Symbolic Computation of Sequential Equilibria

Moritz Graf
University of Freiburg
Freiburg, Germany
grafm@cs.uni-freiburg.de

Thorsten Engesser
IRIT
Toulouse, France
thorsten.engesser@irit.fr

Bernhard Nebel
University of Freiburg
Freiburg, Germany
nebel@uni-freiburg.de

ABSTRACT

The sequential equilibrium is a standard solution concept for extensive-form games with imperfect information that includes an explicit representation of the players' beliefs. An assessment consisting of a strategy and a belief is a sequential equilibrium if it satisfies the properties of sequential rationality and consistency.

Our main result is that both properties together can be written as a single finite system of polynomial equations and inequalities. The solutions to this system are exactly the sequential equilibria of the game. We construct this system explicitly and describe an implementation that solves it using cylindrical algebraic decomposition. To write consistency as a finite system of equations, we need to compute the extreme directions of a set of polyhedral cones. We propose a modified version of the double description method, optimized for this specific purpose. To the best of our knowledge, our implementation is the first to symbolically solve general finite imperfect information games for sequential equilibria.[1]

KEYWORDS

game theory; extensive-form games; sequential equilibrium

ACM Reference Format:

Moritz Graf, Thorsten Engesser, and Bernhard Nebel. 2024. Symbolic Computation of Sequential Equilibria. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

The sequential equilibrium was proposed by Kreps and Wilson [9] as a solution concept for extensive-form games with imperfect information. In extensive-form games, players have multiple decision points in sequence, at which they must decide how to act. A strategy for a player specifies what action the player takes at each of their decision points. When an extensive-form game has imperfect information, it means that the players do not always know the exact state of the game. The reason is that there may be actions of other players (or random events) that the players do not observe.

The sequential equilibrium is a generalization of subgame perfect equilibrium, which is the standard solution concept for extensive-form games with perfect information. In addition to a strategy, a sequential equilibrium specifies a set of beliefs for each player, assigning probabilities to states that the player cannot distinguish between. Intuitively, two properties must be satisfied: strategies should be rational given the players' beliefs (*sequential rationality*), and beliefs should be reasonable given the players' strategies (*consistency*). Having beliefs as an explicit part of the game's equilibria can provide additional insight. The strategies specify what the players are doing, and the beliefs provide an explanation why.

While finding all Nash equilibria or all subgame perfect equilibria of a game has been implemented in tools like *Gambit* [12] or *Game Theory Explorer* [18], there are no implemented solvers that symbolically compute all sequential equilibria of a finite game. Azhar et al. [2] have outlined an algorithm, enumerating so-called "consistent bases" of a game (of which there might be exponentially many) and characterizing the sequential equilibria for each basis by a system of polynomial equations and inequalities.

We take a similar approach, characterizing all equilibria of the game by a single such system. To this end, we combine existing results from the literature [7, 8]. Most importantly, our paper details all the steps necessary to generate and solve the system.

Polynomial systems of equations can in general have an infinite number of solutions, and indeed games with imperfect information often have an infinite number of sequential equilibria. Therefore, solving such a system does not require simply enumerating all the solutions, but rather finding a description of the solutions that allows any one of them to be easily extracted. For this purpose we will use the cylindrical algebraic decomposition algorithm provided by the computer algebra system *Mathematica*. We obtain a list of intervals, one for each of the variables, which are stratified in the sense that the boundaries of the intervals depend only on the variables before them. With this, any sequential equilibria of the game can be obtained by successively choosing a value for each variable. Using symbolic computation has the added advantage of allowing us to solve families of games where some outcomes or probabilities of random events are controlled by a set of parameters.

In Section 2 we recapitulate the basic definitions and solution concepts for extensive-form games. In Section 3 we study the properties of sequential equilibria and show how both sequential rationality and consistency can be written as a system of polynomial equations and inequalities. Section 4 shows the steps required to implement a sequential equilibrium solver, as well as some strategies for reducing computation time. We conclude in Section 5.[2]

2 THEORETICAL BACKGROUND

In Section 2.1, we introduce the notation used throughout this paper, mostly following Osborne and Rubinstein [14]. In Section 2.2 we

[1]Our implementation is based on the open source Game Theory Explorer [18] and can be downloaded at https://github.com/tengesser/GTE-sequential. See also our system demonstration paper [6].

[2]An extended version of this paper, with an appendix containing more detailed proofs and a full description of our modified version of the double description method, is available on arXiv [5].

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

then recapitulate the most important solution concepts, leading to the definition of sequential equilibria.

2.1 Extensive-Form Games

Game. A game consists of a set of *players* $N = \{1, \dots, n\}$, a set of *histories* H of which a subset $Z \subseteq H$ are *terminal* histories, a *player function* $N(h)$ that assigns an acting player to each non-terminal history, a function $A(h)$ specifying the set of *actions* available at each non-terminal history, and a *utility function* $u_i(h^*)$ that assigns to each player $i \in N$ a utility for each of the terminal histories $h^* \in Z$. Imperfect information is represented by a set of *information sets* $\mathcal{I}$ that partitions the set of non-terminal histories. We consider only games where the set of histories H is finite.

Actions and Histories. Histories can be thought of as nodes in a game tree. Each history $h = \langle a_1, \dots, a_k \rangle \in H$ encodes the sequence of actions leading to that node. We say that $h = \langle a_1, \dots, a_k \rangle$ is a *prefix* of $h' = \langle a'_1, \dots, a'_l \rangle$ if $k \le l$ and $a'_i = a_i$ for all $i \in \{1, \dots, k\}$ (that is, if $h = h'$ or if h is an ascendant of h' in the game tree). If $h \in H$, it must thus be the case that $h' \in H$ for all prefixes h' of h. The terminal histories $h^* \in Z$ are exactly those histories which are not prefix of some other $h' \neq h^*$ from H (the leaf nodes of the tree).

Information Sets. The information sets $I \in \mathcal{I}$ are used to represent imperfect information in the game. After some history $h \in I$ is reached, the acting player cannot distinguish whether h or any other history $h' \in I$ is the actual history. This implies that in each history of an information set both the acting player and the set of available actions must be identical. We thus usually write $N(I)$ and $A(I)$ instead of $N(h)$ and $A(h)$. We say that a game has *perfect recall* if for any two histories h, h' in the same information set I of player i, the sequence of information sets encountered and actions played by i are identical for both h and h'. This poses a restriction on game trees: players cannot 'forget' information about their position in the game tree or about the actions they played. Sequential equilibria are only defined for games with perfect recall.

Strategies and Beliefs. Solution concepts usually contain a strategy profile $\beta = (\beta_1, \beta_2, \dots, \beta_n)$, consisting of a behavioral strategy β_i for each player i. These β_i assign to each information set I of player i a probability distribution over all the actions $a \in A(I)$, such that player i plays action $a \in A(I)$ at I with probability $\beta_i(I)(a)$.

