Lecture 13 - Foundations of Extensive Form Games

Iosif Sakos

November 6, 2025

Abstract

These notes cover extensive form games (EFGs), a fundamental framework in game theory for modeling sequential interactions with imperfect information. We define EFGs formally, discuss key concepts such as information sets and perfect recall, and explore strategies within this framework, including mixed and behavioral strategies. A central result, Kuhn's theorem, establishes the equivalence between mixed strategies and behavioral strategies in EFGs with perfect recall. This equivalence has significant implications for the existence of Nash equilibrium (NE) in such games.

Contents

1 Extensive Form Games

Extensive form games (EFGs) are games that capture the sequential nature of decisions, the timing of moves, and the information available to players at each decision point. They are typically represented with game trees, where nodes represent decision points, edges represent actions, and terminal nodes represent outcomes with associated payoffs.

1.1 Game Trees

A game tree is *a directed, rooted tree* that represents the sequential structure of a game. Each node in the tree corresponds to a decision point for a player, and each outgoing edge represents an action that a player can take at that node. Terminal nodes (leaves) represent the possible outcomes of the game, along with the payoffs for each player. For example, consider the following simple game tree shown in Figure 1, where Player 1 (P1) makes the first move, followed by Player 2 (P2).

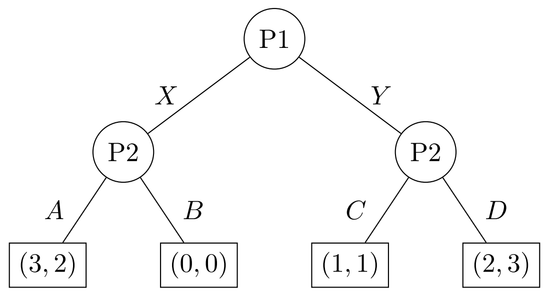

Figure 1: A simple game tree.

A game tree can also have *chance* nodes that represent stochastic events in the game, such as random card deals or dice rolls. These are typically drawn with outgoing edges labeled with outcome probabilities (labels may be omitted when the probabilties are equal). An example of a game tree with chance nodes is shown in Figure 2.

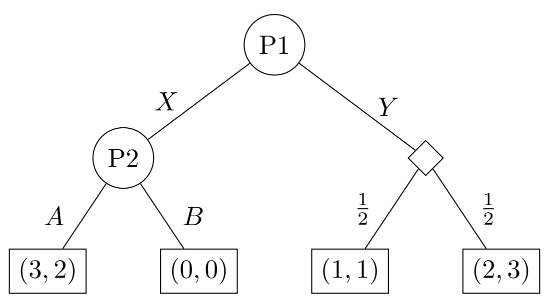

Figure 2: A game tree with chance nodes.

Imperfect information. Imperfect information can be represented in EFGs using information sets, which group together nodes that a player *cannot distinguish among* when making a decision. A player's move must be the same at all nodes within an information set, so nodes in the same set share the same available actions. This captures the idea that a player may not have complete knowledge of the history of play (including their own past actions) when making a decision. For example, if Player 2 cannot distinguish between two decision nodes after Player 1's move, those nodes are grouped into a single information set for Player 2, as depicted in Figure 3.

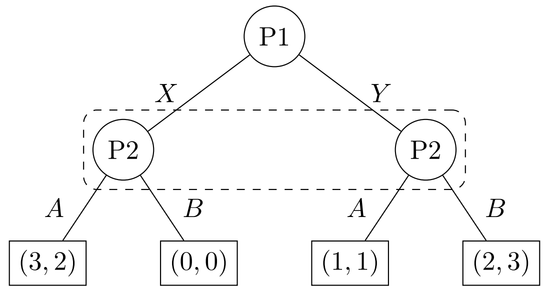

Figure 3: A game tree with an information set.

Kuhn poker. A classic example of an EFG with imperfect information is *Kuhn poker* [1], a simplified version of poker involving two players and a deck of three cards: King (K), Queen (Q), and Jack (J). Each player is dealt one card, and there is a single round of betting. The game tree for Kuhn poker is shown in Figure 4. This game already captures key elements of imperfect information since neither player knows the other's card when making betting decisions.

