

Extensive-Form Perfect Equilibrium Computation in Two-Player Games

Gabriele Farina

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
gfarina@cs.cmu.edu

Nicola Gatti

Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano
Piazza Leonardo da Vinci, 32
I-20133, Milan, Italy
nicola.gatti@polimi.it

Abstract

We study the problem of computing an Extensive-Form Perfect Equilibrium (EFPE) in 2-player games. This equilibrium concept refines the Nash equilibrium requiring resilience with respect to a specific vanishing perturbation, representing mistakes of the players at each decision node. The scientific challenge is intrinsic to the EFPE definition: it requires a perturbation over the agent form, but the agent form is computationally inefficient due to the presence of highly non-linear constraints. We show that the sequence form can be exploited in a non-trivial way and that, for general-sum games, finding an EFPE is equivalent to solving a suitably perturbed linear complementarity problem. We prove that Lemke’s algorithm can be applied, showing that computing an EFPE is PPAD-complete. In the notable case of zero-sum games, the problem is in FP and can be solved by linear programming. Our algorithms also allow one to find a Nash equilibrium when players cannot perfectly control their moves, being subject to a given execution uncertainty, as is the case in most realistic physical settings.

Introduction

Computing solutions of games is currently one of the hottest problems in computer science, as providing optimal strategies to *autonomous agents* interacting strategically is central in Artificial Intelligence (Shoham and Leyton-Brown 2008). Finding a Nash Equilibrium (NE)—the basic solution concept for non-cooperative games—is PPAD-complete even in 2-player games (Chen, Deng, and Teng 2009) and it is unlikely that there is a polynomial-time algorithm, since it is commonly believed that $FP \subset PPAD \subset FNP$. We recall that a search problem is in PPAD if there exists a *path-following* algorithm whose iterations have a polynomial-time cost. In the case of 2-player normal-form games, this algorithm is provided by (Lemke and Howson 1964).

Extensive-form games provide a richer representation of strategic interaction situations than the normal form. The study of extensive-form games is much more involved than that of normal-form games. A variation of Lemke-Howson’s algorithm, called *Lemke’s algorithm*, finds an NE in a 2-player extensive-form game showing that the problem is in

the PPAD class (Koller, Megiddo, and von Stengel 1996). However, the concept of NE is not satisfactory in extensive-form games, and NE refinements are studied (Selten 1975). When information is perfect, the concept of Subgame Perfect Equilibrium (SPE) is satisfactory, while it is not when information is imperfect. In this latter case, refinements are usually based on the idea of *perturbations* representing mistakes of the players. In a Quasi-Perfect Equilibrium (QPE)—proposed by van Damme—a player maximizes their utility in each decision node taking into account the future mistakes of the opponents *only*, whereas, in an Extensive-Form Perfect Equilibrium (EFPE)—proposed by Nobel prized Selten—, players maximize their utility in each decision node keeping into account the future mistakes of *both* themselves and their opponents (Hillas and Kohlberg 2002). The sets of QPEs and EFPEs may be disjoint, requiring different techniques. Given a specific perturbation, computing a QPE is PPAD-complete (Miltersen and Sørensen 2010) and can be done by summing the perturbation to the constant terms in the linear constraints of the sequence form (von Stengel 1996); due to this reason, we say that this perturbation is *additive*. However, the problem of efficiently computing an EFPE is still open. The scientific challenge is intrinsic to the EFPE definition: it is based on a perturbation over the agent form, but the agent form is computationally inefficient, presenting highly non-linear equilibrium constraints. The only previous attempt is (Gatti and Iuliano 2011), but no proof is provided about neither the soundness (specifically, the perturbations proposed do not match the definition of EFPE) nor polynomial-time cost of each algorithm iteration.

We show that finding an EFPE is PPAD-complete in 2-player general-sum games and can be done by means of Lemke’s algorithm with an extra polynomial computation cost due to a numeric perturbation, and that it is in FP in 2-player zero-sum games and can be done by linear programming with the same perturbation for the general-sum case. The table below summarizes the results known so far. ‘(*)’ denotes original contribution discussed in this paper.

Solution concept	General-sum	Zero-sum
Nash (NE)	PPAD-complete	FP
Subgame Perfect (SPE)	PPAD-complete	FP
Quasi Perfect (QPE)	PPAD-complete	FP
Extensive-Form Perfect (EFPE)	PPAD-complete (*)	FP (*)

In order to prove our main result, we provide also two original results of broader interest. First, we show that a perturbation over the agent form can be formulated as a specific symbolic perturbation over the coefficients of the variables of the sequence form (due to this reason, we say that this perturbation is *multiplicative*). This shows that computing an equilibrium when a player does not have perfect control over the execution of their moves along the game tree, as is customary for physical agents (e.g., robots) whose actions are subject to execution uncertainty, is PPAD-complete or in FP in general-sum and zero-sum games, respectively. Second, we show that we can turn the symbolically perturbed problem above into a numerically perturbed problem. We believe our approach to be particularly interesting, in that it not only applies to the computation of EFPEs, but rather is a more general framework, that can be used to derive, e.g., the results on QPEs in a more natural fashion. All omitted proofs can be found in (Farina and Gatti 2016).

