CS 405/6001: Game Theory and Algorithmic Mechanism Design	Jul–Nov 2026

Lecture 6: Mixed Strategy Nash Equilibrium

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

6.1 Recap

Last lecture we saw that a matrix game need not have a PSNE – the Ghost-versus-Sentinel two-channel duel is a case in point, where $\overline{v} \neq \underline{v}$ – and that mixed strategies, probability distributions over pure strategies, are the natural way to restore an equilibrium notion. Today we define **mixed strategy Nash equilibrium (MSNE)**, characterize it in a way that is actually computable, and prove that it always exists in finite games.

6.2 Mixed Strategy Nash Equilibrium

Definition 6.1 (Mixed Strategy Nash Equilibrium) *A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a* ***mixed strategy Nash equilibrium (MSNE)*** *if*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(\sigma_i, \sigma_{-i}^*), \qquad \forall \sigma_i \in \Delta S_i,\ \forall i \in N.$$

Since a pure strategy is the special case of a mixed strategy that places all probability mass on a single element, every PSNE is trivially an MSNE. The converse is false: the Ghost-Sentinel duel from Lecture 5 (Example 5.2, matrix $u = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$) has no PSNE at all, yet, as we will see shortly, it has an MSNE.

Theorem 6.2 (Pure-deviation suffices) *A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE if and only if, for every $i \in N$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(s_i, \sigma_{-i}^*), \qquad \forall s_i \in S_i.$$

Proof: ($\Rightarrow$) A pure strategy s_i is itself a mixed strategy (probability 1 on s_i), so the inequality is a special case of the MSNE definition.

($\Leftarrow$) Suppose the inequality holds for every pure $s_i \in S_i$. Pick an arbitrary mixed strategy σ_i of player i. By definition of expected utility,

$$u_i(\sigma_i, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i(s_i)\, u_i(s_i, \sigma_{-i}^*) \leqslant \sum_{s_i \in S_i} \sigma_i(s_i)\, u_i(\sigma_i^*, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*) \sum_{s_i \in S_i} \sigma_i(s_i) = u_i(\sigma_i^*, \sigma_{-i}^*),$$

where the inequality applies the hypothesis to every s_i in the sum, term by term (each term's coefficient $\sigma_i(s_i) \geqslant 0$), and the last equality uses $\sum_{s_i} \sigma_i(s_i) = 1$. Since σ_i was arbitrary, σ_i^* is a best response to σ_{-i}^* even among mixed strategies, and $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE. ■

Theorem 6.2 is useful precisely because it reduces checking an infinite family of deviations (all of ΔS_i) to checking only the finitely many pure strategies in S_i.

6.2.1 Checking candidate profiles

Return to the Ghost-Sentinel duel from Lecture 5, with matrix $u = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$. Suppose Ghost mixes $(2/3, 1/3)$ over (L, R) and Sentinel mixes $(4/5, 1/5)$. Is this an MSNE? Sentinel's expected utility from watching L is $\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot (-1) = \frac{1}{3}$, while from watching R it is $\frac{2}{3} \cdot (-1) + \frac{1}{3} \cdot 1 = -\frac{1}{3}$. Since these differ, Sentinel is not indifferent between its two pure strategies while playing this mix, so by Theorem 6.2 Sentinel has a profitable pure deviation (to L, which gives strictly higher expected utility): this profile is *not* an MSNE.

This suggests some balance in the utilities is needed for a genuinely mixed equilibrium. The natural next candidate is the fully balanced profile, $(1/2, 1/2)$ for both players; we verify in Section 6.4 that this is indeed the MSNE.