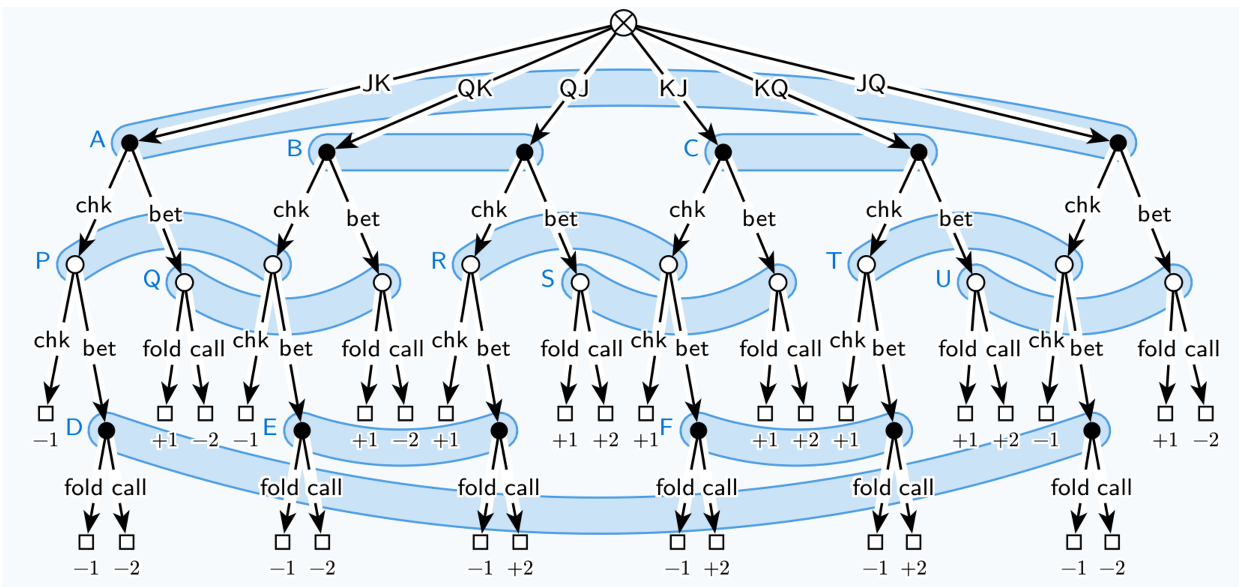

Figure 4: The game tree for Kuhn poker. Adapted from Gabriel Farina, MIT 6.S890—Topics in Multiagent Learning (lecture notes), Massachusetts Institute of Technology, October 3, 2024.

1.2 Histories, Actions, and Payoffs

The nodes of the game tree represent the *histories* of actions taken by players up to that point in the game. In an n-player EFG, each nonterminal history is associated with an *acting player* $i \in [\![n]\!]$ or with the nature player in the case of chance nodes. Players choose from the available actions (or, for the nature player, from the possible outcomes) at that history, and play advances to the next history according to the game tree structure. The set of all terminal nodes, or *terminal histories*, is denoted by $\mathcal{Z}$. Terminal histories are not associated with any player. Instead, each terminal history $z \in \mathcal{Z}$ is associated with a payoff vector $u(z) \in \mathbb{R}^n$ that assigns a real-valued payoff $u_i(z)$ to each player $i \in [\![n]\!]$.

To model imperfect information, the histories associated with each player $i \in [\![n]\!]$ are partitioned into a collection $\mathcal{J}_i$ of information sets. Each information set $\mathcal{I} \in \mathcal{J}_i$ is a set of histories that player i cannot distinguish among when making a decision. If all information sets are singletons, the game is said to have *perfect information*. In contrast, if any information set contains multiple histories, the game has *imperfect* *information*. All histories in the same information set $\mathcal{I}$ share the same set of available actions $\mathcal{A}_{\mathcal{I}}$ available to player i.

Simplifications. For simplicity, we assume that, for each player, the available actions at each information set are disjoint; that is, $\mathcal{A}_{\mathcal{I}} \cap \mathcal{A}_{\mathcal{I}'} = \emptyset$ for all $i \in [\![n]\!]$ and $\mathcal{I}, \mathcal{I}' \in \mathcal{J}_i$ with $\mathcal{I} \neq \mathcal{I}'$. Furthermore, we assume that there are no chance nodes in the game tree. All results in this section extend to EFGs with chance nodes by treating nature as an additional player with fixed distributions at its histories.

