Turning defense into offense in $O(\log 1/\epsilon)$ steps: Efficient constructive proof of the minimax theorem

GABRIELE FARINA
Massachusetts Institute of Technology

Von Neumann's minimax theorem asserts that the ability to defend against any opponent strategy implies the existence of an offensive strategy that guarantees the same value. This note revisits that symmetry from a constructive, oracle-based point of view. Given access to a *defense oracle* that, for any opponent strategy x, returns a response y guaranteeing payoff at least v, we ask how efficiently one can compute an *offense* strategy $y^\star$ that guarantees value at least $v - \epsilon$ against all x. A classical construction via no-regret learning yields such a $y^\star$ after $O((1/\epsilon)^2)$ calls to the defense oracle. In this note, I describe a different construction that uses only $O(\log(1/\epsilon))$ calls to the oracle (up to polynomial factors in the dimension and encoding size). I then illustrate this primitive through three applications: computing Φ-equilibria in convex and extensive-form games beyond polynomial type, computing expected solutions to variational inequalities, and computing expected fixed points of possibly discontinuous maps.

Categories and Subject Descriptors: Mathematical optimization [**Theory of computation**]: Design and analysis of algorithms

General Terms: Algorithms; Theory

Additional Key Words and Phrases: Minimax theorem; equilibrium computation; variational inequalities

1. INTRODUCTION

Suppose we are playing a two-player zero-sum game, and strategies are allowed to be mixed: each (mixed) strategy is represented as a point in a convex and compact set $\mathcal{X}$ or $\mathcal{Y}$. If for any mixed strategy $x \in \mathcal{X}$ that my opponent could pick, I know how to construct a counterstrategy $y = h(x) \in \mathcal{Y}$ that secures me a certain score v, then von Neumann's minimax theorem states that I must have a strategy $y^\star$ that secures me the same score v no matter what my opponent could pick.

This symmetry between defense and offense is anything but obvious. In early work on game theory in the 1920s, Borel was able to prove special cases of this symmetry for small symmetric zero-sum games, but wrongly conjectured that the property could not hold beyond 3×3 games (Borel, 1921; Weinstein, 2022). He was proven wrong in 1928, when von Neumann proved his minimax theorem (v. Neumann, 1928), marking a turning point for modern game theory, and much more.

Since von Neumann's original proof, a repertoire of conceptually different proofs of the minimax theorem have been proposed. Many of them proceed by way of heavy machinery, including fixed-point theorems from topology, compactness arguments in functional analysis, and separation results in convex geometry (Borwein, 2016, §1.1). However, those arguments are typically nonconstructive, in the sense that they prove the *existence* of a good offense policy, but shed little light on how one could compute it starting from knowledge of how to play defense, that is, from the function h.

Author email: `gfarina@mit.edu`

This note concerns *constructive* proofs of the minimax theorem:

> *Given access to an oracle h that constructs good defense policies, how can we construct a good offense policy?*

Crucially, we will not require that $\mathcal{Y}$ be efficiently representable: we will simply require that h can be queried. This will enable constructing good offense policies even if $\mathcal{Y}$ is too complex to optimize over, as will be the case in the three applications mentioned in Sections 5–7.

As we recall in Section 3, a classical (and beautiful) such constructive proof is possible via a connection to regret minimization. However, that construction is quite slow: to find an offense strategy that gives value at least $v - \epsilon$ starting from defense strategies of value v, the construction requires evaluating $O((1/\epsilon)^2)$ defensive strategies.

In contrast, I will present a different construction that only requires $\log(1/\epsilon)$ calls to h. This rate improves exponentially over the regret-based one, and shows that the cost of computing near-optimal offense strategies can grow polynomially in the number of accuracy *bits*, rather than the inverse of the accuracy itself. As a consequence, under standard assumptions, such as polyhedrality of the strategy sets, this improvement leads to *exact* computation of offensive strategy in polynomial time. The construction extends the ellipsoid-against-hope algorithm of Papadimitriou and Roughgarden (2008) for computing correlated equilibria in succinct games, generalizing it beyond correlated equilibria, beyond polyhedrality, and beyond the need for a polynomial number of vertices, and is presented in a recent paper that Charis Pipis and I wrote for NeurIPS (Farina and Pipis, 2024).

