Computing lexicographically safe Nash equilibria in finite two-person games with tight game forms given by oracles

Lecture series

National Research University Higher School of Economics, Moscow, Russia

Vladimir Gurvich[1, 2]

July 18,23,25 2024

[1] National Research University Higher School of Economics, Moscow, Russia,
[2] Rutgers University, NJ, USA

Gurvich, V. and Naumova, M., 2023.
Computing lexicographically safe Nash equilibria
in finite two-person games with tight game forms given by oracles.
Discrete Applied Mathematics, 340, pp.53-68.

Vladimir Gurvich

Mariya Naumova

Introduction

Here we outline main results. Precise definitions will be given later.

Consider a finite n-person game *in normal form* representing it as a pair (g, u), where u is the *payoff* function of n players and g is the so-called *game form*.

The latter can be viewed as a game without payoffs, which are not given yet.

Such approach is standard and convenient:

game form g “is responsible" for structural properties of game (g, u), which hold for any payoff u.

For example, game form g is called *Nash-solvable* if game (g, u) has a Nash equilibrium (NE) in pure strategies for every payoff u.

In 1950 Nash proved that every n-person normal form game has a NE in *mixed strategies* [**Nas50**, **Nas51**].
Yet, there are large families of games solvable in *pure strategies*, for example, *finite n-person positional (graphical) games with perfect information*.

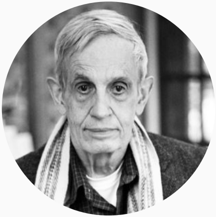

John Forbes Nash, Jr.

Bargaining with Gods. https://www.mathnet.ru/php/archive.phtml?wshow=paper&jrnid=mp&paperid=860&option_lang=rus

Its game structures Γ uniquely defines a finite n-person game form $g(\Gamma)$ that is Nash-solvable.

We expand the set of outcomes including not only terminal positions but also other strongly connected components of the corresponding directed graph.

Doing so, we also expand substantially the corresponding family of game forms, which remain Nash-solvable, but only in case of two players, $n = 2$.

Yet, perfect information is only sufficient but not necessary for Nash-solvability. A concept of *tightness* fits much better. This property is of algebraic nature.

It was introduced in [**Gur73**, **Gur75**] and in the latter paper it was shown that a finite two-person game form is Nash-solvable if and only if it is *tight*.

Note that already for $n = 3$ tightness is neither necessary nor sufficient for Nash-solvability. These results were obtained in [**Gur75**, **Gur89**]; several different proofs were given later [**BBM90**, **BG03**, **DS91**, **Gur18**, **GK18**].

Tightness remains necessary (and, of course, sufficient) for Nash-solvablity in the zero-sum case too. This result was obtained earlier: it follows easily from the so-called *Bottleneck Extrema Theorem* by Edmonds and Fulkerson [**EF70**]; see also [**Gur73**].

Jack Edmonds

Delbert Ray Fulkerson

Here we suggest a new (and much simpler) proof of the general result.

We introduce a concept of *lexicographically safe (lexsafe)* pure strategy of a player in a given game (g, u).

This is a refinement of the standard concept of a safe (maxmin) strategy that maximizes the worst possible outcome, while the lexsafe strategy realizes the lexicographical maximum of all possible outcomes.

Thus, the lexsafe strategies are most safe, but may be not rational.

(For comparison, recall that NE may be not Pareto optimal.)

One can view this as a price of stability.

We prove that a NE appears whenever one player applies a lexsafe strategy, while the opponent chooses some special *best response* to it. Yet, if both players choose their lexsafe strategies then the obtained pair may be not a NE.

Thus, there are two types of NE: lexsafe for one or for the other player.

These NE may coincide. For example, it happens in the zero-sum case, or when the considered game has a unique NE.

By definition, lexsafe strategies of a player do not depend on the payoff of the opponent; the player may be just unaware of it.

This is an interesting property important for applications.

In the proof of [**Gur75**, **Gur89**] the lexsafe strategies were implicitly constructed by an iterative algorithm increasing strategies in a lexicographical order.

Here we suggest a simple polynomial algorithm searching for a lexsafe strategy of a player and for a corresponding NE.

Such algorithm is obvious when a game form g is given explicitly.

Yet, in applications g is frequently given by an oracle $\mathcal{O}$, which size may be logarithmic in the size of g. We assume that this oracle solves in polynomial time any two-person game (g, u^0) in which payoff u^0 is zero-sum and takes only values ± 1; oracle $\mathcal{O}$ tells us who wins and determines a winning strategy.

Based on this assumption, we provide an algorithm computing a lexsafe NE in an arbitrary game (g, u) in time polynomial in the size of $\mathcal{O}$.

In the last section we consider four examples of such oracles from different areas of game theory and show that all four satisfy the above assumption.

Basic definitions

Game forms We consider finite, not necessarily zero-sum, normal form games of two players, Alice and Bob.

They have finite sets of strategies X and Y, respectively.

A *game form* is a mapping $g : X \times Y \to \Omega$, where Ω is a finite set of outcomes.

A two-person game (in normal form) is a triplet $(g; u, w)$,
where $u : \Omega \to \mathbb{R}$ and $w : \Omega \to \mathbb{R}$ are payoff functions of Alice and Bob.

To separate game forms and payoffs is an efficient idea.

Game forms are responsible for the so-called "structural properties" of the games, which hold for all payoffs, for example, Nash-solvability.

Nine examples are given in the next figure, where game forms are represented by tables with rows $x \in X$, columns $y \in Y$, and outcomes $\omega \in \Omega$.

