

Lecture 3

Properties and relaxations of the Nash equilibrium

Instructor: Prof. Gabriele Farina

In this lecture, we will continue analyzing the properties of Nash equilibria in normal-form games. We will then introduce the concept of correlated equilibrium, a relaxation of Nash equilibrium with desirable properties.

L3.1 Further properties of the Nash equilibrium

We ended Lecture 1 with the definition of a Nash equilibrium. Recall that a strategy profile is a Nash equilibrium if no player can unilaterally deviate from their strategy to improve their payoff. We also discussed the existence of Nash equilibria in finite games, which is guaranteed by the Brouwer fixed-point theorem.

L3.1.1 Nash equilibrium in two-player zero-sum games

As mentioned in the previous lecture, in two-player zero-sum games the Nash equilibria are exactly those strategy profiles for which both players are playing a maximin strategy. We formalize this in the next theorem. First, though, we introduce some notation which will make our life easier when dealing with two-player games.

Definition L3.1 (Matrices U_1 and U_2 for two-player games). Consider a generic two-player zero-sum game, as shown next. As usual, we denote the sets of actions for player by A_1 and A_2 .

	action 1	action 2	...	action m
action 1	a_{11}, b_{11}	a_{12}, b_{12}	...	a_{1m}, b_{1m}
action 2	a_{21}, b_{21}	a_{22}, b_{22}	...	a_{2m}, b_{2m}
$\vdots$	$\vdots, \vdots$	$\vdots, \vdots$	$\ddots, \ddots$	$\vdots, \vdots$
action n	a_{n1}, b_{n1}	a_{n2}, b_{n2}	...	a_{nm}, b_{nm}

Let $x \in \Delta(A_1)$ denote a strategy of Player 1, and $y \in \Delta(A_2)$ a strategy of Player 2. We can express the expected utilities for the players according to the bilinear expressions

$$u_1(x, y) = x^\top U_1 y, \quad u_2(x, y) = x^\top U_2 y,$$

where

$$U_1 := \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{pmatrix}, \quad U_2 := \begin{pmatrix} b_{11} & b_{12} & \cdots & b_{1m} \\ b_{21} & b_{22} & \cdots & b_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ b_{n1} & b_{n2} & \cdots & b_{nm} \end{pmatrix}.$$

From now on, we will assume that a two-player game has been defined, and we will use the notation with U_1 and U_2 defined above to refer to the utility matrices of the players.

Theorem L3.2. Consider a two-player zero-sum game, that is, one for which $U_2 = -U_1$. Then, a strategy profile $(x^*, y^*) \in \Delta(A_1) \times \Delta(A_2)$ is a Nash equilibrium if and only if it is a maxmin strategy, *i.e.*, if and only if

$$x^* \in \arg \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y, \quad \text{and} \quad y^* \in \arg \max_{y \in \Delta(A_2)} \min_{x \in \Delta(A_1)} x^\top U_2 y.$$

Proof. We prove the result assuming we trust von Neumann's minimax theorem, which states that

$$\max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y = \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y.$$

($\Rightarrow$) Suppose that (x^*, y^*) is a Nash equilibrium. Then, by the definition of Nash equilibrium and using the fact that $U_2 = -U_1$, we have that

$$(x^*)^\top U_1 y^* = \max_{x \in \Delta(A_1)} x^\top U_1 y^*, \quad \text{and} \quad (x^*)^\top U_1 y^* = \min_{y \in \Delta(A_2)} (x^*)^\top U_1 y.$$

Hence, we can write the chain of equalities and inequalities

$$\min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y \leq \max_{x \in \Delta(A_1)} x^\top U_1 y^* = \min_{y \in \Delta(A_2)} (x^*)^\top U_1 y \leq \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y.$$

By the minimax theorem, all inequalities must be equalities; hence, (x^*, y^*) satisfies

$$\begin{aligned} \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y &= \max_{x \in \Delta(A_1)} x^\top U_1 y^* \Leftrightarrow y^* \in \arg \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y \\ \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y &= \min_{y \in \Delta(A_2)} (x^*)^\top U_1 y \Leftrightarrow x^* \in \arg \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y. \end{aligned}$$

($\Leftarrow$) Conversely, suppose that x^* and y^* are maxmin strategies. Let v^* be the common value of both sides of the minimax theorem, that is,

$$v^* := \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y = \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y.$$