We found this construction to be a helpful computational primitive for a number of problems that bear a connection with the minimax theorem, including applications in equilibrium computation, fixed points, and variational inequalities, as briefly recounted in Sections 5–7.

2. SETUP AND NOTATION

Let $\mathcal{X}$ and $\mathcal{Y}$ be convex and compact domains for which we have oracle access (*e.g.*, an efficient separation oracle, or a projection oracle), and $u : \mathcal{X} \times \mathcal{Y} \to [0, 1]$ be a bilinear utility function, that is, $u(x, y) = x^\top Ay$ for some matrix A. The x-player will want to minimize u, while the y-player will want to maximize it.

Von Neumann's minimax theorem states that

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) = \max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} u(x, y).$$

The x-player is our opponent, while we are the y-player.

For any strategy $x \in \mathcal{X}$ of the opponent, the value $\max_{y\in\mathcal{Y}} u(x, y)$ is the best score that we can secure for ourselves by responding to x. Hence, the statement $\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) \geq v$ means that no matter what strategy x the opponent picks, we can always respond with a strategy that guarantees us a score of at least v. In other words, the statement $\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} u(x, y) \geq v$ is equivalent to the existence of a function $h : \mathcal{X} \to \mathcal{Y}$ such that for all $x \in \mathcal{X}$, $u(x, h(x)) \geq v$. Such a function h is what we will call a *defense oracle*.

On the other hand, the statement $\max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} u(x, y) \geq v$ means that there is a strategy $y^\star \in \mathcal{Y}$ for us, such that no matter how the opponent responds, our score will be at least v. Von Neumann's minimax theorem states the rather unintuitive fact that if a defense oracle h exists, then such a strategy $y^\star$ must also exist.

In the following we will be interested in the following question: *given access to a defense oracle h, how can we construct a strategy $y^\star$ that guarantees us a score of at least $v-\epsilon$ against any opponent's strategy, in $O(\log(1/\epsilon))$ time (ignoring polynomial factors in the dimension and diameter of the sets $\mathcal{X}$ and $\mathcal{Y}$)?*

3. WARMUP: A SLOW CONSTRUCTION USING REGRET MINIMIZATION

A classical, neat, and deceptively short constructive proof of the minimax theorem can be derived from the mere fact that no-regret learning algorithms exist.

The idea is very simple. Suppose we are a y-player concerned with constructing an *offense* strategy $y^\star$ for ourselves, but so far have only figured out how to play *defense* via a defense oracle h. We do not know how our opponent, the x-player, will behave, so we'd better do some training. We will do so by simulating a plausible, fictitious x-player who plays according to a no-regret learning algorithm. In each round t, the x-player selects a strategy $x_t \in \mathcal{X}$, and we respond by playing the defense strategy $y_t = h(x_t)$, inducing a loss of Ay_t to the x-player.

After T rounds of this simulation, the regret incurred by the x-player is by definition

$$\mathrm{Reg}_T = \max_{\hat{x}\in\mathcal{X}} \sum_{t=1}^{T} (x_t^\top A y_t - \hat{x}^\top A y_t).$$

If the simulated x-player is no-regret, then by definition the regret is sublinear in the time horizon T. For concreteness, several algorithms are known to guarantee regret on the order of $\sqrt{T}$, ignoring polynomial factors in the dimension and diameter of the strategy set $\mathcal{X}$.

Dividing by T, and plugging in the bound $\mathrm{Reg}_T \le \sqrt{T}$ to get an estimate, we find that

$$\max_{\hat{x}\in\mathcal{X}} \frac{1}{T}\sum_{t=1}^{T} \hat{x}^\top A y_t \ge \frac{1}{T}\sum_{t=1}^{T} x_t^\top A y_t - \frac{1}{\sqrt{T}}.$$

The value of $x_t^\top A y_t$ is at least v for all t, since y_t is a defense strategy by construction. Furthermore, the average $(1/T)\sum_{t=1}^{T} \hat{x}^\top A y_t$ is equal to $\hat{x}^\top A\bar{y}$, where $\bar{y} = (1/T)\sum_{t=1}^{T} y_t$ is the average of the defense strategies played. Hence, we have shown that for any $\hat{x} \in \mathcal{X}$,

$$\hat{x}^\top A\bar{y} \ge v - \frac{1}{\sqrt{T}}.$$

In other words, the average strategy $\bar{y}$ is an offense strategy that guarantees us a score of at least $v - 1/\sqrt{T}$ against any opponent's strategy. Choosing $T = 1/\epsilon^2$, we find that $\bar{y}$ is an offense strategy that guarantees us a score of at least $v - \epsilon$.