6.3 Support and the Characterization Theorem

Definition 6.3 (Support) *For a mixed strategy σ_i, its* ***support*** *is $\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$, the set of pure strategies played with positive probability.*

Theorem 6.4 (MSNE characterization) *A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE if and only if, for every $i \in N$:*

(C1) $u_i(s_i, \sigma_{-i}^*)$ is identical for every $s_i \in \delta(\sigma_i^*)$ (indifference on *the support), and*

(C2) $u_i(s_i, \sigma_{-i}^*) \geqslant u_i(s_i', \sigma_{-i}^*)$ for every $s_i \in \delta(\sigma_i^*)$ and every $s_i' \notin \delta(\sigma_i^*)$ (no profitable deviation off *the support).*

We defer the proof of Theorem 6.4 to Section 6.5, and first put it to work.

6.3.1 Solving the Ghost-Sentinel duel

We search over possible support profiles for $u = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$.

Case 1: $(\{L\}, \{L\})$**.** Ghost's only pure best response would need to make $s_1' = R$ unprofitable, but Ghost playing R against Sentinel's pure L gives $u_1(R, L) = 1 > -1 = u_1(L, L)$: this violates (C2). By symmetry every singleton-support profile fails.

Case 2: $(\{L, R\}, \{L\})$ (and symmetric cases with one player mixing, the other pure). Condition (C1) would require $u_1(L, L) = u_1(R, L)$, i.e., $-1 = 1$, which is false: violates (C1).

Case 3: $(\{L, R\}, \{L, R\})$**.** Condition (C2) is vacuous here (there is no strategy outside the support). For (C1), let Ghost play L with probability p and Sentinel play L with probability q. Indifference for Ghost requires

$$u_1(L, (q, 1-q)) = u_1(R, (q, 1-q)) \implies (-1)q + 1 \cdot (1-q) = 1 \cdot q + (-1)(1-q) \implies q = \frac{1}{2},$$

and indifference for Sentinel (by the symmetric computation on $u_2 = -u_1$) requires $p = \frac{1}{2}$. Both conditions are satisfiable, so

$$\text{MSNE} = \left(\left(\tfrac{1}{2}, \tfrac{1}{2}\right), \left(\tfrac{1}{2}, \tfrac{1}{2}\right)\right)$$

is the unique MSNE of the Ghost-Sentinel duel, confirming the guess from Section 6.2.

6.3.2 Exercises

As practice, revisit Riya and Karan's Chess-or-Football coordination game from Lecture 3 (Week 2), payoffs

$$\begin{pmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{pmatrix}, \qquad \text{rows/columns } (F, C) \text{ for each player,}$$

and find its (genuinely mixed) MSNE using the same case analysis – note that this game, unlike the Ghost-Sentinel duel, *also* has two PSNE, which by Theorem 6.2 are themselves (degenerate) MSNE. As a further exercise, suppose Karan gets a third option D (a documentary neither prefers, but not a total mismatch either), with payoffs

$$\begin{pmatrix} (2,1) & (0,0) & (1,1) \\ (0,0) & (1,2) & (2,0) \end{pmatrix}, \qquad \text{rows } (F, C) \text{ for Riya, columns } (F, C, D) \text{ for Karan,}$$

and determine which support profiles can possibly support an MSNE.

6.4 Proof of the Characterization Theorem

We now prove Theorem 6.4. First, two observations.

- Maximizing a linear functional over a probability distribution is the same as maximizing over the underlying pure strategies: $\max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i})$, since expected utility is a weighted average of the $u_i(s_i, \sigma_{-i})$'s and a weighted average of numbers never exceeds their maximum, with equality achieved by placing the entire probability mass on a maximizing pure strategy.
- If $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE, the maximizing pure strategy above must lie in $\delta(\sigma_i^*)$: if some $s_i' \notin \delta(\sigma_i^*)$ achieved the maximum utility instead, moving all probability mass to s_i' would give player i a strictly higher payoff, contradicting that σ_i^* was already optimal.