We now show that (x^*, y^*) is a Nash equilibrium. By definition, this means we need to show that

$$(x^*)^\top U_1 y^* = \max_{x \in \Delta(A_1)} x^\top U_1 y^* \quad \text{and} \quad (x^*)^\top U_1 y^* = \min_{y \in \Delta(A_2)} (x^*)^\top U_1 y.$$

Using the hypothesis,

$$\begin{aligned} x^* \in \arg \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y, & \implies v^* = \min_{y \in \Delta(A_2)} (x^*)^\top U_1 y, \\ y^* \in \arg \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y, & \implies v^* = \max_{x \in \Delta(A_1)} x^\top U_1 y^*. \end{aligned}$$

These equalities imply that $v^* \leq (x^*)^\top U_1 y^*$ and $v^* \geq (x^*)^\top U_1 y^*$, and thus $v^* = (x^*)^\top U_1 y^*$. This shows that the players are best responding to the strategy of the opponent, completing the proof that (x^*, y^*) is a Nash equilibrium. $\square$

Computation As we will see shortly, Theorem L3.2 gives us nontrivial information about the structure of Nash equilibria in two-player zero-sum games. But it also gives us a computational tool. Indeed, the theorem above tells us that finding a Nash equilibrium in a two-player zero-sum game can be expressed as an optimization problem. Let's show that this optimization problem is a linear program. Without loss of generality, let's focus on Player 1's optimization problem, that is,

$$x^* \in \arg \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y.$$

The key insight is that this problem can be rewritten as

$$\begin{cases} \max_v v \\ \text{s.t. } v \leq x^\top U_1 a_2 \quad \forall a_2 \in A_2 \\ 1^\top x = 1 \\ x \geq 0. \end{cases}$$

which is a linear program with a linear number of constraints in the number of actions of Player 2. We can use any linear programming solver to find such a solution. We will see more scalable methods to compute maxmin strategies from repeat play starting next week.

Connection with linear programming It is worth pausing for a moment to appreciate some historical context. We started the proof by assuming von Neumann's minimax theorem, which we justified as a consequence of linear programming duality. However, historically, von Neumann did not have the luxury of linear programming to prove his theorem.

- The proof of von Neumann's minimax theorem essentially hides an optimization duality argument. Indeed, we have the following:

$$\begin{array}{ccc} \max_{x \in \Delta(A_1)} \min_{y \in \Delta(A_2)} x^\top U_1 y & & \min_{y \in \Delta(A_2)} \max_{x \in \Delta(A_1)} x^\top U_1 y \\ \updownarrow & & \updownarrow \\ \begin{cases} \max v \\ v \leq x^\top U_1 a_2 \quad \forall a_2 \\ 1^\top x = 1 \\ x \geq 0. \end{cases} & \begin{array}{c} \text{linear programming} \\ \longleftrightarrow \\ \text{duality} \end{array} & \begin{cases} \min w \\ w \geq a_1^\top U_1 y \quad \forall a_1 \\ 1^\top y = 1 \\ y \geq 0. \end{cases} \end{array}$$

- The connection between linear programming and two-player zero-sum games is bidirectional: as it turns out, *solving linear programming* and *finding a Nash equilibrium in a two-player zero-sum game* are *computationally equivalent*. This means that *any* linear programming problem (with arbitrary constraints, variables, etc.) can be efficiently converted into a two-player zero-sum game. This is less obvious than it may seem. For one, the strategy sets in games are probability simplexes, while linear programs might have arbitrary linear equality and inequality constraints. Furthermore, linear programs might be unbounded or unfeasible; yet, a Nash equilibrium of a game always exists. It would have been perfectly reasonable to believe that linear programming was a significantly more general tool than equilibrium solvers for two-player zero-sum games, and we know today that that belief would have been wrong. For more on this, see [Adl13; BR23; Ste23].
- All of this seems simple with the luxury of hindsight. But the two fields were not as closely connected as we might think. We know this from a transcript of the first encounter between Dantzig, one of the fathers of linear programming, and von Neumann, one of the fathers of game theory. And, perhaps in what is a plot twist, it was von Neumann to teach Dantzig about duality!

