¹.
3. The results of Blair *et al* and those of Koller and Megiddo give a setting where finding a pure equilibrium is **NP**-hard while finding an equilibrium in behavior strategies can be done in polynomial time. Considering games without perfect recall, we give an example of the opposite. We show that *determining whether a one-player game in extensive form with imperfect information, imperfect recall and no moves of chance has a behavior strategy that yields a given expected payoff is NP-hard*. In contrast, it is easy to see that finding an optimal *pure* strategy for such a game can be done in linear time. Our result strengthens a result of Koller and Megiddo [3, Proposition 2.5] who showed **NP**-hardness of finding a maximin behavior strategy in a *two-player* game with imperfect recall and no moves of chance. Koller and Megiddo [3, Example 2.12] also showed that a maximin behavior strategy in such a two-player game may require irrational behavior probabilities. We give a one-player example with the same property.
4. Finally, we show that determining whether a Nash equilibrium in behavior strategies exists in a two-player extensive form zero-sum game with no moves of chance but without perfect recall is **NP**-hard.

1.1 Outline of paper

In Section 2, we formally define the objects of interest and introduce the associated terminology (for a less concise introduction, see the paper by Koller and Megiddo, or any textbook on game theory). In Sections 3,4,5 and 6, we prove each of the four results mentioned above.

¹ Chen and Deng actually show hardness for the corresponding problem for *normal form* games, but a normal form game can be easily converted to a two-move extensive form game with no chance nodes.

2 Preliminaries

A two-player *extensive form game* is given by a finite rooted tree with pairs of payoffs (one payoff for each of the two players) at the leaves, and information sets partitioning nodes of the tree. For an example, see Figure 1.




Fig. 1. A two-player extensive form game

In a *zero-sum* game, the sum of each payoff pair is zero. A *general-sum* game is a game without this requirement. In this paper, we do not consider games with nodes of chance, so every node in the tree is owned by either Player 1 or to Player 2. All nodes in an information set belong to the same player. Intuitively, the nodes in an information set are indistinguishable for the player they belong to. In a one-player game, all nodes belong to Player 1.

Actions of a player are denoted by labels on edges of the tree. Given a node u and an action c that can be taken in u , we let $\text{apply}(u, c)$ be the unique successor node v of u with the edge (u, v) being labeled c . Each node in an information set has the same set of outgoing actions. The set of possible actions in information set h we denote C_h . The actions belong to the player owning the nodes of the information set. *Perfect recall* means that all nodes in an information set belonging to a player share the sequence of actions and information sets belonging to that player that are visited on the path from the root to each of the nodes.

A *pure strategy* for a player assigns to each information set belonging to that player a chosen action. A *behavior strategy* assigns to each action at each information set belonging to that player a probability. A pure strategy can also be seen as a behavior strategy that only uses the probabilities 0 and 1. Thus, concepts defined below for behavior strategies also apply to pure strategies. A (pure or behavior) *strategy profile* is a pair of (pure or behavior) strategies, one for each player. Given a

pure strategy profile for a game without chance nodes, there is a unique path in the tree from the root to a leaf formed by the chosen actions of the two players. The leaf is called the *outcome* of the profile. A behavior strategy profile defines in the natural way a probability distribution on the leaves of the tree and hence a probability distribution on payoffs for each of the two players. So given a behavior strategy profile we can talk about the expected payoff for each of the two players.

A *maximin pure strategy* for a player is a pure strategy that yields the maximum possible payoff for that player assuming a worst case opponent, i.e., the maximum possible guaranteed payoff. A *maximin behavior strategy* for a player is a behavior strategy that yields the maximum possible expected payoff for that player assuming a worst case opponent, i.e., the maximum possible guaranteed expected payoff. A *Nash equilibrium* is a strategy profile (s_1, s_2) so that no strategy s'_1 yields strictly better payoff for Player 1 than s_1 when Player 2 plays s_2 and no strategy s'_2 yields strictly better payoff for Player 2 than s_2 when Player 1 plays s_1 .

Kuhn [5] showed that for an extensive form two-player zero-sum game with perfect recall, a pair of maximin behavior strategies is a Nash equilibrium. The expected payoff for Player 1 is the same in any such equilibrium and is called the *value* of the game. Any extensive form general-sum game with perfect recall in fact possesses a Nash equilibrium in behavior strategies. Without the perfect recall assumption, not only is a pair of maximin behavior strategies of a zero-sum game not necessarily a Nash equilibrium, but Nash equilibria in behavior strategies are not even guaranteed to exist.