Similarly, a system of beliefs μ assigns to each information set I a probability distribution over all histories $h \in I$, such that player i believes to be in history $h \in I$ with probability $\mu_i(I)(h)$.

It is often not important which player a strategy or belief belongs to. Since each information set has a distinct acting player i, we will often write $\beta(I)(a)$ and $\mu(I)(h)$ without the explicit subscript i.

Based on the players strategies β, we can formulate the probabilities of reaching a specific history $h = \langle a_1, \dots, a_k \rangle$ as

$$P_\beta(h) = \prod_{i=1}^{k} \beta(I_i)(a_i),$$

where I_i is the information set containing $\langle a_1, \dots, a_{i-1} \rangle$. We can expand this to the probability of reaching an information set I:

$$P_\beta(I) = \sum_{h \in I} P_\beta(h)$$

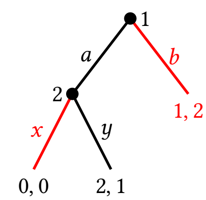

Figure 1: A non-credible threat in an extensive-form game. The actions that are played are marked in red. Each terminal node is annotated with the utilities of player 1 and 2.

We will also use the conditional probabilities of reaching history $h' = \langle a_1, \dots, a_k \rangle$ starting from a prefix $h = \langle a_1, \dots, a_l \rangle$, $l < k$:

$$P_\beta(h'|h) = \frac{P_\beta(h')}{P_\beta(h)} = \frac{\prod_{i=1}^{k} \beta(I_i)(a_i)}{\prod_{i=1}^{l} \beta(I_i)(a_i)} = \prod_{i=l+1}^{k} \beta(I_i)(a_i)$$

Note that $P_\beta(h|h') = 1$, since h' can only be reached after h is reached. If h is not a prefix of h', then $P_\beta(h'|h) = 0$.

Utilities. From the utilities $u_i(h^*)$ assigned to terminal histories, we define player i's expected utility given strategy profile β:

$$U_i^E(\beta) = \sum_{h^* \in Z} u_i(h^*) P_\beta(h^*)$$

We generalize this to player i's expected utility of the subgame starting at h:

$$U_i^E(\beta|h) = \sum_{h^* \in Z} u_i(h^*) P_\beta(h^*|h)$$

To obtain the utility that player i assigns to an information set I, we need the concept of believed utility:

$$U_i^B(\beta, \mu|I) = \sum_{h^* \in Z} u_i(h^*) \sum_{h \in I} \mu(I)(h) P_\beta(h^*|h)$$

Note that at most one term of the inner sum is nonzero, because for every terminal history h^* there is at most one $h \in I$ such that h is a prefix of h^* and therefore $P_\beta(h^*|h) \neq 0$.

2.2 Solution Concepts

The most fundamental solution concept in game theory is the Nash equilibrium. Intuitively, a strategy profile is a Nash equilibrium if no player can improve their utility by deviating from their strategy.

Definition 1 (Nash equilibrium). A strategy profile β is a Nash equilibrium if $U_i^E(\beta') \le U_i^E(\beta)$ for all players i and strategy profiles $\beta' = (\alpha_i, \beta_{-i})$ which only deviate from β in player i's strategy.

For sequential games, where players do not take their actions simultaneously, Nash equilibria can run into the problem of *non-credible threats*. An example is shown in Figure 1. The highlighted strategy profile, where player 1 plays b and player 2 plays x, is a Nash equilibrium. However, playing x is irrational for player 2: If history $\langle a \rangle$ was reached, player 2 should play y to get a higher payoff. This is not captured by Nash equilibria. Its definition considers only the expected utility of the whole game, which does not change when strategies in unreached parts of the game tree change.

This leads to the subgame perfect equilibrium, a refinement of the Nash equilibrium in which deviations from the strategy profile must also not increase the payoffs received by players in any subgame.

Definition 2 (Subgame Perfect Equilibrium). A strategy profile β is a subgame perfect equilibrium if $U_i^E(\beta'|h) \leq U_i^E(\beta|h)$ for all players i, all histories h where that player acts, and all strategy profiles $\beta' = (\alpha_i, \beta_{-i})$ that only deviate from β in player i's strategy.

The marked strategy profile β in Figure 1 is not subgame perfect, since $U_2^E(\beta'|\langle a\rangle) = 1 > 0 = U_2^E(\beta|\langle a\rangle)$, where β' is the strategy profile in which player 2 plays y at $\langle a\rangle$ instead of x.

In imperfect information games, when players reach some information set I, they must decide which action to play without knowing which $h \in I$ is the actual history. To preserve the concept of subgame perfection, players must form a belief $\mu(I)(h)$ about which history they are in. Pairs (β, μ) of a strategy profile β and a belief system μ are called *assessments*. Intuitively, strategies must be believed to be optimal from each information set, and beliefs must reflect the probabilities of histories being reached based on the agents' strategies. This leads us to the following definition.

Definition 3 (Sequential Equilibrium). Let β be a strategy profile and μ a belief system. The assessment (β, μ) is

(i) *sequentially rational* if for any information set I of player i and any strategy profile $\beta' = (\alpha_i, \beta_{-i})$ with a different strategy for player i, we have
$$U_i^B(\beta', \mu|I) \leq U_i^B(\beta, \mu|I), \tag{1}$$
(ii) *consistent* if there exists a series of assessments (β^n, μ^n) such that $\lim_{n\to\infty}(\beta^n, \mu^n) = (\beta, \mu)$, and for all $h \in H$ and $I \in \mathcal{I}$,
$$P_{\beta^n}(h) > 0 \tag{2}$$
$$\mu^n(I)(h) = \frac{P_{\beta^n}(h)}{P_{\beta^n}(I)}, \tag{3}$$
(iii) a *sequential equilibrium* if it is both sequentially rational and consistent.

3 EXPLORING SEQUENTIAL EQUILIBRIA

In the following section we will explore these two properties of sequential equilibria. Since our goal is to compute all sequential equilibria of a given game, we will attempt to find ways to transform both properties into a form that can be used computationally.

For sequential rationality, we will reduce the problem to a more local one, where we only need to consider alternative strategies β' that differ from β only at a single information set I, and we will end up with sequential rationality described by a system of equations and inequalities that are linear in the beliefs $\mu(I)(h)$ and polynomial in the action probabilities $\beta(I)(a)$. This largely follows the work of Hendon et al. [7] and requires that the beliefs are consistent.

We then follow the work of Kohlberg and Reny [8] to also express consistency as a finite set of polynomial equations. This requires us to compute the extreme directions of a set of polyhedral cones, a problem we will discuss further in Section 4.2.

3.1 Sequential Rationality

Sequential rationality is the natural extension of subgame perfection. At every information set, the acting player must believe that no deviating strategy can improve their utility. That is, the acting player cannot achieve a higher payoff by changing their own action probabilities in that information set or further down the game tree.