In summary, the mere existence of regret minimization (for which many readily implementable algorithms exist) gives rise to a constructive proof of the minimax theorem. However, this construction is quite slow, requiring on the order of $1/\epsilon^2$ calls to the defense oracle h to construct an offense strategy that guarantees us a score of at least $v - \epsilon$.

4. TURNING DEFENSE INTO OFFENSE IN $O(\log(1/\epsilon))$ STEPS

In this section, we present a different construction that turns a defense oracle h into an offense strategy $y^\star$ in only $O(\log(1/\epsilon))$ calls to the defense oracle h. We presented this construction in Farina and Pipis (2024). It borrows several important ideas from the ellipsoid against hope algorithm for computing correlated equilibria in succinct games (Papadimitriou and Roughgarden, 2008), but generalizes them into a general framework disconnected from correlated equilibria. Furthermore, unlike Papadimitriou and Roughgarden (2008), our construction does not require polyhedral sets with a polynomial number of vertices, and as we will show later applies to games of non-polynomial type as well.

A main ingredient of the construction is the following. Consider the set of "unbeatable" strategies for the opponent, defined as

$$\Omega \coloneqq \{x \in \mathcal{X} : x^\top A y < v \quad \forall y \in \mathcal{Y}\}.$$

Being defined by linear inequalities, this set is convex. More importantly, this set is *empty*, since by assumption we have access to an oracle h that for any $x \in \mathcal{X}$ constructs a strategy $y = h(x)$ that guarantees us value at least v against x. For any strategy $x \in \mathcal{X}$, we can certify that $x \notin \Omega$ by considering the linear constraint corresponding to $y = h(x)$, which is violated, and hence we have a separation oracle for the set Ω.

Step 1: Reduction of constraints via the ellipsoid method. Ignoring details about numerical precision and bit representations, having access to a separation oracle for Ω allows us to run the ellipsoid method to certify that Ω is near empty, in the sense of having arbitrarily low volume $\le \epsilon$, in a number of steps polynomial in the dimension and diameter of $\mathcal{X}$ and logarithmic in $1/\epsilon$. Each step of the ellipsoid method requires a call to the separation oracle for the ellipsoid center x, which in turn requires a call to the defense oracle h if $x \in \mathcal{X}$, or a call to the separation oracle for $\mathcal{X}$ if $x \notin \mathcal{X}$.

One can think of the ellipsoid method as a method for selecting, out of the infinitely many that define Ω, a polynomial subset of constraints that already certify (near) emptiness of the set. Under suitable assumptions on the bit representation of the problem, this can be easily also turned into a polynomial certificate of emptiness.

Step 2: An application of Farkas lemma. Once the set of constraints has been appropriately sparsified over the course of $T = O(\log(1/\epsilon))$ steps, we are left with a set of strategies

$$\tilde{\Omega} \coloneqq \{x \in \mathcal{X} : x^\top A y_t < v \quad \forall t = 1, \dots, T\}$$

defined by a *finite* subsystem of linear inequalities that is already empty (or near-empty).

The Farkas lemma fully characterizes what it means for a set of linear inequalities to be infeasible: the only way that can happen is if there is a nonnegative, nontrivial linear combination of the inequalities that gives rise to a contradiction. In our case, this means that there must exist nonnegative multipliers $\lambda_1, \dots, \lambda_T \ge 0$, not all zero, such that

$$\sum_{t=1}^{T} \lambda_t x^\top A y_t \ge v \sum_{t=1}^{T} \lambda_t \quad \forall x \in \mathcal{X}.$$

Dividing both sides by $\sum_{t=1}^{T} \lambda_t$ (which is positive by nontriviality), we find that there must exist a convex combination of the strategies y_t (which were obtained by calling the oracle h on the feasible centers of the ellipsoids constructed during the ellipsoid method) that guarantees us value at least v against any opponent's strategy $x \in \mathcal{X}$.