3 Maximin pure strategies in games with perfect recall

Consider a two-player extensive form game G with perfect recall and without chance nodes. We shall consider computing a maximin pure strategy for one of the players, say, Player 1. For the purpose of computing such a strategy, we can consider G to be a zero-sum game where Player 1 (henceforth the max-player) attempts to maximize his payoff and Player 2 (henceforth the min-player) attempts to minimize the payoff of Player 1. Let G' be the zero-sum game obtained from G by dissolving all information sets of the min-player into singletons.

Note that the set of strategies for the max-player is the same in G and G' . For the min-player, however, the set of strategies is larger in G' thereby making G' a better game than G for the min-player, so its *value* as a zero-sum game is at most the value of G . However we have the following key lemma. Note that the lemma fails badly for games containing chance nodes.

Lemma 1. *A pure strategy π for the max-player has the same payoff against an optimal counter strategy in G as it has against an optimal counter strategy in G' (note that the statement makes sense as the max-player has the same set of strategies in the two games).*

Proof. Let σ be a pure best counter strategy against π in G' . As there are no chance nodes, σ and π defines a single path in the tree of G' from the root to a leaf. Due to perfect recall, none of the choices made by the min-player along the path are choices of the same information set. Thus, the same sequence of choices can also be made by a strategy in G . Thus, there is a counter strategy in G that achieves the same payoff against π as σ does in G' , and since the set of possible counter strategies is bigger in G' , the best in each game each achieves exactly the same payoff.

To compute the best payoff that can be obtained by a pure strategy in G' , we define for information set h of G' a value $\text{pval}(h)$ (“pure value”) inductively² in the game-tree as follows.

- If h belongs to the min-player, and therefore consists of a single node u , define

$$\text{pval}(h) = \min_{c \in C_h} \text{pval}(\text{apply}(u, c))$$

- If h belongs to max-player, define

$$\text{pval}(h) = \max_{c \in C_h} \min_{u \in h} \text{pval}(\text{apply}(u, c))$$

Lemma 2. *For every pure strategy π for the max-player, there exists a pure strategy σ for the min-player with the following property. For every information set h of the max-player there is some node $u \in h$ such that that play from u using the pair of strategies (π, σ) yields payoff at most $\text{pval}(h)$. Similarly, for every information set h of the min-player, play from the single node u of h using the pair of strategies (π, σ) yields payoff at most $\text{pval}(h)$*

Proof. Given a pure strategy π for the max-player, we construct the strategy σ inductively in the game tree. Let h be a given information set of the max-player. Then, by definition of $\text{pval}(h)$ there must be a path from some node $u \in h$ using the action chosen by π out of u (say, L), then going through min-nodes to an information set g of the max-player with $\text{pval}(g) \leq \text{pval}(h)$, or to a leaf l with payoff less than or equal to $\text{pval}(h)$.

In the latter case we simply let σ take the choices defining the path to the leaf l . In the former case, by induction, we know we have constructed a pure strategy σ for min from g onwards so that for some node $v \in g$, play from g using π and σ leads to payoff at most $\text{pval}(g)$. Note that we have a path from u to some (possibly) other node $v' \in g$ using min-nodes. We claim that there is a path from some node $\bar{u} \in h$ to v using min-nodes and also choosing the action L in \bar{u} (see Figure 2).




Fig. 2. Finding \bar{u}

Indeed, assume that this is not the case. Then the sequence of information sets and own actions encountered by max on the way to v differs from the corresponding sequence in some of other node

² The induction is well-founded due to perfect recall and the fact that there are no chance nodes, see [6, Lemma 3.2].

(namely v') in the information set of v , contradicting perfect recall. But then, the node \bar{u} establishes the induction claim, with the desired strategy σ taking the choices defining the path from \bar{u} to v .

It remains to provide the first actions for the min-player in case the root node belongs to the min-player. In this case there is a path from the root r , going through min-nodes to an information set h of max with $\text{pval}(h) \leq \text{pval}(r)$, or to a leaf l with payoff equal to $\text{pval}(r)$. As before we let σ take the choices defining this path.

With this we can now obtain the following result.