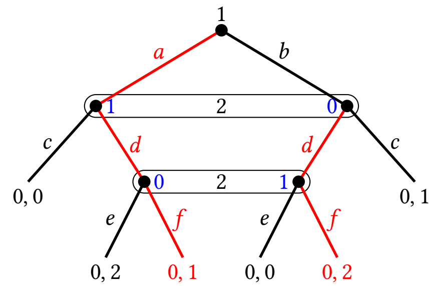

Figure 2: A game with an inconsistent assessment. Beliefs are depicted in blue, strategies in red. Although the assessment is *locally* sequentially rational, it is not sequentially rational.

We recapitulate the result of Hendon et al. [7], that if beliefs are consistent, we can reduce sequential rationality to a more local property. We also show how this property can be described by a set of polynomial equations and inequalities in β and μ. These results are similar to two concepts in the analysis of Nash equilibria and subgame perfect equilibria: the one-shot deviation principle, which states that we only need to consider local deviations to verify subgame perfection, and the fact that players may only assign positive probabilities to actions which are best responses to the opponents' strategy.

3.1.1 One-Shot Deviation Principle. Since sequential rationality is described as a property that holds for each information set, it is natural to define *sequential rationality at I* to mean that this property holds for a given information set I. In addition, we define *local sequential rationality at I* to mean that the property of sequential rationality at I holds for all strategies that differ from β only at I. These concepts are used in [7], but not defined by these names.

Definition 4 (Local Sequential Rationality). Let (β, μ) be an assessment and $I \in \mathcal{I}$ an information set of acting player i. The assessment (β, μ) is *locally sequentially rational at I* if $U_i^B(\beta', \mu|I) \leq U_i^B(\beta, \mu|I)$ for any strategy profile $\beta' = (\alpha_i, \beta_{-i})$ where agent i plays some strategy α_i, with $\alpha_i(I') = \beta_i(I')$ for all $I' \in \mathcal{I} \setminus \{I\}$. We say (β, μ) is *locally sequentially rational* if it is locally sequentially rational at every information set $I \in \mathcal{I}$.

We restate the one-shot deviation principle from Hendon et al. [7]:

Theorem 5 (One-Shot Deviation Principle, [7]). *Let (β, μ) be a locally sequentially rational assessment. If (β, μ) is consistent, then (β, μ) is sequentially rational and therefore a sequential equilibrium.*

Since sequential rationality implies local sequential rationality, we can follow from Theorem 5 that whenever an assessment is consistent, it is a sequential equilibrium if and only if it is also locally sequentially rational. Note that while our version of the theorem requires that the assessment be consistent, Hendon et al. [7] use a weaker form of consistency called *pre-consistency* that is sufficient for local sequential rationality to imply sequential rationality.

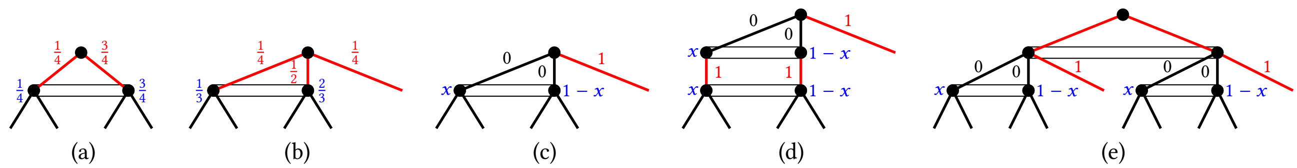

Figure 3: Five examples of game trees annotated with consistent assessments. Beliefs are marked in blue, strategies in red.

Perea [16] has further reduced the requirement to so-called *updating consistency*. Since consistency is a necessary condition for sequential equilibria, we will not use these weaker concepts.

Figure 2 shows why, without a consistency requirement, local sequential rationality is not a sufficient condition for sequential rationality. The problem is that even if an assessment is locally sequentially rational, a deviation in one information set might be believed to be profitable given the beliefs at an earlier information set. Consider the strategy where agent 2 plays e instead of f in the lower information set. While the believed utility in this information set decreases from 2 to 0, the believed utility at the upper information set increases from 1 to 2. Thus, while the depicted assessment is locally sequentially rational, it is not sequentially rational.

3.1.2 Best Responses. The following proposition provides a necessary and sufficient condition on local sequential rationality. If an action a is played with probability $\beta(I)(a) > 0$, then it must be a best response to the other players' actions. Our proposition generalizes Lemma 33.2 from the book by Osborne and Rubinstein.

Proposition 6. *An assessment (β, μ) is locally sequentially rational if and only if for all $I \in \mathcal{I}$ and $a \in A(I)$ the following holds:*

$$\text{if } \beta(I)(a) > 0, \text{ then } \sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle) = U_i^B(\beta, \mu | I) \quad (4)$$

$$\text{if } \beta(I)(a) = 0, \text{ then } \sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle) \leq U_i^B(\beta, \mu | I) \quad (5)$$

Proof sketch. We apply the proof idea for best responses in strategic games: If there is an action that violates (4) or (5), then we can construct a local deviation of β that results in a higher utility, violating local sequential rationality. If the assessment is not locally sequentially rational, then there must exist a local deviation with higher utility, which is only possible if (4) or (5) are violated. □

We can finally rewrite equations (4) and (5) from Proposition 6 as a system of polynomial equations and inequalities without case distinctions. For all $I \in \mathcal{I}$, $a \in A(I)$, and $i = N(I)$ we obtain

$$\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle) \right) - U_i^B(\beta, \mu | I) \quad \leq 0, \text{and}$$

$$\beta(I)(a) \cdot \left(\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle) \right) - U_i^B(\beta, \mu | I) \right) = 0.$$

An assessment is locally sequentially rational if and only if it satisfies this system of equations. By Theorem 5, the sequential equilibria of a game are exactly the consistent assessments that are locally sequentially rational. In the following, we will show how consistency can be similarly characterized as a system of polynomial equations, following the results of Kohlberg and Reny [8].

3.2 Consistency

While sequential rationality enforces that strategies are optimal given players' beliefs, consistency enforces that beliefs correctly reflect the conditional probabilities of each history being reached, given players' strategies. In particular, for information sets that are reached with probability $P_\beta(I) > 0$, we have

$$\mu(I)(h) = \frac{P_\beta(h)}{P_\beta(I)} = P_\beta(h|I).$$

Note that consistency depends only on the structure of the game tree. In particular, whether a given assessment is consistent does not depend on the utilities of the game, nor does it depend on the acting player at each information set. In the case of $P_\beta(I) = 0$, the restrictions imposed by consistency can become more complex.

Figure 3 shows the structures of five different game trees, together with possible consistent assessments. In the simplest cases, the beliefs correspond directly to the action probabilities (a), or to the conditional probabilities of the strategies leading to each history (b). Sometimes, if an information set is not reached, the assessment is consistent for any belief (c). However, such arbitrary beliefs may further constrain the beliefs at the next information set (d), and even at different parts of the game tree (e). In the last two examples, to satisfy consistency, the beliefs at both information sets must be identical, since they would have to be identical in any fully mixed assessment (i.e., with only positive action probabilities) that converges to (β, μ). For a more detailed discussion of how beliefs can be constrained by consistency, see the paper by Pimienta [17].