1.3 Perfect Recall

A standard assumption in EFGs is that players have perfect recall, meaning they remember all their previous actions and the information available to them when those actions were taken. The condition can be formalized as follows:

Definition 1 (Perfect Recall)**.** A player $i \in [\![n]\!]$ has perfect recall if, for any information set $\mathcal{I} \in \mathcal{J}_i$ and any two histories $h, h' \in \mathcal{I}$, the sequences of actions taken by player i that lead to h and h' are identical. An EFG has perfect recall if all players have perfect recall.

Without the assumption of perfect recall, solving EFGs can be *intractable* [2].

1.4 Strategies in Extensive Form Games

1.4.1 Normal-Form Representation

The *normal-form representation* of an EFG is obtained by expressing the game in terms of players' *strategies* and *payoffs*, rather than the game tree structure. As the decision nodes within any information set are indistinguishable to the player, it suffices to define a strategy as a sequence of actions that the player takes at each information set; i.e., the set of *pure strategies* of player $i \in [\![n]\!]$ is given by

$$\mathcal{S}_i = \prod_{\mathcal{I} \in \mathcal{J}_i} \mathcal{A}_{\mathcal{I}}. \tag{1}$$

Accordingly in an EFG we define the set of *mixed strategies* for player i via the normal-form representation; specifically, it is the probability simplex $\Delta(\mathcal{S}_i)$ over the set of pure strategies $\mathcal{S}_i$.

Given a terminal history $z \in \mathcal{Z}$, let $\left(\alpha_{i,1}, \ldots, \alpha_{i,m_i}\right)$ denote the sequence of actions taken by player i to reach z, where m_i is the length of the action sequence. Let $\mathcal{I}_{i,j}$ denote the information set at which action $\alpha_{i,j}$ was taken. Then the probability of reaching terminal history z under a mixed-strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is

$$\begin{aligned}\Pr_{\sigma}(z) &= \sum_{s_1 \in \mathcal{S}_1} \cdots \sum_{s_n \in \mathcal{S}_n} \prod_{i=1}^{n} \sigma_i(s_i) \cdot \mathbf{1}\left\{s_{i,\mathcal{I}_{i,j}} = \alpha_{i,j} \, \forall j \in [\![m_i]\!]\right\} \\ &= \prod_{i=1}^{n} \sum_{s_i \in \mathcal{S}_i} \sigma_i(s_i) \cdot \mathbf{1}\left\{s_{i,\mathcal{I}_{i,j}} = \alpha_{i,j} \, \forall j \in [\![m_i]\!]\right\}.\end{aligned} \tag{2}$$

The last equality follows from the independence of players' strategies and the per-player factorization of the indicator. It follows from Nash's theorem that a mixed-strategy Nash equilibrium (NE) always exists in EFGs.

As an example, consider the EFG shown in Figure 3. Player 1 has two possible actions at the root node (X or Y), and Player 2 has two possible actions (A or B) at each of the two decision nodes. Thus, Player 1's set of pure strategies is $\{X, Y\}$. Furthermore, since both decision nodes of Player 2 belong to the same information set, Player 2's set of pure strategies is $\{A, B\}$. Then the normal-form representation of the game can be expressed as the payoff matrix:

$$\begin{array}{c|cc} & A & B \\ \hline X & (3,2) & (0,0) \\ Y & (1,1) & (2,3) \end{array} \tag{3}$$

As another example, consider the EFG shown in Figure 1. Again, Player 1 has two possible actions at the root node (X or Y). However, Player 2 has two information sets, each with two possible actions (A or B in the first set, and C or D in the second). Thus, Player 1's set of pure strategies is still $\{X, Y\}$, but Player 2's set of pure strategies is now $\{AC, AD, BC, BD\}$, where, for example AC denotes choosing action A in the first information set and action C in the second. Then the normal-form representation of the game can be expressed as thepayoff matrix:

$$\begin{array}{c|cccc} & AC & AD & BC & BD \\ \hline X & (3,2) & (3,2) & (0,0) & (0,0) \\ Y & (1,1) & (2,3) & (1,1) & (2,3) \end{array} \tag{4}$$

Observe that the size of the normal-form representation grows exponentially in the number of information sets and available actions, making it impractical for large EFGs.