Proof:[Proof of Theorem 6.4] ($\Rightarrow$) Suppose $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE. By the two observations above,

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*). \tag{6.1}$$

On the other hand, by definition of expected utility restricted to the support,

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i)\, u_i(s_i, \sigma_{-i}^*). \tag{6.2}$$

Equations (6.1) and (6.2) say that the *maximum* of the values $\{u_i(s_i, \sigma_{-i}^*)\}_{s_i \in \delta(\sigma_i^*)}$ equals a strictly positive weighted *average* of exactly those same values (the weights $\sigma_i^*(s_i)$ are positive on $\delta(\sigma_i^*)$ by definition of support, and sum to 1). A weighted average of numbers equals their maximum only when all the numbers being averaged are equal to that maximum. This proves (C1).

For (C2): suppose, for contradiction, that there exist $s_i \in \delta(\sigma_i^*)$ and $s_i' \notin \delta(\sigma_i^*)$ with $u_i(s_i, \sigma_{-i}^*) < u_i(s_i', \sigma_{-i}^*)$. Shifting the probability mass $\sigma_i^*(s_i)$ from s_i to s_i' produces a new mixed strategy with strictly higher expected utility for player i, contradicting that σ_i^* was a best response. Hence no such pair exists, proving (C2). This completes the necessary direction.

($\Leftarrow$) Suppose (C1) and (C2) hold. By (C1), let $m_i(\sigma^*_{-i}) := u_i(s_i, \sigma^*_{-i})$ for every $s_i \in \delta(\sigma^*_i)$ (a common value). By (C2), this common value is at least as large as $u_i(s'_i, \sigma^*_{-i})$ for every s'_i outside the support, and it trivially equals itself on the support, so

$$m_i(\sigma^*_{-i}) = \max_{s_i \in S_i} u_i(s_i, \sigma^*_{-i}).$$

Now compute, using the definition of expected utility restricted to the support, then the common-value observation, then the max-over-support-equals-max-over-all-pure-strategies identity just derived, then the first bulleted observation above:

$$u_i(\sigma^*_i, \sigma^*_{-i}) = \sum_{s_i \in \delta(\sigma^*_i)} \sigma^*_i(s_i)\, u_i(s_i, \sigma^*_{-i}) = m_i(\sigma^*_{-i}) = \max_{s_i \in S_i} u_i(s_i, \sigma^*_{-i}) = \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma^*_{-i}) \geqslant u_i(\sigma_i, \sigma^*_{-i})$$

for every $\sigma_i \in \Delta S_i$. Since i was arbitrary, $(\sigma^*_i, \sigma^*_{-i})$ is an MSNE. ∎

6.5 An Algorithm for Finding an MSNE

Theorem 6.4 converts the search for an MSNE into a search over **support profiles**. For a game $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$, the total number of (non-empty) support profiles is

$$K = \prod_{i \in N} \left(2^{|S_i|} - 1\right).$$

For each candidate support profile $X_1 \times \cdots \times X_n$ (with $X_i \subseteq S_i$), solve the feasibility program

$$\begin{aligned}
w_i &= \sum_{s_{-i} \in S_{-i}} \Big(\prod_{j \neq i} \sigma_j(s_j)\Big) u_i(s_i, s_{-i}), && \forall s_i \in X_i,\ \forall i \in N,\\
w_i &\geqslant \sum_{s_{-i} \in S_{-i}} \Big(\prod_{j \neq i} \sigma_j(s_j)\Big) u_i(s_i, s_{-i}), && \forall s_i \in S_i \setminus X_i,\ \forall i \in N,\\
&\sigma_j(s_j) \geqslant 0\ \forall s_j \in S_j,\ \forall j \in N, \qquad \sum_{s_j \in X_j} \sigma_j(s_j) = 1\ \forall j \in N.
\end{aligned}$$

The first set of constraints encodes (C1) (common value w_i on the support), the second encodes (C2) (no profitable deviation off the support), and the last two encode that σ_j is a valid distribution supported on X_j. Any feasible point is an MSNE with that support.