We will now follow the work of Kohlberg and Reny [8] to represent consistency by a finite set of polynomial equations. The reduction consists of several steps: An assessment is consistent if and only if a special system of linear equations $Ax = b$, where A depends on the structure of the game tree and b depends on the assessment, has a positive approximate solution. Such a solution exists if and only if a certain property holds for all vectors p with $pA = 0$. This property can be written as an equation in β and μ. Finally, the set of relevant vectors p can be reduced so that the system of equations becomes finite and polynomial without changing the solution set. To obtain this subset of relevant vectors, we have to compute the extreme directions of a set of polyhedral cones.

Our contribution is to provide more details on the approach by Kohlberg and Reny. We explicitly construct the linear system $Ax = b$ (Theorem 8) and formalize and prove all the necessary intermediate steps to derive the coefficients and exponents of the polynomial equations (Propositions 10-14). In Kohlberg and Reny's paper, the underlying ideas are stated informally and without proof.

3.2.1 Positive Approximate Solutions. The concept of positive approximate solutions is the first step in our process of representing consistency as a system of polynomial equations. These are solutions to systems of equations $Ax = b$, where the vector b may contain values such as ∞ or $-\infty$, or ill-defined expressions such as $\infty - \infty$ or $\frac{0}{0}$, using the conventions of the *extended real number line* $\overline{\mathbb{R}}$ (see, e.g., [1]). A positive approximate solution is then a series x^n such that each component of each vector x^n is positive and Ax^n converges to b for each component of b that is well-defined.

Definition 7 (Positive Approximate Solution of a Linear System). Let $Ax = b$ be a linear system where each $b_i \in \overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ or ill-defined. A positive approximate solution to this system is a series $(x^n)^{n \in \mathbb{N}}$ where $x^n > 0$ for all $n \in \mathbb{N}$, and $\lim_{n\to\infty}(Ax^n)_i = b_i$ for all i such that b_i is well-defined.

We will now construct a linear system from an assessment (β, μ) in such a way that it has a positive approximate solution if and only if the assessment is consistent.

Theorem 8 (Linear System for consistency). *Let $\{a_i \mid i \in \{1, \ldots, N_{actions}\}\}$ be the set of all actions in an extensive-form game (we assume without loss of generality that each action can be played in exactly one information set), and let $\{(h_i^1, h_i^2) \mid i \in \{1, \ldots, N_{pairs}\}\}$ be the set of all history pairs such that and both histories h_i^1 and h_i^2 are in the same information set. Furthermore, let $M = N_{actions} + N_{pairs}$. Then an assessment (β, μ) is consistent if and only if the linear system $Ax = b$ which is defined as follows has a positive approximate solution.*

$$A = \begin{bmatrix} \tilde{A} \\ I_{N_{actions}} \end{bmatrix}, \text{ where } \tilde{A} \in \mathbb{R}^{N_{pairs} \times N_{actions}} \text{ is defined as below:}$$

$$\tilde{A}_{i,j} = \begin{cases} 1 & \text{if } a_j \in h_i^1 \text{ and } a_j \notin h_i^2 \\ -1 & \text{if } a_j \notin h_i^1 \text{ and } a_j \in h_i^2 \\ 0 & \text{otherwise} \end{cases}$$

$$b_i = log(\alpha_i) - log(\gamma_i) \text{ for } i \in \{1, \ldots, M\}, \text{ where}$$

$$\alpha_i = \begin{cases} \mu(h_i^1) & \text{if } i \in \{1, \ldots, N_{pairs}\} \\ \beta(a_{i-N_{pairs}}) & \text{otherwise} \end{cases}$$

$$\gamma_i = \begin{cases} \mu(h_i^2) & \text{if } i \in \{1, \ldots, N_{pairs}\} \\ 1 & \text{otherwise} \end{cases}$$

Note that $A \in \mathbb{R}^{M \times N_{actions}}$, $b \in \overline{\mathbb{R}}^M$, and $x \in \mathbb{R}^{N_{actions}}$.

Proof sketch. Both positive approximate solutions and consistency depend on the existence of a convergent series. For consistency, we need a series of fully-mixed assessments (β^n, μ^n), where the μ^n are determined by the β^n via Bayes' rule. We define a multiplicative system for which the positive approximate solutions (which are defined similarly as for linear systems) are exactly these series. For $I \in \mathcal{I}$, each pair $h_1, h_2 \in I$ and action $a_i \in I$, we have:

$$\frac{\prod_{a_i \in h_1} x_i}{\prod_{a_i \in h_2} x_i} = \frac{\mu(I)(h_1)}{\mu(I)(h_2)} \tag{6}$$

$$x_i = \beta(I)(a_i) \tag{7}$$

The linear system is obtained by taking the logarithm of this system. It has a positive approximate solution if and only if the multiplicative system has one. The full proof is in the appendix [5]. □

Note that some of the b_i can be ill-defined. This is the case if $\alpha_i = \gamma_i = 0$ and thus $b_i = \log(0) - \log(0) = \infty - \infty$. Since the existence of positive approximate solutions only depends on the equations where b_i is well-defined, we can reduce the system such that the equations where b_i is ill-defined are omitted. From now on, we will assume that all b_i are well-defined.

3.2.2 *Existence of a Positive Approximate Solution.*

Kohlberg and Reny [8] give a result for the existence of a positive approximate solution to a linear system. We restate this here without proof:

Theorem 9 (Solution Existence for Linear Systems, [8]). *A linear system $Ax = b$ has a positive approximate solution if and only if the following property holds for all $p \in \mathbb{R}^M$ where $pA = 0$:*

$$\sum_{p_i \neq 0} p_i b_i = 0 \text{ or } \sum_{p_i \neq 0} p_i b_i \text{ is ill-defined} \tag{8}$$

The sum can be ill-defined if it contains the expressions $0 \cdot \infty$ or $\infty - \infty$. Since we only sum over $p_i \neq 0$, we only need to consider the second case. Importantly, we can write Property (8) for a given vector $p \in \mathbb{Z}^M$ as a polynomial equation.