Preliminaries

In the following, we adopt the notation introduced by (Shoham and Leyton-Brown 2008). We invite the reader unfamiliar with the topic to refer to (Shoham and Leyton-Brown 2008) or any other classic textbook on the subject for further information and context.

An *extensive-form game* Γ is defined over a game tree. In each non-terminal node a single player moves and each edge corresponds to an action available to the player. As customary, N denotes the set of players, A_i denotes the set of actions available to player i and a is an action, \mathbf{a} denotes the action profile of all the players and \mathbf{a}_{-i} denotes the action profile of the opponents of player i . Furthermore, H_i denotes the set of information sets of player i and h is an information set. Finally, $\iota(h)$ is the player that moves at h , $\rho(h)$ is the set of actions available at h to player $\iota(h)$, and function u_i returns the utility of player i from each terminal node.

The agent form (Selten 1975) of an extensive-form game is a tabular representation in which, for every player i and information set $h \in H_i$, there is a fictitious player called *agent* and all the agents of player i have the same utility from the terminal nodes. Player i 's strategy over action a , called *behavioral*, is denoted by $\pi_i(a) \geq 0$ and is such that for each h it holds $\sum_{a \in \rho(h)} \pi_{\iota(h)}(a) = 1$. The strategy of the agent playing at h is the restriction of $\pi_{\iota(h)}$ to actions $\rho(h)$. A behavioral strategy profile is denoted by π .

The concept of Extensive-Form Perfect Equilibrium (Selten 1975), also known as ‘‘Trembling hand perfect equilibrium’’, is defined on the agent form. We initially introduce the definitions of *perturbed game* (over the agent form) and Nash equilibrium of the agent form since they are necessary to introduce the definition of EFPE.

Definition 1. Let Γ be an extensive-form game and let $l(a)$ be a positive real called perturbation such that $\sum_{a \in \rho(h)} l(a) < 1$ for every h . An (agent-form) perturbed game (Γ, l) is an extensive-form game with the constraint that $\pi_{\iota(h)}(a) \geq l(a)$ for every h and $a \in \rho(h)$.

Definition 2. A behavioral strategy profile π is a Nash equilibrium of the agent form of Γ if, for every information set h , the behavioral strategy of the agent playing at h is best response to the strategies of all the other agents.

The problem of finding a Nash equilibrium of the agent form can be formulated as a non-linear complementarity problem (NLCP). This formulation is not useful in practice since the high non-linearity raises a number of computational issues. We now introduce the definition of EFPE.

Definition 3. A strategy profile π is an EFPE of Γ if it is a limit point of a sequence $\{\pi(l)\}_{l \downarrow 0}$ where $\pi(l)$ is a Nash equilibrium of the agent form of the perturbed game (Γ, l) .

Finally, we introduce the sequence form (von Stengel 1996). It provides a computationally efficient representation of an extensive-form game. The set of players of the sequence form is the same of that of the extensive form and each player i plays sequences $q \in Q_i$ of actions $a \in A_i$ over the game tree. There is a special sequence, denoted by q_0 and available to all the players, and all the other sequences $q \in Q_i$ are defined by induction extending some sequence $q' \in Q_i$, starting from q_0 , with an action $a \in A_i$. As customary, $qa \in Q_i$ denotes the sequence obtained by extending sequence $q \in Q_i$ with action $a \in A_i$. A sequence is called *terminal* if, combined with some sequence of the other players, leads to a terminal node, and *non-terminal* otherwise. With 2 players, U_i is the utility matrix of player i and $U_i(q_i, q_{-i})$ returns, when sequence profile (q_i, q_{-i}) leads to a terminal node (here $q_i \in Q_i$ and $q_{-i} \in Q_{-i}$), the utility of the node, and zero otherwise. The strategy of player i over sequence q is denoted by $r_i(q) \geq 0$ and is called *realization plan*. Finally, strategies are subject to some constraints: $r_i(q_0) = 1$ and, for every information set h and sequence q leading to h , $r_i(q) = \sum_{a \in \rho(h)} r_i(qa)$. For notational convenience, these constraints can be written as: $F_i r_i = f_i$, where F_i is an opportune matrix and f_i is a vector of zeros except for the first position whose value is one. Finally, we recall that a strategy profile r is *realization equivalent* to a strategy profile π when r and π induce the same probability distribution on the terminal nodes. Given profile r , a realization-equivalent π can be derived as $\pi_i(a) = r_i(qa)/r_i(q)$ if $r_i(q) > 0$ and $\pi_i(a)$ is any otherwise.

The definition of *Nash equilibrium of the sequence form* is standard, requiring each player to play their best response. In contrast to what happens in the agent form, the problem of finding a Nash equilibrium in the sequence form can be formulated as a linear complementarity problem (LCP) (Koller, Megiddo, and von Stengel 1996) and can be solved by means of Lemke’s algorithm (Lemke 1970). In particular, (Koller, Megiddo, and von Stengel 1996) show that by applying an opportune affine transformation of the players’ utility matrices the LCP satisfies two properties that allow Lemke’s algorithm to always terminate with a Nash equilibrium. This result, combined with the fact the computational cost of each pivoting step of Lemke’s algorithm is polynomial, shows that the problem of finding a Nash equilibrium of the sequence form is in the PPAD class. Importantly, it is known that, without any perturbation, any Nash equilibrium of the

sequence form is also a Nash equilibrium of the agent form, while, in presence of perturbations, this result may hold or not, depending on the definition of the specific perturbation used.