Theorem 1. *Given a two-player extensive form game with perfect recall G without chance nodes, we can compute a maximin pure strategy for a player in linear time in the size of the game tree.*

Proof. We describe how to compute a maximin strategy for one of the players, say Player 1. By Lemma 1 we can compute this by computing a pure maximin strategy in the game G' . We compute the pval function of the information sets in G' and let the strategy of the max-player be the choices that obtains the maximum in the definition of pval for every information set, i.e., the choice in information set h is $\arg\max_{c \in C_h} \min_{u \in h} \text{pval}(\text{apply}(u, c))$. We claim that the value $\text{pval}(r)$ assigned to the root is the best guaranteed payoff the max-player can get in G' using some pure strategy. Indeed the max-player is guaranteed payoff $\text{pval}(r)$, where r is the root of G' , playing this strategy, and Lemma 2 establishes this is the best he can be guaranteed.

Note also that having computed the maximin pure strategy, we can determine whether it is also maximin as a behavior strategy by computing the value v of the game in polynomial time using, e.g., the algorithm of Koller and Megiddo [3] or the more practical one by Koller, Megiddo and von Stengel [4] and checking if the computed pure value $\text{pval}(r)$ of the root equals v .

4 Enumerating all pure equilibria of games with perfect recall

Let G be a 2-player *general sum* extensive form game with perfect recall and without chance nodes. Let (π, σ) be a pair of pure strategies. For (π, σ) to be a pure equilibrium we must have that π is a best response to σ and vice versa. Play using the pair (π, σ) will lead to a unique leaf of G , since there are no chance nodes. Consider now a leaf l of G , as a potential outcome of a pure equilibrium. Clearly the actions along the path from the root r of G to the leaf must be such that they follow the path. Hence what remains are to find the actions of the remaining information sets. Player 1 must find pure actions in his remaining information sets such that Player 2 can not obtain greater payoff than she receives at l . Similarly Player 2 must find pure actions in her information sets such that Player 1 can not obtain greater payoff than he receives at l . Given l , we can define zero-sum games G_1 and G_2 by modifying G such that such actions, if they exist, can be found in linear time using Theorem 1.

We can simply construct G_1 from G as follows (the construction of G_2 being the same with Player 1 and Player 2 exchanged). Player 1 will be the max-player of G_1 and Player 2 will be the min-player. For every information set of Player 1 along the path from the root to l we remove all choices (and the subgames below) except the ones agreeing with the path. The payoff at a leaf in G_1 is the negative of the payoff that Player 2 receives in the corresponding leaf in G . The following lemma is immediate.

Lemma 3. *There is a pure strategy for Player 1 in G leading towards l ensuring that Player 2 can obtain at most payoff p if and only if there is a pure strategy for the max-player of G_1 ensuring payoff at least $-p$.*

Using this lemma, it is easy to check in linear time if a given leaf l with payoffs (p_1, p_2) is a pure equilibrium outcome: We check that the maximin pure strategy for Player 1 in G_1 ensures payoff at least $-p_2$ and we check that the maximin pure strategy for Player 2 in G_2 ensures payoff at least $-p_1$. Also, given such an outcome, we can in linear time construct a pure strategy equilibrium with this outcome: The equilibrium is the profile consisting of transferring in the obvious way to G the maximin pure strategies for Player 1 in G_1 and for Player 2 in G_2 .

Since we can check in linear time if a given leaf is an outcome, we can enumerate the set of outcomes in quadratic time. To get a linear time algorithm, we will go one step further and work with a derived game that is independent of the leaf l .

Let G'_1 be the zero-sum game obtained from G by dissolving the information sets of Player 2 and letting payoff at a leaf in G'_1 be the negative of the payoff that Player 2 receives in the corresponding leaf in G . We define the pval function on G'_1 as in section 3.

Let T_1 be a tree on the information sets of Player 1 and the leaves together with a root, such that the parent of an information set or leaf is the first information set on the path to the root in G'_1 or the root itself.

Define a *point of deviation* with respect to a given leaf l , to be a node in T not on the path from the root to l , but having its parent on the path from the root to l . The following lemma is easy to establish.