Proposition 10. *Consider a linear system $Ax = b$ where $b_i = log(\alpha_i) - log(\gamma_i)$ for some $\alpha_i, \gamma_i \in \mathbb{R}$. Then Property (8) holds for some vector $p \in \mathbb{Z}^M$ if and only if the following equation is satisfied:*

$$\prod_{p_i>0} \alpha_i^{p_i} \prod_{p_i<0} \gamma_i^{-p_i} = \prod_{p_i>0} \gamma_i^{p_i} \prod_{p_i<0} \alpha_i^{-p_i} \tag{9}$$

Proof. Consider first the case where $\sum_{p_i \neq 0} p_i b_i$ is ill-defined. Here, we know that Property (8) always holds. Therefore, we only need to show that equation (9) is satisfied. For $\sum_{p_i \neq 0} p_i b_i$ to be ill-defined, there must exist indices i, j such that $p_i b_i = \infty$ and $p_j b_j = -\infty$. Here the sum $p_i b_i + p_j b_j = \infty - \infty$ is ill-defined. This happens if either $p_i > 0$ and $\gamma_i = 0$, or $p_i < 0$ and $\alpha_i = 0$. In any case, $\prod_{p_i>0} \gamma_i^{p_i} \prod_{p_i<0} \alpha_i^{-p_i} = 0$. Similarly, we know that either $p_j > 0$ and $\alpha_i = 0$, or $p_j < 0$ and $\gamma_j = 0$, which means that $\prod_{p_i>0} \alpha_i^{p_i} \prod_{p_i<0} \gamma_i^{-p_i} = 0$. Therefore, equation (9) is satisfied.

In the case where $\sum_{p_i \neq 0} p_i b_i$ is well-defined, Property (8) holds for p if and only if $\sum_{p_i \neq 0} p_i b_i = 0$. We obtain equation (9) by taking the exponential function and then multiplying by all the terms with a negative exponent.

$$\sum_{p_i \neq 0} p_i b_i = 0 \iff \sum_{p_i \neq 0} p_i(log(\alpha_i) - log(\gamma_i)) = 0$$
$$\iff \prod_{p_i \neq 0} \left(\frac{\alpha_i}{\gamma_i}\right)^{p_i} = 1$$
$$\iff \prod_{p_i>0} \alpha_i^{p_i} \prod_{p_i<0} \gamma_i^{-p_i} = \prod_{p_i>0} \gamma_i^{p_i} \prod_{p_i<0} \alpha_i^{-p_i}$$

Note that some of the terms we multiply by can be equal to zero. If this is the case, all of the terms with positive exponents are nonzero, since otherwise $\sum_{p_i \neq 0} p_i b_i$ would be ill-defined. Here, neither equation is satisfied and their equivalency still holds. □

3.2.3 *A Finite System of Equations.*

To write consistency as a finite system of polynomial equations, we have to solve two problems: In Theorem 9, we consider vectors which can have non-integer components. This means that we cannot use Proposition 10 to obtain an equivalent polynomial equation. Furthermore, there are infinitely many vectors p with $pA = 0$ (except in perfect information games where A is the identity matrix $I_{N_{actions}}$).

We now reduce the set of relevant vectors to a finite one.

Let $W_b^A \subseteq \{p \mid pA = 0\}$ be the set of all p such that $\sum_{p_i \neq 0} p_i b_i$ is well-defined. We then only need to check Property (8) for all $p \in W_b^A$, since we already know that it holds for all $p \notin W_b^A$.

Consider again that $\sum_{p_i \neq 0} p_i b_i$ is well-defined if there are no indices i and j such that $p_i b_i = \infty$ and $p_j b_j = -\infty$, For any $p \in W_b^A$, either all infinite terms of the sum must be positive, or all infinite terms must be negative. We can thus write $W_b^A = C_b^A \cup -C_b^A$ where

$$C_b^A = \{p \mid pA = 0 \wedge p_i \geq 0 \text{ if } b_i = \infty \wedge p_i \leq 0 \text{ if } b_i = -\infty\}, \text{ and}$$
$$-C_b^A = \{p \mid pA = 0 \wedge p_i \leq 0 \text{ if } b_i = \infty \wedge p_i \geq 0 \text{ if } b_i = -\infty\}.$$

We now show that it is sufficient to check Property (8) for all $p \in C_b^A$. As we will see, we do not need to consider $p \in -C_b^A$.

Proposition 11. *Let $Ax = b$ be the linear system from Theorem 8. Then the assessment (β, μ) is consistent if and only if Property (8) is satisfied for all $p \in C_b^A$.*

Proof. By Theorem 9, the assessment is consistent if and only if Property (8) is satisfied for all p where $pA = 0$. Since the property is satisfied if $\sum_{p_i \neq 0} p_i b_i$ is ill-defined, we do not need to consider vectors $p \notin W_b^A$. For $p \in W_b^A$, note that $\sum_{p_i \neq 0} p_i b_i = 0 \iff \sum_{p_i \neq 0} -p_i b_i = 0$. This means that Property (8) holds for p if and only if it holds for $-p$. The assessment is thus consistent if and only if Property (8) is satisfied for all $p \in C_b^A$. □

As we can see, C_b^A is an intersection of half spaces and thus a pointed polyhedral cone:

$$C_b^A = \{p \mid pA = 0\} \cap \bigcap_{b_i = \infty} \{p \mid p_i \geq 0\} \cap \bigcap_{b_i = -\infty} \{p \mid p_i \leq 0\}$$

We can alternatively represent C_b^A as the set of all conical combinations of finitely many vectors $\{e_1, \ldots, e_k\}$ such that

$$C_b^A = \{\lambda_1 e_1 + \ldots + \lambda_k e_k \mid \lambda_k \in \mathbb{R}^+\}.$$

These vectors are called extreme directions (or conical basis) of C_b^A. Transforming one representation into the other can be done with the double description method [22], which we will discuss in Section 4. Note that in the cases where all b_i are infinite, the extreme directions of C_b^A are unique modulo scaling. Otherwise, this is not necessarily the case. Furthermore, because the entries of A are always integers, each of the extreme directions can be scaled to have integer components. This is another result by Kohlberg and Reny [8]. We chose an arbitrary conical basis $ED(C_b^A)$ which has this property. This will allow us to use Proposition 10 to obtain a system of polynomial equations. We show that if Property (8) holds for two vectors of a cone, it also holds for arbitrary conical combinations. This allows us to reduce the system to a finite one.