Extensive-Form Perfect Equilibria and LCPs

We initially show that introducing a specific perturbation over the realization-plan strategies is equivalent to introducing a perturbation over the behavioral strategies. For the sake of simplicity, we study the case in which $l(a) = \epsilon$ for every a , thus leading to a specific EFPE. All the results discussed in this section and in the following ones can be extended to the general case in which $l(a)$ is a polynomial in ϵ potentially different for each action a —we recall that considering only polynomial functions of ϵ is sufficient to find any EFPE, as discussed in (Blum, Brandenburger, and Dekel 1991; Govindan and Klumpp 2003).

Theorem 1. *A realization-plan strategy profile r is an EFPE of Γ if it is a limit point of a sequence $\{r(\epsilon)\}_{\epsilon \downarrow 0}$, where $r(\epsilon)$ is a Nash equilibrium of the sequence form of Γ under the constraint $r_i(qa) \geq \epsilon r_i(q)$ for every player i , sequence q , and action a .*

Proof. The proof is structured into two steps. In the first step, we show that requiring $r_i(qa) \geq \epsilon r_i(q)$ for every player i , sequence q , and action a in sequence form is equivalent to considering the perturbed game (Γ, l) where $l(a) = \epsilon$ for every action a . In the second step, we show that any Nash equilibrium in the sequence form of such a perturbed game is a Nash equilibrium in the agent form.

Focus on the first step. The empty sequence q_\emptyset is played with probability one and then, by induction, every sequence q is played with a strictly positive probability of at least $\epsilon^{|q|}$ where $|q|$ is the length in terms of actions of sequence q . Since the behavioral strategy $\pi_i(a)$ is defined as $\pi_i(a) = r_i(qa)/r_i(q)$, we have that requiring $r_i(qa) \geq \epsilon r_i(q)$ is equivalent to require $\pi_i(a) \geq \epsilon$ for every a . This completes the proof of the first step.

Focus on the second step. The proof follows from the definition of sequence form. Nevertheless, we report all the details. The expected utility (in the agent-form representation) $EU_{a_h}^{\text{AF}}$ provided by action $a_h \in \rho(h)$ to player $\iota(h) = i$ given π_{-h} is:

$$EU_{a_h}^{\text{AF}}(\pi_{-h}) = \sum_{\mathbf{a}_{-h} \in A_{-h}} U_i^{\text{AF}}(a_h, \mathbf{a}_{-h}) \prod_{h' \in H \setminus \{h\}} \pi_{\iota(h')}(a_{h'}).$$

We denote by U_i^{AF} the utility function of player i in the agent-form representation and by \mathbf{a}_{-h} the action profile in which only the action played at h is excluded. The expected utility (in the sequence-form representation) $EU_{a_h}^{\text{SF}}$ provided by sequence $qa_h \in Q_i$ where $a_h \in \rho(h)$ to player i given r_{-i} is:

$$EU_{a_h}^{\text{SF}}(r_{-i}) = \sum_{q': qa_h \in Q_i} \sum_{q'' \in Q_{-i}} U_i(q', q'') r_{-i}(q').$$

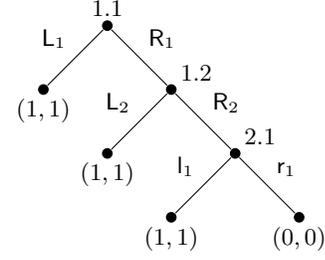


Figure 1: A sample game.

In the agent form, for each information set h , a Nash equilibrium assures that action a_h is played with probability $\sigma_{\iota(h)}(a_h) > \epsilon$ only if $EU_{a_h}^{\text{AF}}(\pi_{-h})$ is the maximum among the $EU_{a'_h}^{\text{AF}}(\pi_{-h})$ s for all $a'_h \in \rho(h)$. In the sequence form, for each information set h , a Nash equilibrium assures that sequence qa_h is played with $r_{\iota(h)}(qa_h) > \epsilon r_{\iota(h)}(q)$ only if $EU_{qa_h}^{\text{SF}}(r_{-i})$ is the maximum among the $EU_{qa'_h}^{\text{SF}}(r_{-i})$ s for all $a'_h \in \rho(h)$.