To find such multipliers λ_t, it suffices to maximize the concave function

$$d : \Delta^T \to \mathbb{R}, \qquad d(\lambda) \coloneqq \min_{x \in \mathcal{X}} \sum_{t=1}^{T} \lambda_t x^\top A y_t$$

over the T-dimensional simplex Δ^T. This is typically straightforward to carry out efficiently using standard convex optimization methods.

Compared to the regret minimization construction of Section 3, the construction in this section does not use uniform averaging, but at the same time it requires simulating defending against only $T = O(\log(1/\epsilon))$ strategies of the opponent before an offensive strategy can be extracted.

Remark 4.1. In the above, the convexity and compactness of $\mathcal{X}$ were needed to run the ellipsoid method on Ω. However, only convexity (and not compactness) was used for $\mathcal{Y}$. In other words, the same argument yields a constructive proof of the slightly stronger statement that

$$\min_{x \in \mathcal{X}} \sup_{y \in \mathcal{Y}} u(x, y) = \sup_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} u(x, y)$$

whenever $\mathcal{X}$ is convex and compact, $\mathcal{Y}$ is convex, and u is bilinear. This will come in handy in our applications, where $\mathcal{Y}$ will often be the space of distributions over a convex compact set.

5. APPLICATION 1: EFFICIENT COMPUTATION OF EQUILIBRIA IN CONVEX GAMES

We can combine the previous construction with an observation by Hart and Schmeidler (1989) to produce algorithms that compute exact Φ-equilibria in multiplayer games, beyond the polynomial type requirement of the original paper of Papadimitriou and Roughgarden (2008).

To fix ideas, let's suppose we would like to compute an exact coarse correlated equilibrium in a convex game. Each player i has a convex and compact strategy set $\mathcal{X}_i$, and a utility function $u_i : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \to [0, 1]$ that is multilinear. A coarse correlated equilibrium is a joint distribution over strategy profiles such that no player can gain in expectation by unilaterally deviating to any fixed strategy. Formally, a distribution y over $\mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ is a coarse correlated equilibrium if for all players i and all strategies $x_i \in \mathcal{X}_i$,

$$\mathbb{E}_{z \sim y}[u_i(z)] \geq \mathbb{E}_{z \sim y}[u_i(x_i, z_{-i})].$$

We can think of coarse correlated equilibria as a minimax strategy for the game between a *mediator* who picks y, and an *opponent* who picks a player i and a deviation strategy $x_i \in \mathcal{X}_i$ to check whether the equilibrium conditions are satisfied. It is direct to check that the preconditions (Remark 4.1) for the construction of

Section 4 hold. To construct an optimal strategy for the mediator player y we can then start from determining whether y has a defense oracle against any x.

> In the case of coarse correlated equilibria, a defense oracle is quite simple: if the mediator knows the strategy of the opponent, that is, the distribution over pairs (i, x_i), they can respond by picking a product distribution y whose moments match the marginals of the opponent's distribution. This trivially guarantees that no player can gain by deviating to any fixed strategy, since the expected utility of each player under y is equal to the expected utility when the opponent picks their deviation.

The construction of Section 4 then allows us to turn this defense oracle into an offense strategy (distribution over strategy profiles) that works no matter what the opponent picks, and that is a coarse correlated equilibrium by definition.

Since each product distribution returned by the defense oracle admits a succinct representation (via its marginals), all while allowing d to be maximized efficiently, this construction shows that coarse correlated equilibria can be computed in polynomial time well beyond games of polynomial type. In particular, for extensive-form games, which were explicitly excluded by Papadimitriou and Roughgarden (2008), it shows that a coarse correlated equilibrium can be computed in time polynomial in the game tree size.

Finally, the same idea applies more generally to more complicated notions of Φ-equilibria, as long as a defense oracle can be constructed efficiently (usually via computing a suitable fixed point of the deviation functions). For example, this was a key ingredient in Daskalakis et al. (2025). Furthermore, this framework recovers the algorithm for extensive-form correlated equilibria of Huang and von Stengel (2008) as a particular instantiation.

6. APPLICATION 2: VARIATIONAL INEQUALITIES

Another application that my coauthors and I found useful concerns variational inequalities. We studied this in a joint paper with Zhang, Anagnostides, Tewolde, Berker, Farina, Conitzer, and Sandholm (2025a).