« On October 3, 1947, I visited him (von Neumann) for the first time at the Institute for Advanced Study at Princeton. I remember trying to describe to von Neumann, as I would to an ordinary mortal, the Air Force problem. I began with the formulation of the linear programming model in terms of activities and items, etc. Von Neumann did something which I believe was uncharacteristic of him. “Get to the point,” he said impatiently. Having at times a somewhat low kindlingpoint, I said to myself “O.K., if he wants a quicky, then that’s what he will get.” In under one minute I slapped the geometric and algebraic version of the problem on the blackboard. Von Neumann stood up and said “Oh that!” Then for the next hour and a half, he proceeded to give me a lecture on the mathematical theory of linear programs. At one point seeing me sitting there with my eyes popping and my mouth open (after I had searched the literature and found nothing), von Neumann said: “I don’t want you to think I am pulling all this out of my sleeve at the spur of the moment like a magician. I have just recently completed a book with Oskar Morgenstern on the theory of games. What I am doing is conjecturing that the two problems are equivalent. The theory that I am outlining for your problem is an analogue to the one we have developed for games.” Thus I learned about Farkas’ Lemma, and about duality for the first time.»

(from Dantzig, G. B. [Dan82])

- In light of the above you might be wondering how easy it would be to prove the minimax theorem without relying on linear programming duality. As we will show, the mere existence of learning dynamics in games is enough. We will also see a different proof next time.

Topological properties It is important to realize that what Theorem L3.2 is saying is that in two-player zero-sum games, each player can plan their own strategy *independently*. *Any* combination of maxmin strategies for the players forms an equilibrium. This is in contrast

with the general case: there, in order to specify a Nash equilibrium we need to provide a tuple of strategies for all players. In two-player zero-sum games, instead, *any product of maximin strategies is an equilibrium*. We have just arrived at the following corollary.

Corollary L3.3. In a two-player zero-sum game, the set of Nash equilibria is a Cartesian product of nonempty, convex, compact sets.

Since Cartesian products of nonempty, convex, and compact sets are themselves nonempty, convex, and compact, Corollary L3.3 immediately implies the following as well.

Corollary L3.4. The set of Nash equilibria in a two-player zero-sum game is nonempty, convex, and compact.

It is worth remarking again that what does the heavy lifting here is really Theorem L3.2 the rest follows as a direct corollary.

■ L3.1.2 Nash equilibrium in two-player general-sum games

In the general two-player case, often referred to as *two-player general-sum games*, many of the nice properties of the zero-sum case are lost.


Topological properties Perhaps the most striking is that not only the set of Nash equilibria is no longer guaranteed to be convex, but it is not even guaranteed to be contractible. We show this phenomenon with the next example.

Remark L3.5 (Complex topology of Nash equilibria; Kohlberg-Mertens game [\[KM86\]](#)). Beyond two-player zero-sum games, the set of Nash equilibria in a game can be quite complex.

For one, *it is not at all guaranteed that the set is convex*. Even more, the set might be *topologically complex*, e.g., exhibiting holes. This phenomenon was already observed by [Kohlberg, E., & Mertens, J.-F. \[KM86\]](#), who considered the two-player three-action game

	A	B	C
A	1, 1	0, -1	-1, 1
B	-1, 0	0, 0	-1, 0
C	1, -1	0, -1	-2, -2

The figure on the right, similar to the one in [\[Mil+23\]](#), shows the projection of the set of all 0.27-approximate Nash equilibria of this game, i.e., all strategy profiles such that no player has a unilateral deviation that increases their utility by more than 0.27.



Computation In two-player general-sum games, computation of Nash equilibria is not a linear program. However, it is a *linear complementarity problem* (LCP), a more general class of problems than linear feasibility programs, and which are written in the form

$$\text{find } x, w \in \mathbb{R}^d \quad \text{s.t.} \quad w = Mx + q, \quad x, w \geq 0, \quad x^\top w = 0.$$

The Lemke-Howson algorithm is a well-known algorithm to solve LCPs, and it can be used to find Nash equilibria in two-player general-sum games. However, the algorithm is not polynomial-time in the worst case, and it can be hard to find Nash equilibria in practice. An important corollary of the connection between two-player general-sum games and LCPs is the following:

Corollary L3.6. Any two-player general-sum games with rational payoffs admits a Nash equilibrium with rational coordinates.