We show that the two families of constraints are the same except for an affine transformation not depending on the actions available at the information set h and preserving the maximum. At every information set h , it holds

$$EU_{a_h}^{\text{AF}}(\pi_{-h}) = \alpha EU_{a_h}^{\text{SF}}(r_{-i}) + \beta$$

for every $a_h \in \rho(h)$, where α_h, β_h do not depend on the actions available at h . More precisely,

$$\alpha = \prod_{a' \in q} \pi_i(a') = r_i(q) > \epsilon^{|q|} > 0, \text{ and}$$

$$\beta = \sum_{\substack{q': \exists a' \in \rho(h) \\ qa' \in q'}} \sum_{q'' \in Q_{-i}} U_i(q', q'') r_{-i}(q').$$

Therefore, the action a_h that maximizes $EU_{a_h}^{\text{AF}}(\pi_{-h})$ maximizes also $EU_{a_h}^{\text{SF}}(r_{-i})$ when π and r are realization equivalent. \square

The proof of Theorem 1 above shows that requiring the condition $r_i(qa) \geq \epsilon r_i(q)$ for every $i \in N, q \in Q_i, a \in A_i$ is equivalent to considering the perturbed game (Γ, l) where $l(a) = \epsilon$ for every action a . For notational convenience, such a condition can be expressed as:

$$R_i(\epsilon) r_i = \tilde{r}_i \geq 0,$$

where $R_i(\epsilon)$ is a matrix that we call *behavioral perturbation matrix* and \tilde{r}_i is the residual strategy (i.e., the strategy of the player once perturbation has been excluded).

Example 1. *Consider the sample game of Figure 1. Matrices $R_1(\epsilon)$ and $R_2(\epsilon)$ are as follows:*

$$R_1(\epsilon) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ -\epsilon & 1 & 0 & 0 & 0 \\ -\epsilon & 0 & 1 & 0 & 0 \\ 0 & 0 & -\epsilon & 1 & 0 \\ 0 & 0 & -\epsilon & 0 & 1 \end{pmatrix}, R_2(\epsilon) = \begin{pmatrix} 1 & 0 & 0 \\ -\epsilon & 1 & 0 \\ -\epsilon & 0 & 1 \end{pmatrix}.$$

We study the properties of matrix $R(\epsilon)$ and of its inverse.

Remark 1. Behavioral perturbation matrices are lower triangular square matrices having only 0 or $-\epsilon$ as entries.

Lemma 1. Let $R(\epsilon)$ be a $n \times n$ behavioral perturbation matrix. Then $R(\epsilon)$ is invertible, and its inverse is

$$R(\epsilon)^{-1} = I + \epsilon E(\epsilon),$$

where I is the identity matrix, and $E(\epsilon)$ is a lower triangular matrix whose entries are polynomials in ϵ having non-negative integer coefficients.

Example 2. For the matrix $R_1(\epsilon)$ of Example 1, we have:

$$R_1(\epsilon)^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ \epsilon & 1 & 0 & 0 & 0 \\ \epsilon & 0 & 1 & 0 & 0 \\ \epsilon^2 & 0 & \epsilon & 1 & 0 \\ \epsilon^2 & 0 & \epsilon & 0 & 1 \end{pmatrix}.$$

Now we are in the position to formulate the problem of finding an EFPE as a linear complementarity program.

Lemma 2. An EFPE is the limit point as $\epsilon \rightarrow 0$ of any solution of the perturbed standard-form LCP

$$P(\epsilon) : \begin{cases} \text{find} & z, w \\ \text{s.t.} & z^\top w = 0 \\ & w = M(\epsilon)z + b \\ & z, w \geq 0 \end{cases}$$

where (the underlined entries depend on ϵ)

$$z = \begin{pmatrix} \tilde{r}_1 \\ \tilde{r}_2 \\ \underline{v_1^\top} \\ \underline{v_1^\top} \\ \underline{v_2^\top} \\ \underline{v_2^\top} \end{pmatrix}, \quad b = \begin{pmatrix} 0 \\ 0 \\ f_1 \\ -f_1 \\ f_2 \\ -f_2 \end{pmatrix},$$

$$\underline{M} = \begin{pmatrix} 0 & -\underline{R_1}^{-\top} U_1 \underline{R_2}^{-1} & \underline{R_1}^{-\top} F_1^\top \\ -\underline{R_2}^{-\top} U_2^\top \underline{R_1}^{-1} & 0 & 0 \\ -\underline{F_1} \underline{R_1}^{-1} & 0 & 0 \\ \underline{F_1} \underline{R_1}^{-1} & 0 & 0 \\ 0 & -\underline{F_2} \underline{R_2}^{-1} & 0 \\ 0 & \underline{F_2} \underline{R_2}^{-1} & 0 \\ -\underline{R_1}^{-\top} F_1^\top & 0 & 0 \\ 0 & \underline{R_2}^{-\top} F_2^\top & -\underline{R_2}^{-\top} F_2^\top \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Proof. The proof directly follows from Theorem 1, the LCP above expressing the best-response conditions of the two players in the perturbed game. However, we report the complete derivation of the LCP, being useful for our treatment.