ω_1	ω_1
ω_2	ω_3

g_1

ω_1	ω_1	ω_2	ω_2
ω_3	ω_4	ω_3	ω_4

g_2

ω_1	ω_1	ω_3
ω_1	ω_2	ω_2
ω_3	ω_2	ω_3

g_3

ω_1	ω_1	ω_3
ω_1	ω_1	ω_2
ω_4	ω_2	ω_2

g_4

ω_1	ω_2	ω_1	ω_2
ω_3	ω_4	ω_4	ω_3
ω_1	ω_4	ω_1	ω_5
ω_3	ω_2	ω_6	ω_2

g_5

ω_1	ω_1
ω_1	ω_2

g_6

ω_1	ω_2
ω_2	ω_1

g_7

ω_1	ω_1	ω_2
ω_3	ω_4	ω_3

g_8

ω_1	ω_1	ω_2
ω_4	ω_5	ω_2
ω_4	ω_3	ω_3

g_9

Mapping g is assumed to be surjective, but not necessarily injective, that is, an outcome $\omega \in \Omega$ may occupy an arbitrary array in the table of g.

A pair of strategies (x, y) is called a *situation* (term "strategy profile" is also used in literature). Sets

$g(x) = \{g(x, y) \mid y \in Y\}$ and $g(y) = \{g(x, y) \mid x \in X\}$

are called the *supports* of strategies $x \in X$ and $y \in Y$, respectively.

A strategy is called *minimal* if its support is not a proper superset of the support of any other strategy. For example, in g_6 the first strategies of Alice and Bob are minimal, while the second are not; in the remaining eight game forms all strategies are minimal. Moreover, any two strategies of a player have distinct supports, for every game form, except g_7.

A situation (x, y) is called *simple* if $g(x) \cap g(y) = \{g(x, y)\}$. For example, all situations of game forms g_1, g_2, g_8, g_9 are simple (such game forms are called *rectangular*); in contrast, no situation is simple in g_7; in g_3 all are simple, except three on the main diagonal; in g_4 all are simple, except the central one; in g_6 all are simple, except one with the outcome ω_2.

Payoffs and games in normal form

Payoffs of Alice and Bob are defined by real valued mappings

$u : \Omega \to \mathbb{R}$ and $w : \Omega \to \mathbb{R}$, respectively. Both players are maximizers.

Triplet $(g; u, w)$ defines a *finite two-person game in normal form*, or just a *game*, for short. Game $(g; u, w)$ and payoffs (u, w) are called:

- *zero-sum* if $u + w = 0$, that is, $u(\omega) + w(\omega) = 0$ for all $\omega \in \Omega$;
- *zero-sum* ± 1 (or just ± 1, for short) if, in addition, functions u and w take only two values 1 and -1.

Alternatively, a ± 1 payoff can be given by a partition $\Omega = \Omega_A \cup \Omega_B$, where Ω_A and Ω_B are the outcomes preferred by Alice and by Bob, respectively:

$$u(\omega) = 1, w(\omega) = -1 \text{ for } \omega \in \Omega_A \text{ and } u(\omega) = -1, w(\omega) = 1 \text{ for } \omega \in \Omega_B.$$

For a ± 1 game notation $(g; \Omega_A, \Omega_B)$ will be used along with $(g; u, w)$.

There exist only 2^k distinct ± 1 payoffs versus $(k!)^2$ orders, where $k = |\Omega|$.

Nash equilibria and saddle points

Given a game $(g; u, w)$, a situation (x, y) of its game form $g : X \times Y \to A$ is called a *Nash equilibrium* (NE) if

$$u(g(x, y)) \geq u(g(x', y)), \forall x' \in X, \text{ and } w(g(x, y)) \geq w(g(x, y')), \forall y' \in Y;$$

that is, if neither Alice nor Bob can profit replacing her/his strategy provided the opponent keeps his/her one unchanged, or in other words,

if x is a best response for y and y is a best response for x.

Note that a best response may be not unique.

This concept of solution was introduced in 1950 by John Nash [**Nas50**, **Nas51**].

Zero-sum case; Matrix games; maxmin, minmax, and saddle point

In this case, a game is called matrix, and a NE in it is called a *saddle point* (SP).

This concept was known much earlier, it is about 200 years old. NE, its natural extension to the non-zero-sum case, is 75 years old.

Alice is a maximizer, while Bob is a minimizer.

An SP is a **situation** (not a number!) (x, y), where $x \in X, y \in Y$ are the strategies of Alice and Bob such that the payoff $u(x, y)$ realizes a minimum in the row x and a maximum in the column y, respectively.

Here u is a payoff of Alice. The payoff of Bob is by default $w = -u$.

By the above definition a saddle point is a (Nash) equilibrium.

Alice (Bob) has no motivation to change her (his) strategy x (resp. y) provided the opponent, Bob (Alice) will keep his (her) strategy y (resp. x) unchanged.

Lemma

(i) An SP may be not unique; the set of SPs is a box $X^* \times Y^* \subseteq X \rtimes Y$, where X^* and $Y*$ are called the *optimal* strategies of Alice and Bob, respectively.

(ii) The saddle point value $u(x, y)$ is the same for all $x \in X$ and $y \in Y$. It is called the value of the matrix game.

Thus, any pair of optimal strategies $x \in X^*$ and $y \in Y^*$ forms a SP with the same value v, which is called the value of the considered matrix game.

Prove the lemma. Note that it cannot be extended to the non-zero-sum case.

Consider game form g_6; assume that Alice and Bob both prefer ω_2 to ω_1. Then,

(i) NE form the main diagonal, not a box.

(ii) There are two distinct NE values.

In a matrix game, *maxmin* and *minmax* are defined by formulas:

$$\max\min = \max_{x \in X} \min_{\omega \in g(x)} u(\omega);\ \min\max = \min_{y \in Y} \max_{\omega \in g(y)} u(\omega); \tag{1}$$

They are numbers, not situations. In other words, given a matrix game,

take a minimum in each row and then the maximum of all obtained minima; this is maxmin, it is the best result that Alice can guarantee, independently of Bob's actions;

take a maximum in each column and then the minimum of all obtained maxima; this is minmax, it is the best result that Bob can guarantee, independently of Alice's actions.