This follows directly from the way Lemke-Howson works, which is similar to the simplex algorithm. The algorithm moves along edges of a rational polytope until it finds a Nash equilibrium. Since the algorithm only moves along the edges of the polytope, it will only generate rational solutions.

An interesting result about the computation of ϵ -approximate Nash equilibria is due to [Lipton, R. J., Markakis, E., & Mehta, A. \[LMM03\]](#), and is based on the observation that every game admits an ϵ -approximate Nash equilibrium where the strategy of Player 1 is supported on at most $w := O\left(\frac{\log |A_2|}{\epsilon^2}\right)$ strategies. This follows from using a Hoeffding bound on samples from the distribution of Player 1's strategy. One can then check any support for Player 1's strategy of size up to w , and for each such support, solve a linear program to verify if a Nash equilibrium with that support exists. This gives a subexponential-time

algorithm (of order $O(s^{\log s/\epsilon^2})$, where s is the size of input) for computing an ϵ -approximate Nash equilibrium.

L3.1.3 Nash equilibrium in games with more than two players

In games with more than two players, the behavior of Nash equilibria can be even more problematic.

Analytic properties As a start, *rational numbers might not be enough anymore* to store the probabilities of each player's actions at equilibrium.

Example L3.7. In his original paper, Nash showed a three-player game with rational payoffs and with the property that *all* Nash equilibria prescribe probabilities that are irrational numbers [Nas51]. Another simple example is also reported by Nau, R., Canovas, S. G., & Hansen, P. [NCH04], as follows:

	<i>left</i>	<i>right</i>		<i>left</i>	<i>right</i>
<i>top</i>	3, 0, 2	0, 2, 0	<i>top</i>	1, 0, 0	0, 1, 0
<i>bottom</i>	0, 1, 0	1, 0, 0	<i>bottom</i>	0, 3, 0	2, 0, 3
	<i>(action X)</i>			<i>(action Y)</i>	

A simple calculation shows that the only Nash equilibrium (x^*, y^*, z^*) of this game satisfies

$$(x_1^*, y_1^*, z_1^*) = \left(\frac{53}{46} - \frac{\sqrt{601}}{46}, -\frac{13}{24} + \frac{\sqrt{601}}{24}, -\frac{23}{4} + \frac{\sqrt{601}}{4} \right) \approx (0.619, 0.480, 0.379).$$

From a computational point of view, this property raises the question of how a Nash equilibrium solver could even *represent* such an output.

Proof. Homework. □

Remark L3.8. The issues with irrational numbers do not stop at square roots. In fact, *any polynomial root* might be required to represent a Nash equilibrium. This was shown by Bubelis, V. [Bub79], who showed how to construct games with arbitrary polynomial roots.

Beyond the representation, the topology of Nash equilibria is also in general arbitrarily complex in three-player games. In particular, Datta, R. S. [Dat03] showed that for any real algebraic variety, one can come up with some three-player game whose set of Nash equilibria is isomorphic to that variety.

Computation On the computational side, the situation is even more dire. As a first consideration, because Nash equilibria might require irrational numbers, even the question of how to *represent* the output equilibrium needs attention. In general, we cannot hope for an *exact* value. However, even asking for a *constant* approximation turns out to be hard. We will talk about this in more detail at the end of the course, where we relate the computation of (approximate) Nash equilibria to a complexity class called PPAD.

If one is willing to stomach a worst-case superpolynomial runtime, some methods exist. While the Lemke-Howson algorithm cannot be used beyond two-player games, other methods (such as [PNS08]) still apply.

L3.2 Correlated and coarse correlated equilibrium

The discussion above shows that Nash equilibria can be hard to compute and might not form a convex (or even contractible) set. This motivates the study of *correlated equilibria* [Aum74] and *coarse correlated equilibria* [MV78], which are a relaxation of Nash equilibria that are easier to compute, always form a convex set, and for which rational solutions always exist. As we will show starting in a few lectures, another major advantage of correlated equilibria is that they can be learned from repeated play, in a way that is fundamentally incompatible with Nash equilibria.¹

L3.2.1 Coarse correlated equilibrium

Remember that in a Nash equilibrium we are seeking a strategy profile $(x_1, \dots, x_n) \in \Delta(A_1) \times \dots \times \Delta(A_n)$ such that no player can unilaterally deviate to improve their payoff, that is,

$$u_i(a'_i, x_{-i}) \geq u_i(x_i, x_{-i}) \quad \forall i \in [n], a'_i \in A_i.$$

Here, u_i was defined as the expected payoff when all the players randomize *independently*.