The problem of finding the best response of player i in the perturbed game (Γ, l) with $l(a) = \epsilon$ for every a is a linear problem, defined as

$$\text{BR}_i(\epsilon) : \begin{cases} \max_{r_i} & r_i^\top U_i r_{-i} \\ \text{s.t.} & F_i r_i = f_i \\ & R_i(\epsilon) r_i \geq 0 \end{cases}$$

Notice that $R_i(\epsilon)$ is invertible (Lemma 1), hence by changing variable, we find the equivalent problem:

$$\text{BR}_i(\epsilon) : \begin{cases} \max_{\tilde{r}_i} & \tilde{r}_i^\top R_i(\epsilon)^{-\top} U_i R_{-i}(\epsilon)^{-1} \tilde{r}_{-i} \\ \text{s.t.} & \textcircled{1} \quad F_i R_i(\epsilon)^{-1} \tilde{r}_i = f_i \\ & \textcircled{2} \quad \tilde{r}_i \geq 0 \end{cases}$$

Taking the dual:

$$\overline{\text{BR}}_i(\epsilon) : \begin{cases} \min_{v_i} & f_i^\top v_i \\ \text{s.t.} & \textcircled{3} \quad R_i(\epsilon)^{-\top} F_i^\top v_i \geq \\ & \quad R_i(\epsilon)^{-\top} U_i R_{-i}(\epsilon)^{-1} \tilde{r}_{-i} \\ & \textcircled{4} \quad v_i \text{ free in sign} \end{cases}$$

Finally, complementarity slackness requires that

$$\textcircled{5} \quad \tilde{r}_i^\top (R_i(\epsilon)^{-\top} F_i^\top v_i - R_i(\epsilon)^{-\top} U_i R_{-i}(\epsilon)^{-1} \tilde{r}_{-i}) = 0.$$

Solving problem $\text{BR}_i(\epsilon)$ or $\overline{\text{BR}}_i(\epsilon)$ is equivalent to solving the feasibility problem defined by constraints $\textcircled{1}$ to $\textcircled{5}$. It is now easy to see that we can cast the problem of satisfying conditions $\textcircled{1}$ to $\textcircled{5}$ for both players as a standard-form LCP whose parameters are as defined in this lemma. \square

We conclude this section with a couple of lemmas that we will use in the following sections.

Lemma 3. Consider the LCP formulation of Lemma 2, where ϵ is treated as a symbolic variable, so that the entries of $M(\epsilon)$ are polynomials in ϵ . A number of bits polynomial in the input game size is sufficient to store all coefficients appearing in $P(\epsilon)$.

Lemma 4. Let $\nu = \max_{h \in \cup_i H_i} \{|\rho(h)|\}$ be the maximum number of actions available at an information set. If $0 \leq \epsilon \leq 1/n$, there always exists a realization-plan strategy r_i such that $F_i R_i(\epsilon)^{-1} r_i = f_i$.

Perturbed LCPs

Before turning our attention to the *computational* aspects of finding an EFPE, we introduce some general concepts, pertaining to perturbed linear optimization problems. While we target the development of these concepts with our specific use-case in mind, it should be noted that this section's definitions and lemmas are of broader interest, being applicable to any linear program (LP) or LCP. We recall that a *basis* \mathcal{B} for a standard-form LP with constraints $Mx = b$ or a standard-form LCP with linear equality constraints $w = Mz + b$ is a set of linearly independent columns of M such that the associated solution (called *basic solution*) is feasible.

Definition 4 (Negligible positive perturbation (NPP)). Let $P(\epsilon)$ be an LCP dependent on some perturbation ϵ . The value $\epsilon^* > 0$ is a negligible positive perturbation (NPP) if any optimal basis \mathcal{B} for $P(\epsilon^*)$ is optimal for $P(\epsilon)$, for all $\epsilon \in (0, \epsilon^*]$.

Definition 5 (Optimality certificate for a basis). Given an LCP $P(\epsilon)$, and a basis \mathcal{B} for it, we call the finite-dimensional column vector $\mathcal{C}_{\mathcal{B}}(\epsilon)$ an optimality certificate for \mathcal{B} if for all $\epsilon > 0$

$$\mathcal{C}_{\mathcal{B}}(\epsilon) \geq 0 \iff \mathcal{B} \text{ is optimal for } P(\epsilon).$$

Lemma 5. *In the case of a perturbed LCP in standard form*

$$P(\epsilon) : \begin{cases} \text{find} & z, w \\ \text{s.t.} & \textcircled{1} \quad z^\top w = 0 \\ & \textcircled{2} \quad w = M(\epsilon)z + b(\epsilon) \\ & \textcircled{3} \quad z, w \geq 0 \end{cases}$$

an optimality certificate for the complementary basis \mathcal{B} is

$$\mathcal{C}_{\mathcal{B}}(\epsilon) = B(\epsilon)^{-1}b(\epsilon)$$

where B is the basis matrix corresponding to \mathcal{B} .

Proof. Since the basis \mathcal{B} is complementary by hypothesis, constraint $\textcircled{1}$ is always satisfied. Constraint $\textcircled{2}$ is satisfied by the definition of $B(\epsilon)$. Constraint $\textcircled{3}$ is satisfied if and only if $B(\epsilon)^{-1}b(\epsilon) = \mathcal{C}_{\mathcal{B}}(\epsilon) \geq 0$. \square

Finally, before proceeding, we introduce three mathematical lemmas that come in handy when dealing with optimality certificates. Indeed, it is often the case that $\mathcal{C}_{\mathcal{B}}(\epsilon)$ has polynomial or rational functions (with respect to ϵ) as entries.