Remark 6.5 *This program is linear only when $n = 2$ (for more players, the products $\prod_{j \neq i} \sigma_j(s_j)$ are nonlinear in the unknowns). For a general n-player game, no polynomial-time algorithm is known; in fact, the problem of computing an MSNE is* ***PPAD-complete*** *(Polynomial Parity Argument on Directed graphs) [DGP09], which is strong evidence that no efficient general algorithm exists.*

6.6 MSNE and Dominance

Example 6.6 (A gig-platform driver) *A driver picks a shift type: standard rides (T), pooled rides (M), or premium rides (B); the platform's dispatch algorithm runs in mode L or R:*

$$\begin{pmatrix} 4,1 & 2,5 \\ 1,3 & 6,2 \\ 2,2 & 3,3 \end{pmatrix}, \qquad \textit{rows } (T, M, B) \textit{ for the driver, columns } (L, R) \textit{ for the platform.}$$

Neither T nor M strictly dominates B in pure strategies ($T = (4, 2)$ beats $B = (2, 3)$ against L but loses against R; $M = (1, 6)$ loses against L). But an equal mix of T and M gives the driver expected payoff $\frac{1}{2}(4) + \frac{1}{2}(1) = 2.5$ against L and $\frac{1}{2}(2) + \frac{1}{2}(6) = 4$ against R, and $2.5 > 2$ and $4 > 3$: this mix strictly dominates B in both *columns, even though no single pure strategy does.*

Theorem 6.7 *If a pure strategy s_i is strictly dominated by some mixed strategy $\sigma_i \in \Delta S_i$ (i.e., $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ for every $s_{-i} \in S_{-i}$), then in every MSNE of the game, s_i is played with probability zero.*

Theorem 6.7 justifies discarding such strategies before running the algorithm of Section 6.4: doing so shrinks the number of support profiles K without losing any equilibrium.

6.7 Existence of Mixed Strategy Nash Equilibrium

Definition 6.8 (Finite game) *A game is **finite** if the number of players is finite and each player's strategy set S_i is finite.*

Theorem 6.9 (Nash, 1951) *Every finite game has a mixed strategy Nash equilibrium.*

Remark 6.10 *Theorem 6.9 is the theoretical bedrock beneath any training procedure that searches for an equilibrium between adversarial agents – GANs, adversarial training, and self-play in general – since it guarantees* some *solution concept exists to converge towards, regardless of how adversarial the setup: Ghost against Sentinel, a driver against a dispatch algorithm, or two neural networks locked in an arms race.*

The proof needs two ingredients from real analysis, both of which we now state without proof.

Definition 6.11 (Convex, closed, bounded, compact sets) *A set $S \subseteq \mathbb{R}^n$ is:*

- ***convex*** *if $\lambda x + (1 - \lambda)y \in S$ for every $x, y \in S$ and $\lambda \in [0, 1]$;*
- ***closed*** *if it contains all its limit points (every neighborhood of a limit point contains a point of S); e.g., $[0, 1)$ is* not *closed, since every ball around 1 meets $[0, 1)$ but $1 \notin [0, 1)$;*
- ***bounded*** *if $\exists x_0 \in \mathbb{R}^n, R \in (0, \infty)$ such that $\|x - x_0\|_2 < R$ for every $x \in S$;*
- ***compact*** *if it is both closed and bounded.*

Theorem 6.12 (Brouwer's fixed point theorem) *If $S \subseteq \mathbb{R}^n$ is convex and compact, and $T : S \to S$ is continuous, then T has a fixed point: there exists $x^* \in S$ with $T(x^*) = x^*$.*

We now give the proof of Theorem 6.9 in full, for two players (the n-player case is a direct, notationally heavier, extension of the same construction).

Proof:[Proof of Theorem 6.9] Define the standard simplex

$$\Delta^k = \left\{ x \in \mathbb{R}^{k+1}_{\geqslant 0} : \sum_{i=1}^{k+1} x_i = 1 \right\},$$

which is a convex and compact subset of $\mathbb{R}^{k+1}$ (it is an intersection of half-spaces and a hyperplane, hence convex and closed, and it sits inside the unit ball, hence bounded).