1.4.2 Behavioral-Form Representation

In contrast to the normal-form representation, the *behavioral-form representation* of an EFG focuses on players' strategies at each decision point in the game tree. A behavioral strategy for a player specifies a probability distribution over the available actions at each of the player's information sets. Thus the set of behavioral strategies for player $i \in [\![n]\!]$ is given by

$$\mathcal{B}_i = \prod_{\mathcal{I} \in \mathcal{J}_i} \Delta(\mathcal{A}_{\mathcal{I}}), \tag{5}$$

where $\Delta(\mathcal{A}_{\mathcal{I}})$ denotes the probability simplex over the action set $\mathcal{A}_{\mathcal{I}}$. Behavioral strategies are particularly useful in EFGs with imperfect information, as they allow players to randomize at each decision point based on the information available to them.

For example, in the EFG shown in Figure 3, a behavioral strategy for Player 2 could specify choosing A with probability 0.7 and B with probability 0.3 at each history in their information set. The player's behavioral strategy is then characterized by the distribution at that information set, i.e., $\big((0.7, 0.3)\big)$. Similarly, in the EFG shown in Figure 1, a behavioral strategy for Player 2 could specify choosing A with probability 0.6 and B with probability 0.4 in the first information set, and choosing C with probability 0.5 and D with probability 0.5 in the second. Then the behavioral strategy of Player 2 is $\big((0.6, 0.4), (0.5, 0.5)\big)$.

Behavioral Nash equilibrium. Given a terminal history $z \in \mathcal{Z}$, let $\big(\alpha_{i,1}, \ldots, \alpha_{i,m_i}\big)$ denote the sequence of actions taken by player i to reach z, where m_i is the length of the action sequence. Then the probability of reaching terminal history z under a profile of behavioral strategies $b = (b_1, \ldots, b_n)$ is given by

$$\Pr_b(z) = \prod_{i=1}^{n} \prod_{j=1}^{m_i} b_i\big(\alpha_{i,j}\big). \tag{6}$$

Here we write $b_i(\alpha)$ unambiguously for any action α of player i (the player's information sets are disjoint, so action labels do not collide).

The expected payoff of player $i \in [\![n]\!]$ under a profile of behavioral strategies $b = (b_1, \ldots, b_n)$ is

$$u_i(b) = \mathop{\mathbb{E}}_b[u_i(z)] = \sum_{z \in \mathcal{Z}} \Pr_b(z) u_i(z). \tag{7}$$

A behavioral NE $b^\star$ is a profile of behavioral strategies in which no player can improve their expected payoff by unilaterally changing their behavioral strategy; i.e.,

$$u_i(b^\star) \geq u_i(b_i, b_{-i}^\star), \quad \forall i \in [\![n]\!],\ b_i \in \mathcal{B}_i. \tag{8}$$

In general, a behavioral NE might not exist in EFGs. However, under the assumption of perfect recall, we have the following important result known as Kuhn's theorem.

1.4.3 Kuhn's Theorem

Theorem 2 (Kuhn's Theorem [3]). *In any EFG with perfect recall, for any mixed-strategy profile* $\sigma = (\sigma_1, \ldots, \sigma_n)$ *there exists a profile of behavioral strategies* $b = (b_1, \ldots, b_n)$ *such that* $\Pr_\sigma(z) = \Pr_b(z)$ *for all* $z \in \mathcal{Z}$*. Conversely, for any profile of behavioral strategies* $b = (b_1, \ldots, b_n)$ *there exists a mixed-strategy profile* $\sigma = (\sigma_1, \ldots, \sigma_n)$ *such that* $\Pr_\sigma(z) = \Pr_b(z)$ *for all* $z \in \mathcal{Z}$*.*

Proof. For each player $i \in [\![n]\!]$, since player i has perfect recall, for each information set $\mathcal{I} \in \mathcal{J}_i$ there exists a unique sequence of actions taken by player i to reach any history in $\mathcal{I}$. Consequently, for each action $\alpha \in \mathcal{A}_\mathcal{I}$, there exists a unique sequence of actions taken by player i that culminates in choosing α. Let $q_\alpha = (q_{\alpha,1}, \ldots, q_{\alpha,m_\alpha})$ denote this unique sequence, where m_α is the length of the action sequence. Let $\mathcal{I}_{\alpha,j}$ denote the information set at which action $q_{\alpha,j}$ was taken.