Lemma 6. *Let $p(\epsilon) = a_0 + a_1\epsilon^1 + \dots + a_n\epsilon^n$ be a real polynomial such that $a_0 \neq 0$, and let $\mu = \max_i |a_i|$. Then $p(\epsilon)$ has the same sign of a_0 for all $0 \leq \epsilon \leq \epsilon^*$, where*

$$\epsilon^* = |a_0|/(\mu + |a_0|).$$

As a corollary, we can easily extend the result of Lemma 6 to rational functions.

Lemma 7. *Let*

$$p(\epsilon) = \frac{a_0 + a_1\epsilon^1 + \dots + a_n\epsilon^n}{b_0 + b_1\epsilon^1 + \dots + b_m\epsilon^m}$$

be a rational function such that $a_0, b_0 \neq 0$, and let $\mu_a = \max_i |a_i|$, $\mu_b = \max_i |b_i|$. Then $p(\epsilon)$ has the same sign of a_0/b_0 for all $0 \leq \epsilon \leq \epsilon^$, where*

$$\epsilon^* = \min\{|a_0|/(\mu + |a_0|), |b_0|/(\mu + |b_0|)\}.$$

Proof. The proof follows immediately by applying Lemma 6 to the numerator and the denominator of $p(\epsilon)$. \square

Lemma 8. *Let*

$$p(\epsilon) = \frac{a_0 + a_1\epsilon^1 + \dots + a_n\epsilon^n}{b_0 + b_1\epsilon^1 + \dots + b_m\epsilon^m}$$

be a rational function with integer coefficients, where the denominator is not identically zero; let $\mu_a = \max_i |a_i|$, $\mu_b = \max_i |b_i|$, $\mu = \max\{\mu_a, \mu_b\}$ and $\epsilon^ = 1/(2\mu)$. Then exactly one of the following holds:*

- $p(\epsilon^*) = 0$ for all $0 < \epsilon \leq \epsilon^*$,
- $p(\epsilon^*) > 0$ for all $0 < \epsilon \leq \epsilon^*$,
- $p(\epsilon^*) < 0$ for all $0 < \epsilon \leq \epsilon^*$.

Computation of Extensive-Form Perfect Equilibria

We finally delve into the computational details of finding Extensive-Form Perfect Equilibria. The central result of this section (Theorem 2) roughly states that the EFPE LCP (Lemma 2) always admits a “small” NPP. Leveraging this fact, we quickly derive a path-following algorithm for the computation of EFPE in general-sum games in which each pivoting step has a polynomial-time cost (Theorem 3), and a polynomial-time algorithm for the zero-sum counterpart (Theorem 4). These two algorithms put the two search problems in PPAD and FP classes, respectively.

We start by showing that, as long as the perturbation ϵ is “reasonably small”, the LCP defined in Lemma 2 always admits a solution. In particular:

Lemma 9. *If $0 < \epsilon \leq 1/\nu$, where $\nu = \max_{h \in \cup_i H_i} \{|\rho(h)|\}$ is the maximum number of actions available at an information set, Lemke’s algorithm always finds a solution for $P(\epsilon)$.*

We remark that when ϵ is a given value, the task of finding an NE of $P(\epsilon)$ has a powerful interpretation. Indeed, it captures the situation in which the moves of a player are subject to execution uncertainty and therefore a player cannot perfectly control their actions.

Theorem 2. *Given a (general-sum) two-player game Γ with $\nu = \max_{h \in \cup_i H_i} \{|\rho(h)|\}$, the problem $P(\epsilon)$ of determining any EFPE for Γ admits an NPP $\epsilon^* \leq 1/\nu$ that can be computed from Γ in polynomial time. In particular, $\epsilon^* = 1/V^*$, where the integer value V^* can be represented in memory with a number of bits polynomial in the input game size.*

Proof. We illustrate the steps that lead to the determination of such V^* . The central idea is as follows: we want to determine ϵ^* so that, whatever the feasible base \mathcal{B} for $P(\epsilon^*)$ may be, the optimality certificate for ϵ^* is positive for all $\epsilon \in (0, \epsilon^*]$. Indeed, it is immediate to see that such ϵ^* is necessarily an NPP.

Optimality certificate. We begin by studying the optimality certificate for the LCP $P(\epsilon)$, that is, by Lemma 5,

$$B(\epsilon)^{-1}b(\epsilon) \geq 0,$$

where $B(\epsilon)$ is base matrix corresponding to the feasible base \mathcal{B} found by Lemke’s algorithm. Introducing $C(\epsilon) = \text{cof } B(\epsilon)$, the cofactor matrix of matrix $B(\epsilon)$, and leveraging the well-known identity $B(\epsilon)^{-1} = C(\epsilon)^\top / \det B(\epsilon)$, we can rewrite the optimality certificate above as

$$\frac{C(\epsilon)^\top b(\epsilon)}{\det B(\epsilon)} \geq 0.$$

The vectorial condition above is equivalent to a system of n scalar conditions, each of the form

$$f_i(\epsilon) = \frac{c_i(\epsilon)^\top b(\epsilon)}{\det B(\epsilon)} \geq 0,$$

where $c_i(\epsilon)$ is the i -th row of $C(\epsilon)^\top$. Evidently, $f_i(\epsilon)$ is a rational function in ϵ , having only integer coefficients, for all $i = 1, \dots, n$.