Lemma For every matrix (zero-sum) game (g, u, w) we have:

(j) $\textit{maxmin} \le \textit{minmax}$.

(jj) there exists a saddle point if and only if maxmin = minmax.

Prove the lemma.

Remark

*In [***Nas50**, **Nas51***] solvability in mixed strategies is studies. In contrast, we restrict the players to their pure strategies. Such approach is considered, for example, in [***AFPT10**, **AHMS08**, **BBM90**, **BCGM20**, **BFT17**, **BG03**, **BGMP10**, **BGMS07**, **FHKOY20**, **Gur75**, **Gur89**, **Gur17**, **Gur18**, **GK18**, **GN22**, **HKM21**, **Kuk11**, **MS96**, **Ros73**, **SM13**, **THS12**, **Was90***].*

Solvability of game forms

A game form g is called:

(i) *Nash-*, (ii) *zero-sum-*, (iii) ± 1-*solvable*

if the corresponding game $(g; u, w)$ has a NE for

(i) all, (ii) all zero-sum, (iii) all zero-sum ± 1 payoffs, respectively.

Implications $(i) \Rightarrow (ii) \Rightarrow (iii)$ are obvious.

In fact, all three properties are equivalent [**Gur75**, **Gur89**, **Gur09**].

For (ii) and (iii) it was shown earlier by Edmonds and Fulkerson [**EF70**]; see also [**Gur73**].

The list of equivalent properties (i), (ii), (iii) was extended in [**Gur75**] by adding the so-called *tightness*; see below.

Replacing payoffs by preferences and eliminating ties

Given a game $(g; u, w)$, we can assume wlog that payoffs $u : \Omega \to \mathbb{R}$ and $w : \Omega \to \mathbb{R}$ have no ties. Indeed, one can get rid of all ties by arbitrarily small perturbations of values of u and w. In accordance with definition, the set of NE will be either unchanged or reduced by such perturbations.

We focus on Nash-solvability (in pure strategies), that is we study conditions that guarantee the existence of a NE for arbitrary payoffs u and w.

Hence, we can wlog assume that both, u and w, have no ties and replace them by linear orders $\succ_A$ and $\succ_B$ over the set of outcomes Ω, which are called the *preferences* of Alice and Bob, respectively.

Thus, game $(g; u, w)$ can be replaced by $(g; \succ_A, \succ_B)$ and it is enough to study Nash-solvability of the latter.

Althgough some NE of $(g; u, w)$ may disappear in $(g; \succ_A, \succ_B)$, yet, Nash-solvability holds or fails for both games simultaneously.

Remark

Above arguments would fail in the case of mixed strategies.

Tight game forms

Mappings $\phi : X \to Y$ and $\psi : Y \to X$ are called *response strategies* of Bob and Alice, respectively. The motivation for this name is clear:

a player chooses his/her strategy as a function of a known strategy of the opponent.

Standardly, $gr(\phi)$ and $gr(\psi)$ denote the graphs of mappings ϕ and ψ in $X \times Y$.

Game form $g : X \times Y \to \Omega$ is called *tight* if

(I) $g(gr(\phi)) \cap g(gr(\psi)) \neq \emptyset$ for any mappings ϕ and ψ.

It is easy to verify that in Figure 1 the first six game forms ($g_1 - g_6$) are tight, while the last three ($g_7 - g_9$) are not.

In [**EF70**, **Gur73**, **Gur75**, **Gur89**, **Gur18**] the reader can find several equivalent properties characterizing tightness. Here we recall some of them:

(II-A) For every response strategy $\phi : X \to Y$ there exists a strategy $y \in Y$ such that $g(y) \subseteq g(gr(\phi))$.

(II-B) For every response strategy $\psi : Y \to X$ there exists a strategy $x \in X$ such that $g(x) \subseteq g(gr(\phi))$.

It is not difficult to see that (I) and (II-A) are equivalent [**Gur73**, **Gur89**].

Then, by transposing g, we conclude that (I) and (II-B) are equivalent as well.

Hence, all three properties are equivalent. One can verify this for nine examples $g_1 - g_9$.

Properties (II-A) and (II-B) show that playing a zero-sum game $(g; u, w)$ with a tight game form g Bob and Alice do not need non-trivial response strategies but can restrict themselves by the standard ones, that is, by Y and X, respectively.

Given a game form $g : X \times Y \to \Omega$, introduce on the ground set Ω of the outcomes two multi-hypergraphs $\mathcal{A} = \mathcal{A}(g)$ and $\mathcal{B} = \mathcal{B}(g)$ whose edges are the supports of strategies of Alice and Bob, respectively:

$$\mathcal{A}(g) = \{g(x) \mid x \in X\} \text{ and } \mathcal{B}(g) = \{g(y) \mid y \in Y\}.$$

Recall that distinct edges of a multi-hypergraph may contain one another or even coincide. Obviously, the edges of $\mathcal{A}$ and $\mathcal{B}$ pairwise intersect, that is, $g(x) \cap g(y) \neq \emptyset$ for all $x \in X$ and $y \in Y$. Furthermore, g is tight if and only if

(III) hypergraphs $\mathcal{A}(g)$ and $\mathcal{B}(g)$ are *dual*, that is, satisfy also the following two (equivalent) properties:

(III-A) for every $\Omega_A \subseteq \Omega$ such that $\Omega_A \cap g(y) \neq \emptyset$ for all $y \in Y$ there exists an $x \in X$ such that $g(x) \subseteq \Omega_A$;

(III-B) for every $\Omega_B \subseteq \Omega$ such that $\Omega_B \cap g(x) \neq \emptyset$ for all $x \in X$ there exists an $y \in Y$ such that $g(y) \subseteq \Omega_B$.