Given a convex and compact set $\mathcal{Z} \subseteq \mathbb{R}^d$ and a function $F : \mathcal{Z} \to \mathbb{R}^d$, the variational inequality problem asks to find a point $z^\star \in \mathcal{Z}$ such that

$$F(z^\star)^\top (x - z^\star) \geq 0 \quad \forall x \in \mathcal{Z}.$$

When F is the gradient of a function f, the variational inequality captures the first-order optimality conditions for the constrained optimization problem $\min_{z \in \mathcal{Z}} f(z)$. When $\mathcal{Z} = \Delta^m \times \Delta^n$ and F is the linear map

$$F(z_1, z_2) = -\begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix},$$

the variational inequality captures the Nash equilibria of the bimatrix game with payoff matrices $A, B \in \mathbb{R}^{m \times n}$. Given the hardness of approximating Nash equilibria, the previous example immediately shows that, even for *linear* maps F, approximating a solution to the variational inequality is computationally intractable.

However, we can compute a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\mathbb{E}_{z\sim y}[F(z)^\top (x - z)] \geq -\epsilon \quad \forall x \in \mathcal{Z} \tag{1}$$

in $\text{poly}(d, \log(1/\epsilon))$ time, using the construction of Section 4.[1] In other words, as long as we are willing to seek a *distribution* over points in $\mathcal{Z}$ (rather than a single point), we can bypass the hardness results tied to variational inequalities.

To do so, observe that (1) captures a game between the y-player and the x-player, with bilinear utility $(x, y) \mapsto \mathbb{E}_{z\sim y}[F(z)^\top (x - z)]$.

A defense oracle of value 0 for the y-player is easy to construct. If we know the point $x \in \mathcal{Z}$ that the opponent picks, we could respond by selecting the distribution $y = h(x)$ that puts all the mass on the point x, which trivially satisfies

$$\mathbb{E}_{z\sim y}[F(z)^\top (x - z)] = F(x)^\top (x - x) = 0 \geq 0.$$

Hence, we can use our general construction from Section 4 to efficiently upgrade the defense oracle into an "offensive" point $y \in \Delta(\mathcal{Z})$ that works for all x, that is, a solution to the expected variational inequality problem (1).

In a recent paper (Zhang et al., 2025a), we further strengthen this construction to show that we can efficiently compute a stronger notion of expected solution: find a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\mathbb{E}_{z\sim y}[F(z)^\top (Q(z) - z)] \geq -\epsilon \quad \forall \text{ affine endomorphisms } Q : \mathcal{Z} \to \mathcal{Z}.$$

The special case in which Q is taken to be a constant map recovers (1). It can be shown easily that for suitably constructed operators F, these notions recover coarse-correlated equilibria and refinements of correlated equilibria in concave games.

For general operators F (i.e., not necessarily representing utility functions in games), it is not understood what these concepts recover, but to my knowledge, they are the strongest *tractable* relaxation of the notion of solution for general variational inequalities that we can compute today.

7. APPLICATION 3: EXPECTED FIXED POINTS

As a third application of this construction, consider the problem of computing a fixed point of a function $f : \mathcal{Z} \to \mathcal{Z}$. When f is continuous and $\mathcal{Z} \subseteq \mathbb{R}^d$ is convex and compact, Brouwer's fixed point theorem guarantees the existence of a fixed point $z^\star \in \mathcal{Z}$ such that $f(z^\star) = z^\star$. However, computing such a fixed point is in general computationally intractable, given its intimate connection to Nash equilibria. However, in a recent paper with Zhang, Anagnostides, Tewolde, Berker, Farina, Conitzer, and Sandholm (2025b), we show that we can compute a distribution $y \in \Delta(\mathcal{Z})$ over $\mathcal{Z}$ such that

$$\|\mathbb{E}_{z\sim y}[f(z) - z]\|_2 \leq \epsilon \tag{2}$$

in $O(\log(1/\epsilon))$ steps, by reduction to our construction of Section 4.

[1]Problems of the form (1) are called *outgoing minimax problems* by Foster and Hart (2021).