Denominator coefficients. We now give an upper bound on the coefficients of the denominator of $f_i(\epsilon)$, that is $\det B(\epsilon)$. Let V_B be the largest coefficient that could potentially appear in $B(\epsilon)$ and $b(\epsilon)$, and let m be the largest polynomial degree appearing in $B(\epsilon)$. Notice that $m \in O(\text{poly}(n))$. By using Hadamard's inequality, we can write

$$\text{coeff}(\det B(\epsilon)) \leq n^{n/2} V_B^n \text{coeff}((1 + \epsilon + \dots + \epsilon^m)^n),$$

where $\text{coeff}(\cdot)$ is the largest coefficient of its polynomial argument. Since $\text{coeff}((1 + \epsilon + \dots + \epsilon^m)^n) \leq m^n$, we have

$$\text{coeff}(\det B(\epsilon)) \leq V_D := n^{n/2} (m V_B)^n.$$

Notice that this bound is valid for all possible basis matrices $B(\epsilon)$. Furthermore, notice that

$$\log V_D = n/2 \log n + n \log m + n \log V_B,$$

and by Lemma 3 we conclude that V_D requires a number of bits polynomial in the input game size in order to be stored in memory.

Numerator coefficients. Since the elements of $c_i(\epsilon)$ are cofactors for $B(\epsilon)$, they are upper-bounded by $\det B(\epsilon)$, which in turn is upper-bounded by V_D . Therefore,

$$\text{coeff}(c_i(\epsilon)^\top b(\epsilon)) \leq V_N := V_B V_D.$$

Again, it is worthwhile to notice that this bound is valid for all possible basis matrices $B(\epsilon)$.

Wrapping up. Define $V^* = 2 \max\{V_N, V_D\} = 2V_B V_D$. We now argue that $\epsilon^* = 1/V^*$ is an NPP for $P(\epsilon)$. Indeed, let \mathcal{B}^* be a feasible basis¹ for $P(\epsilon^*)$, and let $B(\epsilon^*)$ be the corresponding base matrix. Being \mathcal{B} feasible for $P(\epsilon^*)$, each row f_i in the optimality certificate is non-negative when evaluated at ϵ^* for all i . Therefore, we know from Lemma 8 that $f_i(\epsilon) \geq 0$ in $(0, 1/V^*] = (0, \epsilon^*]$. Hence, the optimality certificate for \mathcal{B} is non-negative for all $0 < \epsilon \leq \epsilon^*$, which is equivalent to say that ϵ^* is an NPP. Finally, note that $V^* = 2V_B V_D$ be stored in memory with a number of bits polynomial in the game size. This completes the proof. \square

Theorem 3. *The problem of determining an EFPE of a general-sum two-player game Γ is PPAD-complete.*

Proof. Let $\epsilon^* = 1/V^*$ be an NPP as defined in Theorem 2, and let \mathcal{B} be a feasible base for the (numerical) problem $P(\epsilon^*)$, found using Lemke's algorithm. Since ϵ^* is an NPP, the pair of strategies (π_1^*, π_2^*) corresponding to \mathcal{B} retain their feasibility with respect to the LCP $P(\epsilon)$ as $\epsilon \rightarrow 0$, meaning that (π_1^*, π_2^*) is in fact an EFPE.

Furthermore, given that V^* requires a number of bits polynomial in the input game size, each iteration of Lemke's algorithm takes time polynomial in the game size. This proves that the algorithm described is a path-following algorithm requiring a polynomial-time cost at each step, and therefore the problem of finding an EFPE in two-player games is in the PPAD class. The hardness easily follows from the fact that EFPE is a refinement of Nash equilibrium, an EFPE always exists, and finding a Nash is PPAD-complete. Therefore, if finding an EFPE were not PPAD-hard, then one could use the EFPE-finding algorithm with

¹Notice that since $\epsilon^* \leq 1/n$, \mathcal{B} always exists (see Lemma 9).

the aim of finding an NE and therefore not even finding an NE would be PPAD-hard. This concludes the proof. \square

We remark that the proof of theorem above also applies for an arbitrary ϵ (potentially non-NPP), showing that finding an NE for any $\epsilon < \max_{h \in \cup_i H_i} \{|\rho(h)|\}$ is in the PPAD class. We summarize the procedure to find an EFPE of general-sum games in Algorithm 1.