Proposition 12. *If Property (8) holds for two vectors $x, y \in C_b^A$ then it must also hold for any conical combination $z = \alpha x + \beta y$, $\forall \alpha, \beta \in \mathbb{R}^+$.*

Proof. Since C_b^A is a cone, it follows that $z \in C_b^A$. Therefore $\sum_{z_i \neq 0} z_i b_i$ is well defined and Property (8) holds if $\sum_{z_i \neq 0} z_i b_i = 0$. We have $\sum_{x_i \neq 0} x_i b_i = \sum_{y_i \neq 0} y_i b_i = 0$. Let $z = \alpha x + \beta y$, $\alpha, \beta \in \mathbb{R}^+$. We split $z_i \neq 0$ into the cases $(x_i = 0, y_i \neq 0)$, $(x_i \neq 0, y_i = 0)$, and $(x_i \neq 0, y_i \neq 0)$.

$$\begin{aligned}
&\sum_{(\alpha x + \beta y)_i \neq 0} (\alpha x + \beta y)_i \cdot b_i \\
&= \sum_{\substack{x_i \neq 0 \\ y_i = 0}} (\alpha x + \beta y)_i \cdot b_i + \sum_{\substack{x_i = 0 \\ y_i \neq 0}} (\alpha x + \beta y)_i \cdot b_i + \sum_{\substack{x_i \neq 0 \\ y_i \neq 0}} (\alpha x + \beta y)_i \cdot b_i \\
&= \sum_{\substack{x_i \neq 0 \\ y_i = 0}} \alpha x_i b_i + \sum_{\substack{x_i = 0 \\ y_i \neq 0}} \beta y_i b_i + \sum_{\substack{x_i \neq 0 \\ y_i \neq 0}} \alpha x_i b_i + \beta y_i b_i \\
&= \sum_{\substack{x_i \neq 0 \\ y_i = 0}} \alpha x_i b_i + \sum_{\substack{x_i \neq 0 \\ y_i \neq 0}} \alpha x_i b_i + \sum_{\substack{x_i = 0 \\ y_i \neq 0}} \beta y_i b_i + \sum_{\substack{x_i \neq 0 \\ y_i \neq 0}} \beta y_i b_i \\
&= \sum_{x_i \neq 0} \alpha x_i b_i + \sum_{y_i \neq 0} \beta y_i b_i = 0
\end{aligned}$$

In the first step of our transformation, there might be some j where $x_j \neq 0, y_j \neq 0$ but $(\alpha x + \beta y)_j = z_j = 0$. The terms $z_j b_j$ would normally not be included in $\sum_{z_i \neq 0} z_i b_i$. In those cases, it follows that $\alpha x_j = -\beta y_j$ where $\alpha, \beta > 0$, therefore x_j and y_j have a different sign. Since $x, y \in C_b^A$, x_i and y_i must have the same sign whenever b_i is infinite. Thus b_j is finite and $z_j b_j = 0$. We can therefore add these terms to $\sum_{z_i \neq 0} z_i b_i$ while preserving equality.[3] □

We can now formalize a finite test for consistency.

Proposition 13 (Finite Consistency Test). *Let $Ax = b$ be the linear system from Theorem 8. Then the assessment (β, μ) is consistent if and only if Property (8) holds for all $p \in ED(C_b^A)$.*

Proof. If the assessment is consistent, then Property (8) must hold for all $p \in \{p \mid pA = 0\}$ due to Theorem (10), so it also holds for all $p \in ED(C_b^A) \subseteq C_b^A \subseteq \{p \mid pA = 0\}$. If Property (8) holds for all $p \in ED(C_b^A)$, then it holds for all $p \in C_b^A$ because of Proposition 12 and because each p can be written as conical combination of $ED(C_b^A)$. The assessment is then consistent due to Proposition 11.

Proposition 12 also implies that the choice of $ED(C_b^A)$ is irrelevant, since if Property (8) holds for one set of extreme directions, then it holds for the whole cone and thus for any other set of extreme directions. □

3.2.4 Finding all Consistent Assessments.

We now have a finite test for proving consistency of a given assessment (β, μ). However, we still cannot easily describe the set of all consistent assessments. This is because the test from Proposition 13 depends on the specific cone C_b^A, which depends on the right-hand side of the linear system $Ax = b$, which depends on the exact values of (β, μ). More precisely, it is the actions with $\beta(I)(a) = 0$ and the beliefs with $\mu(I)(h) = 0$ that determine which b_i are finite, ∞, or $-\infty$. Assuming that A (which only depends on the game tree) is fixed, only the positions of infinite values in b are relevant for C_b^A. Formally, if $b'_i = \infty \iff b_i = \infty$ and $b'_i = -\infty \iff b_i = -\infty$, then $C_{b'}^A = C_b^A$.

Since we want to characterize all sequential equilibria of a game, we need to find a criterion that works for arbitrary values of (β, μ).

[3] The same argument does not work for linear combinations. Assuming $x_j \neq 0$, $y_j \neq 0$, and $(\alpha x + \beta y)_j = z_j = 0$, it is possible that x_j and y_j have the same sign, since α and β can be negative. Then, b_j can be infinite, in which case $z_j b_j = 0 \cdot \infty$ is ill-defined.

As we will see, we can use the extreme directions of all cones $C^A = \{C_b^A \mid b_i \in \{-\infty, 0, \infty\}, \forall i\}$ relevant to A. The set of extreme directions of all cones relevant to A is defined as follows:

$$\mathcal{E}^A = \bigcup_{C \in C^A} ED(C) = \bigcup_{b_i \in \{-\infty,0,\infty\}, \forall i} ED(C_b^A).$$

We now show that we can use $\mathcal{E}^A$ to obtain a more general test.

Proposition 14 (General Consistency Test). *Let $Ax = b$ be the linear system from Theorem 8. Then the assessment (β, μ) is consistent if and only if Property (8) holds for all $p \in \mathcal{E}^A$.*

Proof. "$\Leftarrow$" If Property (8) holds for all $p \in \mathcal{E}^A$, then it holds specifically for all $p \in ED(C_b^A) \subseteq \mathcal{E}^A$. Thus, by Proposition 13, the assessment is consistent. "$\Rightarrow$" If the assessment is consistent, then by Theorem 9, Property (8) must hold for any p such that $pA = 0$. Since $\mathcal{E}^A \subseteq \{p \mid pA = 0\}$, Property (8) holds for all $p \in \mathcal{E}^A$. □

3.2.5 Polynomial Equations. We can now express consistency as a finite system of polynomial equations. This result follows directly from Propositions 10 and 14.

Theorem 15 (Consistency as Polynomial Equations, [8]). *Let A, α, and γ be defined as in Theorem 8, with α and γ containing the strategies $\beta(\cdot)$ and beliefs $\mu(\cdot)$ as variables. Then an arbitrary assessment (β, μ) is consistent if and only if for all $p \in \mathcal{E}^A$,*

$$\prod_{p_i>0} \alpha_i^{p_i} \prod_{p_i<0} \gamma_i^{-p_i} = \prod_{p_i>0} \gamma_i^{p_i} \prod_{p_i<0} \alpha_i^{-p_i}.$$

4 IMPLEMENTATION

In the previous section we have seen how sequential rationality can be expressed as a system of polynomial equations and inequalities if we assume consistency. We have also seen how consistency can be expressed as a system of polynomial equations. Together, these equations characterize the set of all sequential equilibria.