The concept of *coarse correlated equilibrium* is a relaxation of this definition. In a coarse correlated equilibrium, instead of asking for the players to pick *independent* strategies, we allow coordination. In particular, we define the following.

Definition L3.9 (Coarse correlated equilibrium [MV78]). A *coarse correlated equilibrium* (CCE) is a correlated strategy $\mu \in \Delta(A_1 \times \dots \times A_n)$ such that

$$\mathbb{E}_{(a_1, \dots, a_n) \sim \mu} [u_i(a'_i, \dots, a_n)] \leq \mathbb{E}_{(a_1, \dots, a_n) \sim \mu} [u_i(a_1, \dots, a_n)] \quad \forall i \in [n], a'_i \in A_i. \quad (1)$$

Remark L3.10. The definition of a CCE is a relaxation of the definition of a Nash equilibrium. In a Nash equilibrium, the players randomize independently; in a CCE, they can randomize in a correlated way. A *Nash equilibrium is a CCE μ that happens to be a product distribution*, that is, $\mu = x_1 \otimes \dots \otimes x_n$.

This shows that the set of CCEs is a superset of the set of Nash equilibria. Thus, a coarse correlated equilibria always exists in every game.

¹A paradigm that has been successful in applications is to learn a correlated equilibrium from repeated play, and then marginalize it into a profile that is hoped to be close to a Nash equilibrium. This was used for example to reach superhuman performance in multiplayer poker [BS19].

Properties and computation We can turn Definition L3.9 into an optimization problem. The variables are the entries of the probability distribution μ . This is a $(A_1 \times \dots \times A_n)$ -dimensional nonnegative vector whose entries must satisfy the linear equality constraint

$$\sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} \mu_{a_1, \dots, a_n} = 1.$$

Furthermore, expanding the expectation in inequality 1 defines a set of linear constraints

$$\sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} \mu_{a_1, \dots, a_n} u_i(a'_i, a_{-i}) \leq \sum_{a_1 \in A_1} \dots \sum_{a_n \in A_n} \mu_{a_1, \dots, a_n} u_i(a_i, a_{-i})$$

for all $i \in [n]$ and $a'_i \in A_i$. Hence, the set of CCEs is the intersection of a finite set of linear constraints, and so it is a convex polytope. Note that the number of constraints is polynomial in the game (*i.e.*, in the size of the payoff table), and so we can use linear programming to compute and even optimize over the set of CCEs in time polynomial in $|A_1| \times \dots \times |A_n|$.

Corollary L3.11. Since the coefficients of the linear constraints are the payoffs of the game, the set of CCEs is always a rational polytope.

It is worth knowing that a CCE can also be computed in polynomial time in imperfect-information sequential games, despite the number of “actions” there, which is the number of strategies in the tree, is exponential in the input. Unfortunately, we lose the ability to optimize over the set.

L3.2.2 Correlated equilibrium

The concept of *correlated equilibrium* is an intermediate relaxation between Nash equilibrium and coarse correlated equilibrium.

Definition L3.12 (Correlated equilibrium [Aum74]). A *correlated equilibrium (CE)* is a correlated strategy $\mu \in \Delta(A_1 \times \dots \times A_n)$ such that

$$\mathbb{E}_{(a_1, \dots, a_n) \sim \mu} [u_i(\phi_i(a_i), a_{-i})] \leq \mathbb{E}_{(a_1, \dots, a_n) \sim \mu} [u_i(a_i, a_{-i})] \quad \forall i \in [n], \phi_i : A_i \rightarrow A_i,$$

where the function $\phi_i : A_i \rightarrow A_i$ is arbitrary.

Remark L3.13. A CCE is a special case of a CE, where the functions ϕ_i considered are only *constant* functions. Furthermore, it is not hard to show from expanding the definition that any Nash equilibrium is a CE. Thus, the set of CEs is a superset of the set of Nash equilibria and a subset of the set of CCEs.

All remarks made about the computation of CCEs in normal-form games apply to CEs as well. In particular, the set of CEs is a convex polytope, and a CE can be computed in polynomial time using linear programming.