Remark

*Verification of tightness of an explicitly given game form is an important open problem. No polynomial algorithm is known. A quasi-polynomial one was suggested in [**FK96**]; see also [**GK99**].*

Tightness and solvability

Let us recall an old theorem.

Theorem

*([**Gur75**, **Gur89**]) The following properties of a game form are equivalent:*

(i) Nash-, (ii) zero-sum- , (iii) ± 1-solvability, and (iv) tightness.

Proof

As we already mentioned, implications $(i) \Rightarrow (ii) \Rightarrow (iii)$ are obvious.

Also $(iii) \Rightarrow (iv)$ is easily seen.
Indeed, assume for contradiction that a game form g is not tight.
Then, there exists a response strategies $\phi : X \to Y$ and $\psi : Y \to X$ of Bob and Alice such that $g(gr(\phi)) \cap g(gr(\psi)) = \emptyset$.
Then, we can partition Ω into two sets of outcomes Ω_A and Ω_B (winning for Alice and Bob, respectively) in such a way that

$g(gr(\phi)) \subseteq \Omega_B$ and $g(gr(\psi)) \subseteq \Omega_A$.

(Note that for tight g this would not be possible.) Then,

$-1 = maxmin < minmax = 1$

in the obtained ± 1 game $(g; \Omega_A, \Omega_B)$ and, hence, it has no saddle point.

The inverse implication $(iii) \Leftarrow (iv)$,
(as well as $(ii) \Leftarrow (iv)$, which looks stronger)
are proven similarly; see [**EF70**, **Gur73**].

Assume that a zero-sum game $(g; u, w)$ has no saddle point.

Then, (1) fails and maxmin $<$ minmax.

Consider arbitrary best response strategies $\phi : X \to Y$ and $\psi : Y \to X$
of Bob and Alice, respectively.

Obviously, $g(gr(\phi)) \cap g(gr(\psi)) = \emptyset$ and, thus, g is not tight.

The last claim means that a tight game form is SP-solvable.

Moreover, it has a simple SP situation in minimal strategies [**Gur89**].

To finish the proof of the theorem it only remains to show implication
$(i) \Leftarrow (iv)$, that is, tightness implies Nash-solvability.

First, this was done in [**Gur75**], then, with more details, in [**Gur89**].

Several different proofs appeared later [**BBM90**, **BG03**, **DS91**, **Gur18**, **GK18**].
In the next section we suggest a new (and shortest) proof based on an
important general property of dual multi-hypergraphs.

Lexicographical theorem for dual multi-hypergraphs

Summary

Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ and $\mathcal{B} = \{B_1, \ldots, B_n\}$ an arbitrary pair of finite dual multi-hypergraphs on a common ground set Ω.
Each of them may have embedded or equal edges.

An edge is called *containment minimal* (or just *minimal*, for short) if it is not a strict superset of another edge. (minimal edges may still be equal.)

If $\mathcal{A}$ and $\mathcal{B}$ are dual then

(j) $A \cap B \neq \emptyset$ for every pair $A \in \mathcal{A}$ and $B \in \mathcal{B}$;

(jj) if A is minimal then for every $\omega \in A$ there exists a (minimal) $B \in \mathcal{B}$ such that $A \cap B = \{\omega\}$. We will extend claim (jj) as follows. A linear order $\succ$ over Ω uniquely defines a lexicographic order $\succ_\ell$ over the power set 2^Ω.

(jjj-A) Let A be a lexicographically maximal (lexmax) edge of $\mathcal{A}$. Then, edge A is minimal in $\mathcal{A}$ and for every $\omega \in A$ there exists a (minimal) edge $B \in \mathcal{B}$ such that $A \cap B = \{\omega\}$ and $\omega \succeq \omega'$ for each $\omega' \in B$.

By swapping $A, \mathcal{A}$ and $B, \mathcal{B}$, we obtain the dual statement (jjj-B).

These two statements form the *lexicographical theorem* for dual multi-hepergraphs. To formulate it accurately, we will need a few definitions.

Lexicographical orders over the subsets.

A linear order $\succ$ over a set Ω uniquely determines a lexicographical order $\succ_\ell$ over the power set 2^Ω (of all subsets of Ω) as follows. Roughly speaking, the more small elements are out of a set - the better it is. In particular, $\Omega' \succ_\ell \Omega''$ whenever $\Omega' \subset \Omega''$ and, hence, the empty set $\emptyset \subset \Omega$ is the best in 2^Ω.

Remark

Also $\{\omega'\} \succ_\ell \{\omega', \omega''\}$ *for any* $\omega', \omega'' \in \Omega$ *and order* $\succ$*, although set* $\{\omega', \omega''\}$ *gives a chance for a better outcome* ω'' *if* $\omega' \prec \omega''$*; see game form* g_6 *in Figure 1 and subsection 20 for more detail.*

More precisely, to compare two arbitrary subsets $\Omega', \Omega'' \subseteq \Omega$ consider their symmetric difference $\Delta = (\Omega' \setminus \Omega'') \cup (\Omega'' \setminus \Omega')$. Clearly, $\Delta \neq \emptyset$ if and only if sets Ω' and Ω'' are distinct. Let ω be the minimum with respect to $\succ$ element in Δ. If $\omega \in (\Omega' \setminus \Omega'')$ then $\Omega'' \succ_\ell \Omega'$; if $\omega \in (\Omega'' \setminus \Omega')$ then $\Omega' \succ_\ell \Omega''$.

We can reformulate this definition equivalently as follows. Without loss of generality (wlog), set $\Omega = \{\omega_1, \dots, \omega_p\}$ and assume that $\omega_1 \prec \dots \prec \omega_p$; assign the negative weight $w(\omega_i) = -2^{k-i}$ to every $\omega_i \in \Omega$, and set $w(S) = \sum_{\omega \in S} w(\omega)$ for each subset $S \subseteq \Omega$. Then, $\Omega' \succ_\ell \Omega''$ if and only if $w(\Omega') > w(\Omega'')$.