Lemma 4. *There is a pure strategy for Player 1 in G leading towards l ensuring that Player 2 can obtain at most the payoff p if and only if for every point of deviation h with respect to l we have $\text{pval}(h) \geq -p$.*

Theorem 2. *Given a 2-player general-sum extensive form game with perfect recall G without chance nodes, we can in linear time in the size of the game tree enumerate the set of leaves that are outcomes of pure equilibria.*

Proof. Using Lemma 4, we compute the leaves l such that Player 1 has a pure strategy leading towards l ensuring that Player 2 can obtain at most the payoff received at l and conversely Player 2 has a pure strategy leading towards l ensuring that Player 1 can obtain at most the payoff received at l . These sets can be computed separately; we describe how to compute the former.

We construct the game G'_1 and compute the pval function on G'_1 in linear time. In linear time we then construct the tree T_1 and record the computed pval values in the nodes. Finally we traverse the tree T . During this traversal we maintain the minimum pval value that is on any sibling to the nodes on the path to the root, corresponding to the points of deviations relevant for the leaves in the subtree of the current node. Once we visit a leaf we can then directly decide the criteria of Lemma 4 by comparing with the payoff of the leaf.

5 Optimal behavior strategies in one-player games without perfect recall

In this section we consider *one-player* games without perfect recall and no moves of chance and show **NP**-hardness of the problem of determining whether a behavior strategy yielding an expected payoff of at least a given rational number exists. In contrast, it is straightforward to see that the corresponding problem for pure strategies is in **P**: For each leaf of the game, one checks if this leaf can be reached by a sequence of actions so that the same action is taken in all nodes in a given information set. This results strengthens the result of Koller and Megiddo [3, Proposition

2.6] who showed **NP**-hardness of the problem of determining whether some behavior strategy in a *two-player* game without perfect recall guarantees a certain expected payoff (against any strategy of the opponent). Also, our reduction is heavily based on their reduction but uses imperfect recall to eliminate one of the players. Before giving the proof, we give a simple example showing that an optimal strategy may require irrational behavior probabilities (therefore, strictly speaking, “finding” an optimal strategy is not a well-defined computational problem which leads to considering the stated decision problem instead). A corresponding two-player example was given by Koller and Megiddo [3, Example 2.12]. Our one-player game of Figure 3 is in fact somewhat simpler than their example. All nodes in the game are included in the same information set. The player can choose




Fig. 3. A one-player game where the rational behavior is irrational

either L or R. Thus, a behavior strategy is given by a single probability p_L with $p_R = 1 - p_L$. By construction, the expected payoff is $-2p_L^3 - (1 - p_L)^3$. This is maximized for $p_L = \sqrt{2} - 1$.

Theorem 3. *The following problem is **NP**-hard: Given a rational number v and a one-player extensive form game without chance nodes and a rational number v , does some behavior strategy ensure expected payoff at least v ?*

Proof. The proof is by reduction from 3SAT. Given a 3-CNF formula F with m clauses we construct a game G as follows.

Assume without loss of generality that m is a power of 2, $m = 2^k$. First G will consist of a complete binary tree of depth $2k$, whose nodes are contained in a single information set. If on the path from the root to a node, the same choice is made in step $2(i - 1) + 1$ and $2i$ for some $i \in \{1, \dots, k\}$, the game is terminated and the player receives payoff 0. Otherwise, we will associate a clause to the node in the following way: For $i = 1, \dots, k$ we interpret the choices made at step $2(i - 1) + 1$ and $2i$ as defining a binary choice. With the choices (left, right) we associate the bit 0, and with choices (right, left) we associate the bit 1. Having defined in this way k bits, we may associate a uniquely determined clause with the node.

From this node we let the player, for each of the three variables in the clause, select a truth value. If one of these choices satisfies the clause, the player receives payoff 1, and 0 otherwise. We place the nodes corresponding to the same variable in a single information set. In particular, the player does not know the clause.

The proof is now concluded by the following claim: The player can obtain expected payoff $\frac{1}{m}$ if and only if F is satisfiable.

Assume first that F is satisfiable. The player will make the first $2k$ choices by choosing left with probability $\frac{1}{2}$. The rest of the choices are made according to a satisfying assignment to F . With probability $(\frac{1}{2})^k = \frac{1}{m}$, the player gets to a node corresponding to a clause, and will obtain payoff 1. The expected payoff is therefore $\frac{1}{m}$.