4.1 Equations

First, we recapitulate the entire system of equations and inequalities. The variables in our equations are the probabilities $\beta(I)(a)$ for each action a to be played at its information set I, and the beliefs $\mu(I)(h)$ that players assign to each history h at I. The equations are quantified over all $I \in \mathcal{I}$ (with $i = N(I)$), $a \in A(I)$, and $p \in \mathcal{E}^A$:

$$\beta(I)(a) \geq 0 \quad (10a) \qquad \mu(I)(h) \geq 0 \quad (11a)$$

$$\sum_{a \in A(I)} \beta(I)(a) = 1 \quad (10b) \qquad \sum_{h \in I} \mu(I)(h) = 1 \quad (11b)$$

$$\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle)\right) - U_i^B(\beta, \mu | I) \leq 0 \quad (12)$$

$$\beta(I)(a) \cdot \left(\left(\sum_{h \in I} \mu(I)(h) U_i^E(\beta | \langle h, a \rangle)\right) - U_i^B(\beta, \mu | I)\right) = 0 \quad (13)$$

$$\prod_{p_i>0} \alpha_i^{p_i} \prod_{p_i<0} \gamma_i^{-p_i} = \prod_{p_i>0} \gamma_i^{p_i} \prod_{p_i<0} \alpha_i^{-p_i} \quad (14)$$

Equations (10a-11b) ensure that strategies $\beta(I)$ and beliefs $\mu(I)$ are probability distributions. Equations (12) and (13) correspond to the sufficient and necessary conditions for local sequential rationality (Proposition 6). The equations of type (14) ensure consistency. That is, $\mathcal{E}^A$ is the set of extreme directions of all cones from

$$C^A = \{C_b^A \mid b_i \in \{-\infty, 0, \infty\}, i \in \{1, \ldots, M\}\}, \text{ where}$$

$$C_b^A = \{p | pA = 0\} \cap \bigcap_{i:\, b_i = \infty} \{p | p_i \geq 0\} \cap \bigcap_{i:\, b_i = -\infty} \{p | p_i \leq 0\},$$

and A, α, and γ are defined as in Theorem 8. Next, we will briefly detail how to compute these extreme directions.

4.2 Finding all Extreme Directions

A naive approach is to compute the extreme direction of each cone separately. This can be done with the so-called *double description method* [22]. This algorithm computes the extreme directions of a given cone by iteratively considering all constraints, calculating new extreme directions at each iteration based on the current constraint and the previously computed extreme directions.

For example, to determine the extreme directions of $C^A_{(\infty,\infty,\infty)}$, the algorithm computes the extreme directions of $C^A_{(0,0,0)}$, $C^A_{(\infty,0,0)}$, and $C^A_{(\infty,\infty,0)}$ as intermediate steps. As we can see, running the algorithm for each cone separately is inefficient because the extreme directions of some cones are computed exponentially often as intermediate steps. We can avoid this by computing the extreme directions of cones with fewer constraints first and memorizing the results for the computation of cones with more constraints. Consider the following collection of sets:

$$C_i = \{C_b^A \mid b_j = 0, \forall j \geq i\} \qquad \forall i \in \{1, \ldots, M+1\}$$

Our algorithm first computes the extreme directions of $C^A_{(0,\ldots,0)}$ (which is the only cone in C_1) and then iteratively computes the extreme directions for all the cones in the sets $C_2, \ldots, C_{M+1}$. Importantly, each cone in C_{i+1} corresponds to a cone in C_i with at most one constraint added ($p_i \leq 0$ or $p_i \geq 0$). The computation of new extreme directions for that cone thus corresponds to performing a single additional step of the double description method.

The way we iterate over the cones ensures that our algorithm only has to compute the extreme directions of each cone once. However, each cone may still be relevant to the set of extreme directions. In general, there are 3^M cones, where M is the number of actions plus the number of pairs of histories in the same information set. For larger games, the number becomes prohibitively large.

The number of cones can be reduced by identifying and removing actions that are not relevant to consistency. These are the actions such that for all pairs of histories in the same information set, the action is either on the path of both histories, or on neither.

We can further optimize our approach by pruning cones for which we can determine that no additional extreme directions will be introduced. The full algorithm can be found in the appendix [5].

4.3 Cylindrical Algebraic Decomposition

Cylindrical algebraic decomposition is an algorithm that partitions a subset of $\mathbb{R}^n$ specified by a set of polynomials into so-called *cells* (i.e., connected subsets of $\mathbb{R}^n$). For a complete description of the algorithm, see LaValle [11]. As implemented in *Mathematica* [20], it allows us to solve a system of polynomial equations and

$\beta(a)$	$\beta(c)$	$\beta(e)$	$\mu(\langle a\rangle)$	$\mu(\langle a,d\rangle)$	$U_2^E(\langle\rangle)$
$\in [0, \frac{2}{3})$	$= 0$	$= 0$	$= \beta(a)$	$= \beta(a)$	$= 2 - \beta(a)$
$= \frac{2}{3}$	$= 0$	$\in [0, 1]$	$= \frac{2}{3}$	$= \frac{2}{3}$	$= \frac{4}{3}$
$\in (\frac{2}{3}, 1]$	$= 0$	$= 1$	$= \beta(a)$	$= \beta(a)$	$= 2 \cdot \beta(a)$

Figure 4: All sequential equilibria of the game from Figure 2 as returned by the cylindrical algebraic decomposition. Note that, e.g., $\beta(b) = 1 - \beta(a)$, $\mu(\langle b\rangle) = 1 - \mu(\langle a\rangle)$ **and** $U_1^E(\langle\rangle) = 0$.

inequalities. Each cell successively assigns to each variable v_i an algebraic function specifying an interval $I_i(v_1, \ldots, v_{i-1})$ of possible values for that variable. To extract a specific solution, we can first select $v_1 \in I_1$, then $v_2 \in I_2(v_1)$, then $v_3 \in I_3(v_1, v_2)$, and so on.

Note that cylindrical algebraic decomposition is only defined for polynomials with coefficients in $\mathbb{Q}$. Strictly speaking, our approach is therefore restricted to games with rational payoffs.

Cylindrical algebraic decomposition has a computational complexity that is double exponential in the number of variables. This makes it infeasible for large games. Our system of polynomials contains one variable $\beta(I)(a)$ for each action and one variable $\mu(I)(h)$ for each history in each information set. We can reduce the number of variables by defining one belief and one action probability for each information set implicitly through the others.

Another optimization that has proven effective in some cases is to add additional equations describing the Nash equilibria of the game. Since all sequential equilibria are Nash equilibria, this does not change the solution space and we obtain the same solutions. The polynomial equations characterizing the Nash equilibria follow the usual idea of restricting the behavioral strategies to best responses.

An example output can be seen in Figure 4. The sequential equilibria are partitioned into three cells with infinitely many elements.

5 CONCLUSION

The main contribution of this paper is to show that the set of all sequential equilibria of an extensive-form game can be represented by a system of polynomial equations and inequalities. To this end, we combined theoretical results by Hendon et al. [7] and Kohlberg and Reny [8]. We furthermore described how to obtain and solve this system using symbolic computation.