Denote by $supp(S)$ the $0,1$-vector $(s_1, \dots, s_k)$, where $s_i = 1$ if and only if $\omega_i \in S$. Then obviously, $\Omega' \succ_\ell \Omega''$ if and only if $supp(\Omega')$ is less than $supp(\Omega'')$ in the standard lexicographical order.

Dual multi-hypergraphs

Two finite multi-hypergraphs $\mathcal{A}$ and $\mathcal{B}$ on the common ground set Ω are called *dual* if (j) holds: $A \cap B \neq \emptyset$ for every pair $A \in \mathcal{A}$ and $B \in \mathcal{B}$, and also (jv-A) for each $B^T \subseteq \Omega$ such that $B^T \cap B \neq \emptyset$ for every $B \in \mathcal{B}$ there exists an $A \in \mathcal{A}$ such that $A \subseteq B^T$.

If (j) and (jv-A) both hold we say that $\mathcal{A}$ is dual to $\mathcal{B}$ and write $\mathcal{A} = \mathcal{B}^d$.

Swapping $A, \mathcal{A}$ and $B, \mathcal{B}$ in (jv-A) we obtain (jv-B) and an equivalent definition of duality, that is, (j) and (jv-A) hold if and only (j) and (jv-B) hold.

In other words, $\mathcal{A} = \mathcal{B}^d$ if and only if $\mathcal{B} = \mathcal{A}^d$.
So we just say that multi-hypergraphs $\mathcal{A}$ and $\mathcal{B}$ are dual.

Remark

*Dual multi-hypergraphs have numerous applications and appear in different areas under different names, such as “clutters" and “blockers" [**EF70**] or DNFs and CNFs of monotone Boolean functions [**CH11**].*

Lexicographical Theorem

Claims (j) and (jj) are well-known [**CH11**]. Actually, (j) is required by the definition of duality and (jj) is obvious. Indeed, if (jj) fails then edge A cannot be minimal, since its proper subset $A \setminus \{\omega\}$ would still intersect all $B \in \mathcal{B}$.

Our main result is statement (jjj-A).
Fix an arbitrary order $\succ$ over Ω and find a lexmax edge $A^L \in \mathcal{A}$, that is, one maximal with respect to the lexicographical order $\succ_\ell$ over 2^Ω. Note that such A^L may be not unique but all lexmax edges are equal. The lexicographic theorem is formulated as follows:

Theorem

A lexmax edge A^L is minimal in $\mathcal{A}$.
Furthermore, for every $\omega^* \in A^L$ there exists a (minimal) edge $B^M \in \mathcal{B}$ such that $A^L \cap B^M = \{\omega^*\}$ and $\omega^* \succ \omega$ for each $\omega \in B^M \setminus \{\omega^*\}$.

Proof A lexmax edge must be minimal, since a set is strictly less than any its proper subset in order $\succ_\ell$.

Assume for contradiction that there exists an $\omega^* \in A^L$ such that for every (minimal) $B \in \mathcal{B}$ satisfying (jj), $B \cap A^L = \{\omega^*\}$, there exists an $\omega \in B$ such that $\omega \succ \omega^*$. Clearly, this assumption holds for every $B^0 \in \mathcal{B}$ if it holds for each minimal $B^0 \in \mathcal{B}$. Let us show that it contradicts the lexmaximality of A^L. To do so partition all edges $B \in \mathcal{B}$ into two types:

(a) there is an $\omega \in B \cap A^L$ distinct from ω^*;

(b) $B \cap A^L = \{\omega^*\}$.

In case (b), by our assumption, there is an $\omega \in B$ such that $\omega \succ \omega^*$.

In both cases, (a) and (b), choose the specified ω from B,
thus, getting a transversal B^T.

By (jv-A), there exists an $A \in \mathcal{A}$ such that $A \subseteq B^T$ and, hence, $A \succeq_\ell B^T$.

Furthermore, by construction, $B^T \succ_\ell A^L$.

Indeed, $\omega^* \notin B^T$ and it is replaced in B^T by some larger elements, $\omega \succ \omega^*$, in case (b), while all other elements of B^T, if any, belong to $A^L \setminus \{\omega^*\}$, according to case (a).

Thus, by transitivity, $A \succ_\ell A^L$. Yet, by assumption of the theorem, A^L is a lexmax edge of $\mathcal{A}$, which is a contradiction. $\square$

Sperner hypergraphs

A multi-hypergraph is called *Sperner* if no two of its distinct edges contain one another; in particular, they cannot be equal. In this case, we have a hypergraph rather than multi-hypergraph. For a multi-hypergraph there exists a unique dual Sperner hypergraph. If $\mathcal{A}$ and $\mathcal{B}$ are dual and Sperner then

$\mathcal{A}^{dd} = \mathcal{A}$ and $\mathcal{B}^{dd} = \mathcal{B}$; furthermore $\cup_{A \in \mathcal{A}} A = \cup_{B \in \mathcal{B}} B = \Omega$.

In general, $\cup_{A \in \mathcal{A}} A$ and $\cup_{B \in \mathcal{B}} B$ may be different subsets of Ω.

Remark

*Here we assume that the reader is familiar with basic notions related to monotone Boolean functions, in particular, with DNFs and duality. An introduction can be found in [***CH11***].*

*It is well known [***CH11***] that (dual) multi-hypergraphs are in one-to-one correspondence with (dual) monotone DNFs: (prime) implicants of the latter correspond to (minimal) edges of the former. Furthermore, Sperner hypergraphs correspond to irredundant DNFs. However, we do not restrict ourselves to this case. Although the lexicographical theorem would not lose much but its applications to Nash-solvability would.*

Determining edges A^L and B^M in polynomial time.

Preliminaries

Edges A and B mentioned in (jjj-A) can be found in polynomial time.