Assume on the other hand that the player can obtain expected payoff $\frac{1}{m}$. Suppose that the player chooses left with probability p in the first $2k$ choices. The probability that the player reaches a node associated with a given clause is $(p(1-p))^k \leq \frac{1}{m^2}$, independently of the given node. Since the player can in fact obtain expected payoff $\frac{1}{m}$, we have that at every node associated with a clause the player must obtain payoff 1, and thus his strategy gives a satisfying assignment to F .

6 Determining whether a two-player game without perfect recall has an equilibrium

Our final hardness result again uses a reduction very similar to Koller and Megiddo [3, Proposition 2.6]. In this case, we use the imperfect recall to force Player 1 to use an almost pure strategy.

Theorem 4. *The following problem is NP-hard: Given a two-player zero-sum extensive form game without chance nodes, does the game possess a Nash equilibrium in behavior strategies?*

Proof. The proof is by reduction from 3SAT. Given a 3-CNF formula F with m clauses we construct a zero-sum two-player game G as follows.

Player 1 (the max-player) starts the game by making two actions, each time choosing a clause of F . We put all corresponding $m+1$ nodes (the root plus m nodes in the next layer) of the game in one information set. If he fails to choose the same clause twice, he receives a payoff of $-m^3$ and the game stops. Otherwise, Player 2 (the min-player) then selects a truth value for each of the three variables in the clause. We place all nodes of Player 2 corresponding to the same variable in a single information set. If one of the choices of Player 2 satisfies the clause, Player 1 receives payoff 0. If none of them do, Player 1 receives payoff 1.

The proof is now concluded by the following claim: G has an equilibrium in behavior strategies if and only if F is satisfiable.

Assume first that F is satisfiable. G then has the following equilibrium (which happens to be pure): Player 2 plays according to a satisfying assignment while Player 1 uses an arbitrary pure strategy. The payoff is 0 for both players and no player can modify their behavior to improve this so we have an equilibrium.

Next assume that G has an equilibrium. We shall argue that F has a satisfying assignment. We first observe that Player 1 in equilibrium must have expected payoff at least 0. If not, he could switch to an arbitrary pure strategy and would be guaranteed a payoff of at least 0. Now look at the two actions (i.e., clauses) that Player 1 is most likely to choose. Let clause i be the most likely and let clause j be the second-most likely. If Player 1 chooses i and then j he gets a payoff of $-m^3$. His maximum possible payoff is 1 and his expected payoff is at least 0. Hence, we must have that $-m^3 p_i p_j + 1 \geq 0$. Since $p_i \geq 1/m$, we have that $p_j \leq 1/m^2$. Since clause j was the second most likely choice, we in fact have that $p_i \geq 1 - (m-1)(1/m^2) > 1 - 1/m$. Thus, there is one clause that Player 1 plays with probability above $1 - 1/m$. Player 2 could then guarantee an expected payoff of less than $1/m$ for Player 1 by playing any assignment satisfying this clause. Since we are actually

playing an equilibrium, this would not decrease the payoff of Player 1 so Player 1 currently has an expected payoff less than $1/m$. Now look at the assignment defined by the most likely choices of Player 2 (i.e, the choices he makes with probability at least $\frac{1}{2}$, breaking ties in an arbitrary way). We claim that this assignment satisfies F . Suppose not. Then there is some clause not satisfied by F . If Player 1 changes his current strategy to the pure strategy choosing this clause, he obtains an expected payoff of at least $(1/2)^3 \geq 1/m$ (supposing, wlog, that $m \geq 8$). This contradicts the equilibrium property and we conclude that the assignment in fact does satisfy F .

7 Acknowledgments

We would like to thank Daniel Andersson, Lance Fortnow, and Bernhard von Stengel for helpful comments and discussions.

References

1. Jean R. S. Blair, David Mutchler, and Michael van Lent. Perfect recall and pruning in games with imperfect information. *Computational Intelligence*, 12:131–154, 1996.
2. Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*, pages 261–272, 2006.
3. Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992.
4. Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing*, pages 750–759, 1994.
5. H.W. Kuhn. Extensive games and the problem of information. *Annals of Mathematical Studies*, 28:193–216, 1953.
6. Bernhard von Stengel and Francoise Forges. Extensive form correlated equilibrium: Definition and computational complexity. Technical Report LSE-CDAM-2006-04, London School of Economics, Centre for Discrete and Applicable Mathematics, 2006.