For each player $i \in [\![n]\!]$ and each mixed strategy $\sigma_i \in \Delta(\mathcal{S}_i)$, define

$$r_i^{\sigma_i}(\alpha) = \sum_{s_i \in \mathcal{S}_i} \sigma_i(s_i) \cdot \mathbf{1}\big\{s_i(\mathcal{I}_{\alpha,j}) = q_{\alpha,j} \; \forall j \in [\![m_\alpha]\!]\big\}, \quad \forall \alpha \in \mathcal{A}_\mathcal{I},\ \mathcal{I} \in \mathcal{J}_i. \tag{9}$$

Moreover, for each $b_i \in \mathcal{B}_i$, define

$$r_i^{b_i}(\alpha) = \prod_{j=1}^{m_\alpha} b_i(q_{\alpha,j}), \quad \forall \alpha \in \mathcal{A}_\mathcal{I},\ \mathcal{I} \in \mathcal{J}_i. \tag{10}$$

Then, by Equations (2) and (6), to show that $\Pr_\sigma(z) = \Pr_b(z)$ for all $z \in \mathcal{Z}$, it suffices to show that $r_i^{\sigma_i}(\alpha) = r_i^{b_i}(\alpha)$ for all $i \in [\![n]\!]$, $\mathcal{I} \in \mathcal{J}_i$, and $\alpha \in \mathcal{A}_\mathcal{I}$.

For the forward direction, let $\sigma_i \in \Delta(\mathcal{S}_i)$ be any mixed strategy for player $i \in [\![n]\!]$. Define the corresponding behavioral strategy $b_i \in \mathcal{B}_i$ by

$$b_i(\alpha) = \begin{cases} \frac{r_i^{\sigma_i}(\alpha)}{r_i^{\sigma_i}(q_{\alpha,m_\alpha-1})}, & \text{if } m_\alpha > 1 \text{ and } r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) > 0, \\ r_i^{\sigma_i}(\alpha), & \text{if } m_\alpha = 1, \\ \frac{1}{|\mathcal{A}_\mathcal{I}|}, & \text{otherwise}, \end{cases} \tag{11}$$

for all $\alpha \in \mathcal{A}_\mathcal{I}$ and $\mathcal{I} \in \mathcal{J}_i$. Note that $b_i(\alpha)$ is well defined: by (9), $b_i(\cdot)$ is a valid probability distribution over $\mathcal{A}_\mathcal{I}$ for each $\mathcal{I} \in \mathcal{J}_i$.

Then by induction on m_α, it can be shown that $r_i^{\sigma_i}(\alpha) = r_i^{b_i}(\alpha)$ for all $\mathcal{I} \in \mathcal{J}_i$ and $\alpha \in \mathcal{A}_\mathcal{I}$. Indeed, for the base case $m_\alpha = 1$, we have $r_i^{b_i}(\alpha) = b_i(\alpha) = r_i^{\sigma_i}(\alpha)$. For the inductive step, suppose $\alpha \in \mathcal{A}_\mathcal{I}$ with $m_\alpha > 1$. Then, we have that $r_i^{b_i}(\alpha) \overset{(10)}{=} b_i(\alpha) \cdot r_i^{b_i}(q_{\alpha,m_\alpha-1}) = b_i(\alpha) \cdot r_i^{\sigma_i}(q_{\alpha,m_\alpha-1})$ (by the induction hypothesis). Thus, if $r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) = 0$, then $r_i^{b_i}(\alpha) = 0$. Moreover, by (9), it follows that $r_i^{\sigma_i}(\alpha) \leq r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) = 0$;therefore $r_i^{\sigma_i}(\alpha) = r_i^{b_i}(\alpha) = 0$.

On the other hand, if $r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) > 0$, then

$$r_i^{b_i}(\alpha) = b_i(\alpha) \cdot r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) \overset{(11)}{=} \frac{r_i^{\sigma_i}(\alpha)}{r_i^{\sigma_i}(q_{\alpha,m_\alpha-1})} \cdot r_i^{\sigma_i}(q_{\alpha,m_\alpha-1}) = r_i^{\sigma_i}(\alpha). \tag{12}$$

Thus we have shown that, for any mixed-strategy profile σ, there exists a profile of behavioral strategies b that induces the same distribution over terminal histories; i.e., $\Pr_\sigma(z) = \Pr_b(z)$ for all $z \in \mathcal{Z}$.