The problem is trivial when multi-hypergraphs $\mathcal{A}$ and $\mathcal{B}$ are given explicitly.

We will solve it when only $\mathcal{A}$ is given, and not explicitly, but by a polynomial containment oracle.

For an arbitrary subset $\Omega_A \subseteq \Omega$ this oracle answers in polynomial time the question $Q(\mathcal{A}, \Omega_A)$: whether Ω_A contains an edge $A \in \mathcal{A}$.

By duality of $\mathcal{A}$ and $\mathcal{B}$, we have $A \not\subseteq \Omega_A$ for all $A \in \mathcal{A}$ if and only if $B \subseteq \Omega_B = \Omega \setminus \Omega_A$ for some $B \in \mathcal{B}$.

In other words, question $Q(\mathcal{A}, \Omega_A)$ is answered in the negative
if and only if $Q(\mathcal{B}, \Omega_B)$ is answered in the positive.
Thus, we do not need two separate oracles for $\mathcal{A}$ and $\mathcal{B}$;
it is sufficient to have one, say, for $\mathcal{A}$.

Determining a lexmax edge $\mathcal{A}^L$ in polynomial time.

Recall that multi-hypergraph $\mathcal{A}$ may contain several lexmax edges A^L, but obviously they are all equal.

Fix an arbitrary linear order $\succ$ over Ω.

Wlog we can assume that $\Omega = \{\omega_1, \dots, \omega_p\}$ and $\omega_1 \prec \dots \prec \omega_p$.

Step 1: Consider $\Omega_t^1 = \{\omega_t, \dots, \omega_p\}$ and, by asking question $Q(\mathcal{A}, \Omega_t^1)$ for $t = 1, \dots, p$, find the maximum t_1 for which the answer is still positive.

Then, ω_{t_1} belongs to $\mathcal{A}^L$, while $\omega_1, \dots, \omega_{t_1 - 1}$ do not.

Step 2: Consider $\Omega_t^2 = \{\omega_{t_1}, \omega_{t_1+t}, \dots, \omega_p\}$ and, by asking question $Q(\mathcal{A}, \Omega_t^2)$ for $t = 1, \dots, p - t_1$, find the maximum t_2 for which the answer is still positive.

Then, $\omega_{t_1}, \omega_{t_1+t_2} \in \mathcal{A}^L$, while $\omega_t \notin \mathcal{A}^L$ for any other $t < t_1 + t_2$.

Step 3: Consider $\Omega^3_t = \{\omega_{t_1}, \omega_{t_1+t_2}, \omega_{t_1+t_2+t}, \dots, \omega_p\}$ and, by asking question $Q(\mathcal{A}, \Omega^3_t)$ for $t = 1, \dots, p - (t_1 + t_2)$, find the maximum t_3 for which the answer is still positive.

Then, $\omega_{t_1}, \omega_{t_1+t_2}, \omega_{t_1+t_2+t_3} \in \mathcal{A}^L$, while $\omega_t \notin \mathcal{A}^L$ for any other $t < t_1 + t_2 + t_3$; etc.

We obtain a lexmax edge A^L in at most p polynomial iterations.

Note that on each step i we can speed up the search of t_i by applying dichotomy.

Determining an edge B^M

First, find a lexmax edge $A^L \in \mathcal{A}$ and choose an arbitrary $\omega^* \in A^L$.
We look for an edge $B^M \in \mathcal{B}$ such that $A^L \cap B^M = \{\omega^*\}$ and $\omega^* \succ \omega$ for every $\omega \in B^M \setminus \{\omega^*\}$. In other "words",

$$B^M \subseteq \Omega_B = \Omega \setminus [(A^L \setminus \{\omega^*\}) \cup \{\omega \mid \omega \succ \omega^*\}].$$

By Theorem 2, such B^M exists and, hence, the oracle answers $Q(\mathcal{B}, \Omega_B)$ in the positive, or equivalently, $Q(\mathcal{A}, \Omega \setminus \Omega_B)$ in the negative. We could take any $B^M \in \mathcal{B}$ such that $B^M \subseteq \Omega_B$. Yet, multi-hypergraph $\mathcal{B}$ is not given explicitly.

To get B^M we need "to minimize" Ω_B. To do so, let us delete its elements one by one in some order until we obtain a minimum set Ω_B^* for which the answer to $Q(\mathcal{A}, \Omega \setminus \Omega_B^*)$ is still negative, that is, answers to $Q(\mathcal{A}, \Omega \setminus (\Omega_B^* \setminus \{\omega\})$ become positive for every $\omega \in \Omega_B^*$. Then, we set $B^M = \Omega_B^*$.
Again we can speed up the procedure by applying dichotomy.

Note that the above reduction procedure may be not unique, since we can eliminate elements of $\Omega \setminus \Omega_B$ in an arbitrary order. Thus, in contrast to A^L, there may be several not equal edges B^M satisfying all conditions of Theorem 2.

Lexicographically safe NE in games with tight game forms

Summary

First, we apply Theorem 2 to finish the proof of Theorem 1. It remains to show that $(i) \Leftarrow (iv)$, that is, tightness implies Nash-solvability. In other words, a game $(g; u, w)$ has a NE for any payoffs u and w whenever game form g is tight. The proof is constructive: we will obtain two special types of NE.

Given g and u, choose a lexmax strategy $x \in X$ of Alice. By Theorem 2, there is a strategy $y \in Y$ of Bob such that (x, y) is a NE. By definition, y must be a best response to x such that x is also a best response to y. By Theorem 2, the obtain situation (x, y) is simple and both strategies, x and y are minimal. More precisely, x must be minimal, while y can be chosen minimal. These NE will be called *lexsafe NE of Alice* and the set of these NE will denoted by NE-A. Similarly, we define a set NE-B of Bob's *lexsafe NE*.