We can interpret (2) as a game. On one side, we, the y-player, seek to find a distribution with the property above. To keep us honest, our opponent, the x-player, can select a direction x from the unit Euclidean ball $\mathbb{B}^d$ and win the game if they can prove that y does not induce (in expectation) a fixed point in the direction x, that is, if $\mathbb{E}_{z\sim y}[x^\top(f(z)-z)] < 0$. In other words, we are seeking to solve the game

$$\arg\max_{y\in\Delta(\mathcal{Z})}\min_{x\in\mathbb{B}^d}\mathbb{E}_{z\sim y}[x^\top(f(z)-z)],$$

whose value we know is 0. Once again, we are within the hypotheses of the main construction (see also Remark 4.1).

To solve this game, we can again think in terms of defense and offense. If we knew the direction x that the opponent would pick, we could respond by picking as response the distribution $y = h(x)$ defined as

$$h(x) = \text{distribution putting all mass on} \arg\max_{z\in\mathcal{Z}} x^\top z.$$

This is a valid defense oracle: it guarantees $\mathbb{E}_{z\sim y}[x^\top(f(z)-z)] \le 0$ given that $x^\top z \ge x^\top f(z)$ for the maximizer z in the support of y by construction.

We can then use the construction of Section 4 to turn this defense oracle into a distribution $y \in \Delta(\mathcal{Z})$, that works no matter the direction x picked by the opponent. In particular, a convex combination of the defense strategies constructed during the process can be constructed in time polynomial in the dimension and $\log(1/\epsilon)$, and yields an offense strategy $y^\star$ such that

$$-\epsilon \le \min_{x\in\mathbb{B}^d}\mathbb{E}_{z\sim y^\star}[x^\top(f(z)-z)] = -\|\mathbb{E}_{z\sim y^\star}[f(z)-z]\|_2,$$

that is,

$$\|\mathbb{E}_{z\sim y^\star}[f(z)-z]\|_2 \le \epsilon.$$

As a side note, this construction works even if f is not continuous. We found it to be a fundamental building block for sidestepping the computational equivalence barrier of Hazan and Kale (2007) when constructing Φ-regret minimizers for nonlinear deviations Φ (Zhang et al., 2025b).

8. ACKNOWLEDGMENTS

I am grateful to Samuel Sokota, Brian Hu Zhang, Juan Carlos Perdomo, Jon Schneider, Martino Bernasconi, Charis Pipis, and Samuel Taggart for their useful feedback on various drafts of this note.

REFERENCES

Borel, E. 1921. La théorie du jeu et les équations intégralesa noyau symétrique. *Comptes rendus de l'Académie des Sciences 173,* 1304-1308, 58.

Borwein, J. M. 2016. A very complicated proof of the minimax theorem. *Minimax Theory and its Applications 1,* 1, 21–27.

Daskalakis, C., Farina, G., Fishelson, M., Pipis, C., and Schneider, J. 2025. Efficient learning and computation of linear correlated equilibrium in general convex games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 542–553.

Farina, G. and Pipis, C. 2024. Polynomial-time computation of exact Φ-equilibria in polyhedral games. *Advances in Neural Information Processing Systems (NeurIPS) 37*, 4207–4234.

Foster, D. P. and Hart, S. 2021. Forecast hedging and calibration. *Journal of Political Economy 129,* 12, 3447–3490.

Hart, S. and Schmeidler, D. 1989. Existence of correlated equilibria. *Mathematics of Operations Research 14,* 1, 18–25.

Hazan, E. and Kale, S. 2007. Computational equivalence of fixed points and no regret algorithms, and convergence to equilibria. *Advances in Neural Information Processing Systems 20*.

Huang, W. and von Stengel, B. 2008. Computing an extensive-form correlated equilibrium in polynomial time. In *International Workshop on Internet and Network Economics (WINE)*. 506–513.

Papadimitriou, C. H. and Roughgarden, T. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM) 55,* 3, 1–29.

v. Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen 100,* 1, 295–320.

Weinstein, J. 2022. On a mistranslation of a mistake about minimax.

Zhang, B. H., Anagnostides, I., Tewolde, E., Berker, R. E., Farina, G., Conitzer, V., and Sandholm, T. 2025a. Expected variational inequalities. In *International Conference on Machine Learning (ICML)*.

Zhang, B. H., Anagnostides, I., Tewolde, E., Berker, R. E., Farina, G., Conitzer, V., and Sandholm, T. 2025b. Learning and computation of Φ-equilibria at the frontier of tractability. In *ACM Conference on Economics and Computation (ACM EC)*. 1105–1105.