Let player 1 have m pure strategies, labeled $1, \dots, m$, and player 2 have n pure strategies, labeled $1, \dots, n$. A mixed strategy for player 1 is a point $p \in \Delta^{m-1}$, and for player 2 a point $q \in \Delta^{n-1}$; the set of mixed strategy profiles is $\Delta^{m-1} \times \Delta^{n-1}$. Since we have two players, their utilities are captured by matrices $A, B \in \mathbb{R}^{m\times n}$ (the utilities of players 1 and 2, respectively, at the pure profile given by row and column), and for mixed strategies $p \in \Delta^{m-1}$, $q \in \Delta^{n-1}$,

$$u_1(p,q) = p^\top Aq, \qquad u_2(p,q) = p^\top Bq.$$

Define, for $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n\}$,

$$c_i(p,q) = \max\{A_i q - p^\top Aq,\ 0\}, \qquad d_j(p,q) = \max\{p^\top B_j - p^\top Bq,\ 0\},$$

where A_i denotes the i-th row of A and B_j the j-th column of B. Intuitively, $c_i(p,q)$ is the amount by which pure strategy i beats player 1's current mix p against q (zero if i is not an improvement), and symmetrically for d_j. Both quantities are non-negative by construction, and both are continuous in (p,q) since $A_i q$, $p^\top Aq$, $p^\top B_j$, $p^\top Bq$ are all continuous (indeed bilinear), and $\max\{\cdot, 0\}$ preserves continuity.

Define P, Q coordinate-wise by

$$P_i(p,q) = \frac{p_i + c_i(p,q)}{1 + \sum_{k=1}^m c_k(p,q)}, \quad i = 1, \dots, m; \qquad Q_j(p,q) = \frac{q_j + d_j(p,q)}{1 + \sum_{k=1}^n d_k(p,q)}, \quad j = 1, \dots, n.$$

Each $P_i(p,q) \geqslant 0$ (numerator and denominator both non-negative, denominator strictly positive), and

$$\sum_{i=1}^m P_i(p,q) = \frac{\sum_i p_i + \sum_i c_i(p,q)}{1 + \sum_k c_k(p,q)} = \frac{1 + \sum_i c_i(p,q)}{1 + \sum_k c_k(p,q)} = 1,$$

using $\sum_i p_i = 1$. Hence $P(p,q) \in \Delta^{m-1}$, and symmetrically $Q(p,q) \in \Delta^{n-1}$. Define

$$T(p,q) = (P(p,q), Q(p,q)).$$

Then $T : \Delta^{m-1} \times \Delta^{n-1} \to \Delta^{m-1} \times \Delta^{n-1}$ maps a convex, compact set into itself, and T is continuous since each c_i, d_j is continuous and the maps P_i, Q_j are continuous functions of them (the denominators are always at least 1, so there is no division-by-zero issue). By Brouwer's fixed point theorem (Theorem 6.12), there exists $(p^*, q^*) \in \Delta^{m-1} \times \Delta^{n-1}$ with

$$T(p^*, q^*) = (p^*, q^*).$$

Claim 6.13 $\sum_{k=1}^m c_k(p^*, q^*) = 0$ *and* $\sum_{k=1}^n d_k(p^*, q^*) = 0.$

Proof:[Proof of Claim 6.13] Suppose, for contradiction, that $\sum_{k=1}^m c_k(p^*, q^*) > 0$. Since (p^*, q^*) is a fixed point of T, for every i,

$$p_i^* = \frac{p_i^* + c_i(p^*, q^*)}{1 + \sum_{k=1}^m c_k(p^*, q^*)} \implies p_i^* \left(\sum_{k=1}^m c_k(p^*, q^*) \right) = c_i(p^*, q^*). \tag{6.3}$$