However, the remark about computation in imperfect-information sequential games does not apply to CEs. Whether a CE can be computed efficiently in such games is an open question in the field. Some mild evidence suggests that the problem might be hard. Intuitively, the issue is that the number of functions ϕ in those games might be too large to control.

■ L3.2.3 How to think about correlated play in games

We can think of the correlation between the strategies of the players in a correlated or coarse correlated equilibrium as arising from some *correlation device* in the game. This is a trusted mediator that can recommend but not enforce behavior. The distribution μ from which the correlation device samples recommendations is public knowledge, but the players only get to observe the recommended action that was sampled for them. A correlated / coarse correlated equilibrium is then a distribution μ such that no player can unilaterally deviate from the recommended action to improve their payoff.

The distinction between correlated and coarse correlated equilibrium is in when the players decide when to commit to the recommended action. In a coarse correlated equilibrium, the players commit to the recommended action *before* the recommendation is made. In a correlated equilibrium, the players commit to the recommended action *after* the recommendation is made.

■ L3.3 Bibliography for this lecture

- [BR23] B. Brooks and P. J. Reny, “A canonical game—75 years in the making—showing the equivalence of matrix games and linear programming,” *Econ. Theory Bull.*, vol. 11, no. 2, pp. 171–180, 2023, doi: 10.1007/s40505-023-00252-8.
- [Ste23] B. von Stengel, “Zero-Sum Games and Linear Programming Duality,” *Math. Oper. Res.*, Jul. 2023.
- [Adl13] I. Adler, “The equivalence of linear programs and zero-sum games,” *International Journal of Game Theory*, vol. 42, pp. 165–177, 2013.
- [Dan82] G. B. Dantzig, “Reminiscences about the origins of linear programming,” *Oper. Res. Lett.*, vol. 1, no. 2, pp. 43–48, 1982, doi: 10.1016/0167-6377(82)90043-8.
- [KM86] E. Kohlberg and J.-F. Mertens, “On the strategic stability of equilibria,” *Econometrica: Journal of the Econometric Society*, pp. 1003–1037, 1986.
- [Mil+23] J. Millionis, C. Papadimitriou, G. Piliouras, and K. Spendlove, “An impossibility theorem in game dynamics,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 120, no. 41, 2023, doi: 10.1073/pnas.2305349120.
- [LMM03] R. J. Lipton, E. Markakis, and A. Mehta, “Playing large games using simple strategies,” in *Proceedings of the 4th ACM Conference on Electronic Commerce*, New York, NY, USA: Association for Computing Machinery, 2003, pp. 36–41. doi: 10.1145/779928.779933.

- [Nas51] J. Nash, “Non-Cooperative Games,” *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951, [Online]. Available: <http://www.jstor.org/stable/1969529>
- [NCH04] R. Nau, S. G. Canovas, and P. Hansen, “On the geometry of Nash equilibria and correlated equilibria,” *International Journal of Game Theory*, vol. 32, pp. 443–453, 2004.
- [Bub79] V. Bubelis, “On equilibria in finite games,” *International Journal of Game Theory*, vol. 8, pp. 65–79, 1979.
- [Dat03] R. S. Datta, “Universality of Nash equilibria,” *Mathematics of Operations Research*, vol. 28, no. 3, pp. 424–432, 2003.
- [PNS08] R. Porter, E. Nudelman, and Y. Shoham, “Simple search methods for finding a Nash equilibrium,” *Games Econom. Behav.*, vol. 63, no. 2, pp. 642–662, 2008, doi: 10.1016/j.geb.2006.03.015.
- [Aum74] R. J. Aumann, “Subjectivity and correlation in randomized strategies,” *J. Math. Econom.*, vol. 1, no. 1, pp. 67–96, 1974, doi: 10.1016/0304-4068(74)90037-8.
- [MV78] H. Moulin and J. P. Vial, “Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon,” *International Journal of Game Theory*, vol. 7, pp. 201–221, 1978.
- [BS19] N. Brown and T. Sandholm, “Superhuman AI for multiplayer poker,” *Science*, vol. 365, no. 6456, pp. 885–890, 2019, doi: 10.1126/science.aay2400.

Changelog

- 2025-10-05: Fixed typos (thanks Eric Yang Yu!).