Remark

We assume that both players are maximizers and adjective "lexsafe" can be replaced by "lexmax". If both players are minimizers then it can be replaced by "lexmin". In the zero-sum case Alice is the maximizer, while Bob is the minimizer. Flexible term lexsafe may replace both, lexmax or lexmin.

Previous results provide a polynomial algorithm determining at least one NE from NE-A and at least one from NE-B (which may coincide) in a given game $(g; u, w)$ with a tight game form g. This is trivial when g is given explicitly. Yet, the algorithm works when one of two multi-hypergraphs $\mathcal{A}(g)$ or $\mathcal{B}(g)$ is given by a polynomial containment oracle.

Lexicographically safe strategies of players

Given g and preference $\succ_A$ of Alice, let us introduce the lexicographical pre-order over Alice's strategies $x \in X$ as follows. Consider lexicographical order $\succ_A^\ell$ over 2^Ω defined by the linear order $\succ_A$ over Ω. The larger is the support $g(x) \subseteq \Omega$ in order $\succ_A^\ell$, the safer is strategy x for Alice, while strategies with the same support are equally safe. Alice's strategies that maximize support $g(x)$ in order $\succ_A^\ell$ will be called her lexmax (or lexsafe) strategies.

In particular, all lexmax strategies have the same support.

Furthermore, a lexsafe strategy is minimal. Indeed, x is safer than x' whenever $g(x) \subset g(x')$ and containment is strict.

Note also that Alice's lexsafe strategies are defined by g and $\succ_A$, while Bob's preference $\succ_B$ is irrelevant. Alice may be even unaware of it, which is important for applications.

Similarly, using $y \in Y$ and $\succ_B$ instead of x and $\succ_A$, we define Bob's lexsafe strategies. Respectively, they depend only on g and $\succ_B$, while $\succ_A$ is irrelevant.

The concept of a lexsafe strategy can be viewed as a refinement of the classical concept of a safe (maxmin) strategy. The latter optimizes the worst case scenario outcome, while lexsafe strategies optimize the whole set of outcomes in the lexicographical order defined above.

Thus, lexsafe strategies are safest, but sometimes may be not rational. For example, let $g(x) = \{\omega\}$, $g(x') = \{\omega, \omega'\}$ and $\omega \prec_A \omega'$. Then $x \succ_A x'$, although strategy x' is better for Alice than x. Indeed, x' gives her a chance to obtain the better outcome ω', while x excludes ω' and ensures ω; see Remark 5.

Consider a preference over Ω such that outcomes $\omega, \omega' \in \Omega$ are, respectively, the worst and the best outcomes for both Alice and Bob simultaneously. Consider a game form $g : X \times Y \to \Omega$ having two strategies $x^* \in X$ and $y^* \in Y$ such that $g(x, y) = \omega$ if and only if $x = x^*$ or $y = y^*$ and $g(x, y) = \omega'$ otherwise. Then, x^* and y^* are the only lexsafe strategies of Alice and Bob; furthermore, situation (x^*, y^*) is a unique lexsafe NE, but outcome $g(x^*, y^*) = \omega$ is worse than ω' for both players. One can view this as a price of stability. For comparison, recall that NE may be not Pareto optimal.

Lexsafe Nash equilibria in games with tight game forms

Recall that game form $g : X \times Y \to \Omega$ is tight if and only if its hypergraphs $\mathcal{A} = \mathcal{A}(g) = \{g(x) \mid x \in X\}$ and $\mathcal{B} = \mathcal{B}(g) = \{g(y) \mid y \in Y\}$ are dual.

Given a game $(g; \succ_A, \succ_B)$ with a tight game form g, choose any lexsafe strategy x^L of Alice. By Theorem 2 it is minimal. Let us show that there exists a strategy y^M of Bob such that (x^L, y^M) is a NE. By definition, y^M is a best response to x^L, that is, $g(x^L, y^M) \succeq_B g(x^L, y)$ for any $y \in Y$. (Note, however, that the preference is not strict, because for some y two outcomes may coincide: $g(x^L, y^M) = g(x^L, y)$.) Let us apply Theorem 2 setting

$$g(x^L) = A^L,\ g(y^M) = B^M, g(x^L, y^M) = \omega^*,$$

and conclude that there exists a (minimal) strategy y^M such that x^L, in its turn, is a best response to y^M. Thus, (x^L, y^M) is a NE. Theorem 1 is proven.

Moreover, we can strengthen it summarizing remarkable properties of the obtained NE. Recall that in Theorem 2 both edges A^L and B^M are minimal and $A^L \cap B^M = \{\omega^*\}$. Hence, for the obtained NE (x^L, y^M) both strategies x^L and y^M are minimal and situation (x^L, y^M) is simple, that is, $g(x^L) \cap g(y^M) = \{\omega^*\}$; see **[Gur89]**. More precisely, X^L must be minimal, while Y^L can be chosen minimal.

Denote by X^L the set of all lexmax strategies of Alice. By definition, they all have the same support. Let us fix $x^L \in X^L$ and denote by $Y^M(x^L)$ the set of all Bob's best responses to x^L. In fact, $Y^M(x^L)$ does not depend on x^L provided $x^L \in X^L$. Indeed, set $g(x^L)$ is unique, that is, the same for all $x^L \in X_L$ and $g(x^L) \cap g(y^M) = \omega^*$ for all $y^M \in Y^M(x'^L)$ and for all $x'^L \in X^L$. Hence, $y^M(x^L)$ is a best response of Bob to each Alice's lexmax strategy. Denote by y^M the set of all such best responses.

Thus, we obtain $X^L \subseteq X$ and $Y^M \subseteq Y$ such that for any pair $x^L \in X^L$ and $y^M \in Y^M$ situation (x^L, y^M) is simple, $g(x^L, y^M) = \{\omega^*\}$, and (x^L, y^M) is a NE, because X^L is a best response to Y^M and vice versa.