For the converse, let $b_i \in \mathcal{B}_i$ be any behavioral strategy for player $i \in [\![n]\!]$. Define the corresponding mixed strategy $\sigma_i \in \Delta(\mathcal{S}_i)$ by

$$\sigma_i(s_i) = \prod_{\mathcal{I} \in \mathcal{J}_i} b_i(s_{i,\mathcal{I}}), \quad \forall s_i \in \mathcal{S}_i. \tag{13}$$

Observe that $\sigma_i(\cdot)$ is a well-defined product distribution obtained by sampling independently at each information set.

Let $\mathcal{J}_\alpha = \{\mathcal{I}_{\alpha,1}, \dots, \mathcal{I}_{\alpha,m_\alpha}\}$. Then

$$\begin{aligned}
r_i^{\sigma_i}(\alpha) &= \sum_{s_i \in \mathcal{S}_i} \left(\prod_{\mathcal{I} \in \mathcal{J}_i} b_i(s_{i,\mathcal{I}}) \right) \cdot \mathbf{1}\big\{ s_i(\mathcal{I}_{\alpha,j}) = q_{\alpha,j} \, \forall j \in [\![m_\alpha]\!] \big\} && \text{(14a)} \\
&= \prod_{j=1}^{m_\alpha} b_i(q_{\alpha,j}) \sum_{(\beta_\mathcal{I})_{\mathcal{I} \in \mathcal{J}_i \setminus \mathcal{J}_\alpha}} \prod_{\mathcal{I} \in \mathcal{J}_i \setminus \mathcal{J}_\alpha} b_i(\beta_\mathcal{I}) && \text{(14b)} \\
&= \prod_{j=1}^{m_\alpha} b_i(q_{\alpha,j}) \prod_{\mathcal{I} \in \mathcal{J}_i \setminus \mathcal{J}_\alpha} \underbrace{\sum_{\beta \in \mathcal{A}_\mathcal{I}} b_i(\beta)}_{=1} && \text{(14c)} \\
&\overset{(10)}{=} r_i^{b_i}(\alpha). && \text{(14d)}
\end{aligned}$$

Thus we have shown that, for any profile of behavioral strategies b, there exists a mixed-strategy profile σ that induces the same distribution over terminal histories; i.e., $\Pr_\sigma(z) = \Pr_b(z)$ for all $z \in \mathcal{Z}$. $\square$

Kuhn's theorem establishes the equivalence between mixed strategies and behavioral strategies in EFGs with perfect recall. It follows that NE exist in EFGs with perfect recall, since mixed-strategy NE exist by Nash's theorem.

A List of Abbreviations

EFG **E**xtensive **F**orm **G**ame 1, 3–8

NE **N**ash **E**quilibrium 1, 5, 6, 8

B Index

References

[1] Harold William Kuhn. "A Simplified Two-Person Poker." In: *Contributions to the Theory of Games*. Ed. by Harold William Kuhn and Albert William Tucker. Vol. 1. 2 vols. Annals of Mathematics Studies 24. Princeton, New Jersey, United States: Princeton University Press, 1950, pp. 97–104.

[2] Daphne Koller and Nimrod Megiddo. "The complexity of two-person zero-sum games in extensive form." In: *Games and Economic Behavior* 4.4 (1992), pp. 528–552. ISSN: 0899-8256. DOI: 10.1016/0899-8256(92)90035-Q.

[3] Harold William Kuhn. "Extensive Games and the Problem of Information." In: *Contributions to the Theory of Games*. Ed. by Harold William Kuhn and Albert William Tucker. Vol. 2. 2 vols. Annals of Mathematics Studies 24. Princeton, New Jersey, United States: Princeton University Press, 1953, pp. 193–216.

[4] Harold William Kuhn and Albert William Tucker, eds. *Contributions to the Theory of Games*. 2 vols. Annals of Mathematics Studies 24. Princeton, New Jersey, United States: Princeton University Press, 1950.