Algorithm 1

procedure FIND-EFPE

1. Compute ϵ^* from Γ as in the proof of Theorem 2
 2. Determine a basis \mathcal{B} for the numerical LCP $P(\epsilon^*)$
 3. Let $B(\epsilon)$ be the base matrix corresponding to \mathcal{B} in $P(\epsilon)$, as ϵ varies
 - \triangleright Since $B(\epsilon)^{-1}b$ is a rational bounded function in a neighborhood of 0, $B(0)^{-1}b$ exists
 4. $(\tilde{r}_1, \tilde{r}_2, v_1^+, v_1^-, v_2^+, v_2^-)^\top = B(0)^{-1}b$
 - \triangleright Note that $R_1(0) = R_2(0) = I$, so $\tilde{r}_1 = r_1, \tilde{r}_2 = r_2$
 5. **return** the pair of strategies $(\tilde{r}_1, \tilde{r}_2)$
-

The approach we use in Theorems 2 and 3 extends the one used in (Miltersen and Sørensen 2010). More precisely, in (Miltersen and Sørensen 2010) the authors consider a numerical perturbation that sums to the constant terms of an LP to find a QPE of a zero-sum game, while in Theorem 2 we consider perturbations over the coefficient of the variables of an LCP to find an EFPE of general-sum games (and, below, of zero-sum games). The two approaches can be extended to find a QPE in general-sum games by using a numerical perturbation (the description is omitted here, as it is beyond the scope of this paper). We now show that Algorithm 1 requires polynomial time when the game is zero sum.

Theorem 4. *The problem of determining an EFPE of a zero-sum two-player game Γ can be solved in polynomial time in the size of the input game.*

Proof. Like in Theorem 3, we can easily extract an EFPE by looking at the feasible matrix \mathcal{B} which solves the (numerical) LCP $P(\epsilon^*)$. However, in the zero-sum setting, we do not need to use Lemke's algorithm. Indeed, notice that in zero-sum two-player games, matrix $M(\epsilon)$ as defined in Lemma 2 is such that $M(\epsilon) + M(\epsilon)^\top = 0$ for all ϵ , because $U_2 = -U_1^\top$. Therefore, the complementarity condition can be rewritten as

$$\begin{aligned} z^\top (M(\epsilon)z + b(\epsilon)) &= z^\top M(\epsilon)z + z^\top b(\epsilon) \\ &= \frac{1}{2} z^\top (M(\epsilon) + M(\epsilon)^\top) z + z^\top b(\epsilon) \\ &= z^\top b(\epsilon) = 0, \end{aligned}$$

a linear condition instead of a quadratic one. This shows that when the game is zero-sum, the LCP is actually an LP. As such, a basis for the LCP of Algorithm 1 can be computed in polynomial time, leading to an overall polynomial time algorithm. This completes the proof. \square

Conclusion and Future Work

In this paper, we provide a path-following algorithm to find an EFPE in 2-player games. Our algorithm requires the application of Lemke’s algorithm to a numerically perturbed LCP. We show that the computation cost of each iteration of the algorithm is polynomial, and this shows that finding an EFPE in 2-player games is PPAD-complete. We also show that in the notable case of 2-player zero-sum games, linear programming can be used and that the problem is in the FP class. In order to achieve our result, we also develop two accessory results. The first one shows that the problem of finding a Nash equilibrium when a player does not perfectly control her moves being subject to mistakes, as it happens in practice for physical agents, is PPAD-complete and can be done by means of our algorithm. The second one is an extension of the characterization of numerically perturbed LCPs in which even the coefficients of the variables are perturbed.

In future works, we aim to extend our accessory results as well as to study the verification problem for an EFPE in 2-player games (that is, the problem of deciding whether a strategy profile given in input is an EFPE).

References

- Blum, L.; Brandenburger, A.; and Dekel, E. 1991. Lexicographic probabilities and equilibrium refinements. *ECONOMETRICA* 59(1):81–98.
- Chen, X.; Deng, X.; and Teng, S.-H. 2009. Settling the complexity of computing two-player nash equilibria. *J. ACM* 56(3):14:1–14:57.
- Farina, G., and Gatti, N. 2016. Extensive-Form Perfect Equilibrium Computation in Two-Player Games. *ArXiv e-prints*. <https://arxiv.org/abs/1611.05011>.
- Gatti, N., and Iuliano, C. 2011. Computing an extensive-form perfect equilibrium in two-player games. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*.
- Govindan, S., and Klumpp, T. 2003. Perfect equilibrium and lexicographic beliefs. *INT J GAME THEORY* 31(2):229–243.
- Hillas, J., and Kohlberg, E. 2002. Chapter 42 foundations of strategic equilibrium. volume 3 of *Handbook of Game Theory with Economic Applications*. Elsevier. 1597–1663.
- Koller, D.; Megiddo, N.; and von Stengel, B. 1996. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior* 14(2):247–259.
- Lemke, C. E., and Howson, J. T. 1964. Equilibrium points of bimatrix games. *SIAM Journal on Applied Mathematics* 12(2):413–423.
- Lemke, C. 1970. Recent results on complementarity problems. *Nonlinear programming* 349–384.
- Miltersen, P., and Sørensen, T. 2010. Computing a quasi-perfect equilibrium of a two-player game. *ECON THEOR* 42(1):175–192.
- Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International journal of game theory* 4(1):25–55.
- Shoham, Y., and Leyton-Brown, K. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. New York, NY, USA: Cambridge University Press.
- von Stengel, B. 1996. Efficient computation of behavior strategies. *Games and Economic Behavior* 14(2):220–246.