In other words, the direct product NE-A $= (x^L \times y^M) \subseteq X \times Y$ consists of simple NE situations corresponding to the same outcome $\omega^* \in \Omega$. Furthermore, all strategy of X^L are lexsafe and, hence, minimal, while Y^M contains minimal strategies. We will call NE-A the *box of Alice's lexsafe equilibria*.

By construction, X^L depends only on Alice's preference $\succ_A$, while Bob's preference $\succ_B$ is irrelevant and Alice may be just unaware of it, which is important for applications. In contrast, Y^M is a set of some (special) Bob's best responses to X^L, which are the same for all $x^L \in X^L$.

Swapping the players, we obtain the *box of Bob's lexsafe equilibria* NE-B $= (x^M \times y^L) \subseteq X \times Y$ with similar properties. Thus, we can strengthen Theorem 1 as follows:

Theorem

Every game $(g; \succ_A, \succ_B)$ with a tight game form g has two non-empty boxes of lexmax equilibria NE-A $= X^L \times Y^M$ and NE-B $= X^M \times Y^L$ of Alice and Bob satisfying the above properties. □

Boxes NE-A and NE-B may intersect or even coincide. For example, this always happens in the sero-sum case. Then X^L and X^M are maxmin strategies of Alice, while Y^M and Y^L are minmax strategies of Bob. More detail can be found in the first arXiv version of this paper [**GN21**]. NE-A and NE-B may be equal in the non-zero-sum case too. For example a game may have a unique NE.

A pair of lexsafe strategies of Alice and Bob may be not a NE

For example, consider tight game form g_1 in Figure 1.
Define preferences $\succ_A$ and $\succ_B$ such that $\omega_2 \succ_A \omega_1 \succ_A \omega_3$ and $\omega_2 \succ_B \omega_3$.
It is easily seen that x_1 and y_1 are lexsafe strategies of Alice and Bob, respectively. Yet, situation (x_1, y_1) is not an NE. Alice can improve her result $g_1(x_1, y_1) = \omega_1$ by switching to x_2 and getting $g(x_2, y_1) = \omega_2$.
Thus, two lexsafe strategies, of Alice and Bob, do not form an NE.
However, sets NE-A and NE-B are not empty, in accordance with Theorem 3: NE-A $= \{(x_2, y_1)\}$ and NE-B $= \{(x_1, y_2)\}$.
The corresponding NE outcomes are ω_1 and ω_2, respectively.

Note that ω_2 is the best outcome for both players if $\omega_2) \succ_B \omega_1$.
In this case NE-B is not Pareto-optimal.

Remark

*One could conjecture that each player prefers lexsafe NE of the opponent to his/her own. Such result would be similar to the analogous one from the matching theory; see, for example, [***GI89***].*
*There are two types of stable matchings given by the Gale-Shapley algorithm [***GS62***], depending on men propose to women or vice versa.*
Yet, this conjecture is disproved by the above example.

Computing lexsafe NE in polynomial time

If game form $g : X \times Y \to \Omega$ is given explicitly then to find all its NE is simple: one can just consider all situations $(x, y) \in X \times Y$ one by one verifying Nash's definition for each of them. Yet, in applications g is frequently given by an oracle $\mathcal{O}$ such that size of g is exponential in size $|\mathcal{O}|$ of this oracle.
Then, the straightforward search for NE suggested above becomes not efficient.

Four such oracles will be considered in the next section.

The following three properties of oracle $\mathcal{O}$ will allow us to construct an algorithm computing two lexsafe NE (from NE-A and NE-B, respectively) for a given game $(g; \succ_A, \succ_B)$ with tight game form $g = g(\mathcal{O})$ realized by $\mathcal{O}$, in time polynomial in $|\mathcal{O}|$.

(I) Oracle $\mathcal{O}$ contains explicitly all outcomes Ω of g.
(In contrast, the strategies $x \in X$ and $y \in Y$ may be implicit in $\mathcal{O}$; moreover, $|X|$ and $|Y|$ may be exponential in $|\mathcal{O}|$.)

(II) The game form $g = g(\mathcal{O})$ defined by $\mathcal{O}$ is tight.

(III) Every ± 1 game $(g(\mathcal{O}); \Omega_A, \Omega_B)$ can be solved in time polynomial in $|\mathcal{O}|$.

Requirement (III) needs a discussion.

By tightness of g, exactly one of the following two options holds:

(a) there exists $x \in X$ with $g(x) \subseteq \Omega_A$ (Alice wins);
(b) there exists $y \in Y$ with $g(y) \subseteq \Omega_B$ (Bob wins).

Note that (a) (respectively, (b)) holds if and only if the monotone Boolean function corresponding to multi-hypergraph $\mathcal{A}(g)$ (respectively, $\mathcal{B}(g)$) takes value 1; see Remark 6.

To solve a ± 1 game we determine which option, (a) or (b), holds and output a winning strategy, x or y, respectively.

Note that it is possible to output a *minimal* winning strategy whenever (III) holds. Indeed, suppose Alice wins and we output her winning strategy x, with $g(x) \subseteq \Omega_A$. Reduce Ω_A by one outcome ω by moving it to Ω_B, solve the obtained ± 1 game, and repeat the procedure for all $\omega \in \Omega_A$. If Bob wins in all obtained games then x is already minimal. Otherwise we can move an outcome ω from Ω_A to Ω_B and Alice still wins. Repeating, we obtain a minimal winning strategy of Alice (in the original game) in at most $|\Omega_A|$ steps. We can speed up the above procedure using dichotomy. The same works for Bob.

Theorems 1–3 immediately imply the following statement.

Theorem

Given an oracle $\mathcal{O}$ satisfying requirements (I,II,III), a lexsafe NE of Alice (and of Bob, as well) exists and can be computed in time polynomial in $|\mathcal{O}|$. □