For sequential rationality, we used a local form of sequential rationality that considers only the deviations at each information set. Local sequential rationality holds at I if and only if the acting player believes all played actions to be best responses to the other players' strategies. We have shown how to represent this property as a system of equations and inequalities that are linear in the beliefs $\mu(I)(h)$ and polynomial in the strategies $\beta(I)(a)$. Similar to the one-shot-deviation principle for subgame perfect equilibria, local sequential rationality is a sufficient (and necessary) condition for sequential rationality under the assumption that the assessment is already consistent, as shown by Hendon et al. [7].

For consistency, we briefly discussed the restrictions that consistency can pose on unreached information sets, which are not entirely obvious. Kohlberg and Reny [8] proposed a way to represent consistency as a system of polynomial equations. We have elaborated their approach in sufficient detail to be easily implemented. While this gives us a solution to represent consistency, it comes with a new problem: finding the extreme directions of a set of cones.

The double description method provides an established method to calculate the extreme directions of a single cone. However, it is inefficient for calculating the extreme directions of a set of cones. We proposed a modified version of the algorithm, that takes advantage of its iterative nature to avoid calculating the extreme directions of the same cone multiple times. Our method still considers each cone individually, and even though we found a way to prune some cones, the number of cones is still prohibitively large for larger games.

We use symbolic computation to solve our system of polynomial equations and inequalities. This gives us a compact and exact representation of all sequential equilibria of a game, even if there are connected components of infinitely many equilibria. However, both steps of this process, generating the system of equations and solving it, are infeasible for large games. In the first step, the number of cones is exponential in the size of the game tree, and in the second step, the cylindrical algebraic decomposition is double exponential in the number of variables in an assessment. While we have proposed some minor optimizations, pruning some of the cones and substituting some of the variables, these improvements are not sufficient to make the approach feasible for larger games.

Nevertheless, our implementation successfully handles small games such as Selten's Horse or small signaling games. It is already difficult to find all sequential equilibria for these games by hand. To the best of our knowledge, our implementation is the first tool that finds all sequential equilibria in finite imperfect information games, representing them symbolically. Besides for analyzing small games, we see a practical use for teaching the concept of sequential equilibria in game theory courses.

There are approaches from related work for computing sequential equilibria for restricted subsets of games. Miltersen and Sørensen [13] use minimax strategies to compute sequential equilibria for two-player games. Gilpin and Sandholm [4] compute sequential equilibria in a class of games where chance nodes are the only source of uncertainty. Turocy [21] provides a numerical algorithm to compute sequential equilibria for finite imperfect information games, implemented in Gambit. However, the implementation suffers from issues of numerical instability making it unreliable [13]. Finally, Panozzo [15] proposed some approaches for the algorithmic verification of sequential equilibria.

In future work, we plan to investigate techniques to improve the performance of our approach. These could include improving the computation of the extreme directions of the set of cones, or modifying the equations to gain performance in the decomposition algorithm. There are also alternative characterizations of consistency that we could potentially use that do not rely on the extreme direction of polyhedral cones [3, 19]. We believe that this may allow us to solve slightly larger games such as Kuhn Poker [10].

ACKNOWLEDGMENTS

This work was supported by the ANR LabEx CIMI (grant ANR-11-LABX-0040) within the French State Programme "Investissements d'Avenir", and the EU ICT-48 2020 project TAILOR (No. 952215). We thank Robert Mattmüller, Umberto Grandi, Bernhard von Stengel, and Martin Antonov for their helpful feedback.

REFERENCES

[1] Charalambos D. Aliprantis and Owen Burkinshaw. 1990. Fundamentals of Real Analysis. In *Principles of Real Analysis (Second Edition)*. Academic Press, 1–47.

[2] Salman Azhar, Andrew McLennan, and John H. Reif. 2005. Computation of equilibria in noncooperative games. *Computers & Mathematics with Applications* 50, 5 (2005), 823–854.

[3] Francesc Dilmé. 2023. Lexicographic Numbers in Extensive Form Games. Work in progress, accessed via https://www.francescdilme.com on 2023-10-09.

[4] Andrew Gilpin and Tuomas Sandholm. 2006. Finding equilibria in large sequential games of imperfect information. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC 2006)*. ACM Press, 160–169.

[5] Moritz Graf, Thorsten Engesser, and Bernhard Nebel. 2024. Symbolic Computation of Sequential Equilibria. arXiv:2402.04452 [cs.GT]

[6] Moritz Graf, Thorsten Engesser, and Bernhard Nebel. 2024. A Symbolic Sequential Equilibria Solver for Game Theory Explorer. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*. IFAAMAS, 3 pages.

[7] Ebbe Hendon, Hans J. Jacobsen, and Birgitte Sloth. 1996. The One-Shot-Deviation Principle for Sequential Rationality. *Games and Economic Behavior* 12, 2 (1996), 274–282.

[8] Elon Kohlberg and Philip J. Reny. 1997. Independence on Relative Probability Spaces and Consistent Assessments in Game Trees. *Journal of Economic Theory* 75, 2 (1997), 280–313.

[9] David M. Kreps and Robert Wilson. 1982. Sequential Equilibria. *Econometrica* 50, 4 (1982), 863–894.

[10] Harold W. Kuhn. 1950. A simplified two-person poker. *Contributions to the Theory of Games* 1 (1950), 97–103.

[11] Steven M. LaValle. 2006. *Planning Algorithms*. Cambridge University Press.

[12] Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. 2022. Gambit: Software Tools for Game Theory, Version 16.0.2. http://www.gambit-project.org.

[13] Peter B. Miltersen and Troels B. Sørensen. 2006. Computing sequential equilibria for two-player games. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA 2006)*. ACM Press, 107–116.

[14] Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. The MIT Press.

[15] Fabio Panozzo. 2014. *Algorithms for the verification, computation and learning of equilibria in extensive-form games*. Ph.D. Dissertation. Polytechnic University of Milan, Italy.

[16] Andrés Perea. 2002. A note on the one-deviation property in extensive form games. *Games and Economic Behavior* 40, 2 (2002), 322–338.

[17] Carlos Pimienta. 2014. Bayesian and consistent assessments. *Economic Theory* 55, 3 (2014), 601–617.

[18] Rahul Savani and Bernhard von Stengel. 2015. Game Theory Explorer: software for the applied game theorist. *Computational Management Science* 12, 1 (2015), 5–33.

[19] Peter A. Streufert. 2012. *Additive Plausibility Characterizes the Supports of Consistent Assessments*. Technical Report 20123. University of Western Ontario, Department of Economics.

[20] Adam Strzebonski. 2023. Cylindrical Algebraic Decomposition. https://mathworld.wolfram.com/CylindricalAlgebraicDecomposition.html. Accessed on 2023-10-09.

[21] Theodore Turocy. 2010. Computing sequential equilibria using agent quantal response equilibria. *Economic Theory* 42, 1 (2010), 255–269.

[22] N Yu Zolotykh. 2012. New modification of the double description method for constructing the skeleton of a polyhedral cone. *Computational Mathematics and Mathematical Physics* 52, 1 (2012), 146–156.