Let $I = \{i : p_i^* > 0\}$. By (6.3) and the standing assumption $\sum_k c_k(p^*, q^*) > 0$, we have $p_i^* > 0 \iff c_i(p^*, q^*) > 0$, and by definition of c_i, $c_i(p^*, q^*) > 0 \iff A_i q^* > p^{*\top} Aq^*$. So

$$I = \{i : p_i^* > 0\} = \{i : c_i(p^*, q^*) > 0\} = \{i : A_i q^* > p^{*\top} Aq^*\}.$$

Write $u_1^* := p^{*\top} A q^*$. Then

$$u_1^* = \sum_{i=1}^{m} p_i^*(A_i q^*) = \sum_{i\in I} p_i^*(A_i q^*) > \left(\sum_{i\in I} p_i^*\right) u_1^* = u_1^*,$$

where the first equality expands $u_1^* = p^{*\top} A q^*$ over pure strategies, the second uses that $p_i^* = 0$ outside I, the strict inequality uses $A_i q^* > u_1^*$ for every $i \in I$ together with $p_i^* > 0$ on I (so each term in the sum is strictly increased by replacing $A_i q^*$ with the smaller number u_1^*), and the last equality uses $\sum_{i\in I} p_i^* = \sum_{i=1}^{m} p_i^* = 1$. This is a contradiction ($u_1^* > u_1^*$). Hence $\sum_k c_k(p^*, q^*) = 0$. The identical argument, applied to player 2 with the roles of rows/columns and A/B exchanged, gives $\sum_k d_k(p^*, q^*) = 0$. ■

By Claim 6.13 and non-negativity of each c_k, d_k, we get $c_k(p^*, q^*) = 0$ for every $k = 1, \ldots, m$ and $d_k(p^*, q^*) = 0$ for every $k = 1, \ldots, n$. By definition of c_i, $c_i(p^*, q^*) = 0$ means $A_i q^* \leqslant p^{*\top} A q^*$ for every i. Then for *any* mixed strategy p' of player 1,

$$p'^{\top} A q^* = \sum_{i=1}^{m} p_i'(A_i q^*) \leqslant \sum_{i=1}^{m} p_i' \,(p^{*\top} A q^*) = p^{*\top} A q^*,$$

i.e., p^* is a best response to q^* among *all* mixed strategies of player 1. By the identical argument using $d_j(p^*, q^*) = 0$ for every j, q^* is a best response to p^* among all mixed strategies of player 2. Therefore (p^*, q^*) is a mixed strategy Nash equilibrium. ■

Remark 6.14 *The construction generalizes to n players by taking the domain to be $\Delta^{|S_1|-1} \times \cdots \times \Delta^{|S_n|-1}$ and defining an analogous "regret" function $c_i^{(k)}$ for each player k and each of their pure strategies i, measuring how much that pure strategy improves on the player's current mix given everyone else's mix; the rest of the argument – fixed point via Brouwer, then the claim that all regrets vanish at the fixed point, then translating vanishing regret into a best-response condition – goes through unchanged.*

6.8 Summary and Preview

We defined MSNE and showed that checking only pure deviations suffices (Theorem 6.2); characterized MSNE via indifference on the support and no improving deviation off it (Theorem 6.4), and used this to solve the Ghost-Sentinel duel by hand (Section 6.3.1); turned the characterization into a support-enumeration algorithm that is a linear feasibility program only for two players, and is PPAD-complete in general [DGP09]; showed that strategies strictly dominated even by a mixed strategy get probability zero in every MSNE (Theorem 6.7); and proved, in full, Nash's 1951 theorem that every finite game has an MSNE (Theorem 6.9), via a fixed-point construction using Brouwer's theorem.

References

[DGP09] C. Daskalakis, P. Goldberg and C. Papadimitriou, "The complexity of computing a Nash equilibrium," *SIAM Journal on Computing*, 39(1), 2009, pp. 195–259.

[MSZ13] M. Maschler, E. Solan and S. Zamir, *Game Theory*, Chapter 5, Cambridge University Press, 2013.

[Nas51] J. Nash, "Non-cooperative games," *Annals of Mathematics*, 54(2), 1951, pp. 286–295.