Chapter 13

Existence of Equilibrium and Refinements

This chapter supplies the rigorous core of non-cooperative theory: why equilibria exist, and how to discipline the multiplicity of equilibria in dynamic games.

13.1 Fixed-point theorems

Theorem 13.1: Brouwer's fixed-point theorem

Every continuous function $f : K \to K$ from a nonempty compact convex set $K \subset \mathbb{R}^m$ to itself has a fixed point: $f(x^*) = x^*$.

Theorem 13.2: Kakutani's fixed-point theorem (1941)

Let $K \subset \mathbb{R}^m$ be nonempty, compact, convex, and let $F : K \rightrightarrows K$ be a correspondence such that (i) $F(x)$ is nonempty and convex for every x, and (ii) F has a closed graph. Then F has a fixed point: $x^* \in F(x^*)$.

13.2 Nash's theorem

Theorem 13.3: Nash (1950)

Every finite game $G = \langle N, (A_i), (u_i) \rangle$ has at least one Nash equilibrium in mixed strategies.

Proof. Let $\Sigma = \prod_i \Delta(A_i)$, a nonempty compact convex subset of a Euclidean space. Define the best-response correspondence $B(\sigma) = \prod_i \mathrm{BR}_i(\sigma_{-i})$, where $\mathrm{BR}_i(\sigma_{-i}) = \arg\max_{\tau_i \in \Delta(A_i)} u_i(\tau_i, \sigma_{-i})$. Each $\mathrm{BR}_i(\sigma_{-i})$ is nonempty (a continuous function is maximized on a compact set) and convex (the maximizers of a linear function on a simplex form a face, which is convex). The graph of B is closed because u_i is continuous in σ (Berge's maximum theorem, or a direct limit argument). By Kakutani's theorem, B has a fixed point σ^*, which by construction satisfies $\sigma_i^* \in \mathrm{BR}_i(\sigma_{-i}^*)$ for all i: a Nash equilibrium. □

Remark: Nash's second proof via Brouwer

Nash (1951) gave a self-contained proof using Brouwer directly: define the "gain" $g_{ia}(\sigma) = \max\{0,\, u_i(a, \sigma_{-i}) - u_i(\sigma)\}$ and the continuous map

$$T_{ia}(\sigma) = \frac{\sigma_i(a) + g_{ia}(\sigma)}{1 + \sum_{a'} g_{ia'}(\sigma)}.$$

Fixed points of T are exactly the Nash equilibria (if σ is not an equilibrium, probability flows toward better replies and σ moves). This map is also the starting point for computational approaches—and for understanding why computing the fixed point is hard (Chapter 17).

Theorem 13.4: Games with continuous action spaces (Debreu–Glicksberg–Fan 1952)

If each A_i is a nonempty compact convex subset of $\mathbb{R}^m$, and each u_i is continuous in a and *quasi-concave* in own action a_i, then a *pure-strategy* Nash equilibrium exists.

This covers Cournot, Bertrand with differentiated goods, public-good games, etc. Glicksberg (1952) further extends Nash's theorem to continuous games (compact metric action spaces, continuous payoffs) with mixed strategies. Discontinuous games (all-pay auctions, Bertrand) need more delicate arguments (Reny's better-reply security, 1999).

13.3 Trembling-hand perfection

Nash equilibrium can rely on weakly dominated actions and incredible behavior at unreached information sets. Selten's idea: robustness to small mistakes.

Definition 13.1: Trembling-hand perfect equilibrium (Selten 1975)

A mixed profile σ^* of a finite normal-form game is *(trembling-hand) perfect* if there is a sequence of *totally mixed* profiles $\sigma^k \to \sigma^*$ (every action played with positive probability) such that each σ_i^* is a best response to σ_{-i}^k for every k.

Proposition 13.5: Properties of perfection

Every finite game has a perfect equilibrium. Every perfect equilibrium is Nash, and assigns zero probability to weakly dominated actions. In two-player games, an equilibrium is perfect iff both strategies are undominated.

Example 13.1: Ruling out dominated equilibria

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$0, 0$

Both (T, L) and (B, R) are Nash equilibria, but (B, R) uses weakly dominated actions on both sides: any tremble by the opponent makes T (resp. L) strictly better. Only (T, L) is perfect.

13.4 Sequential equilibrium

For extensive-form games, the workhorse refinement pairs strategies with beliefs, like PBE, but disciplines off-path beliefs by a limiting consistency requirement.

Definition 13.2: Sequential equilibrium (Kreps–Wilson 1982)

An *assessment* (σ, μ)—a behavior strategy profile plus a belief system —is a *sequential equilibrium* if:

1. **Sequential rationality**: at every information set, the mover's continuation strategy is optimal given μ and σ;
2. **Consistency**: $(\sigma, \mu) = \lim_k (\sigma^k, \mu^k)$ for some sequence of totally mixed σ^k with μ^k derived from σ^k by Bayes' rule everywhere.

Consistency says off-path beliefs must be explicable as limits of small correlated trembles—ruling out arbitrary belief threats that plain PBE permits. Every finite extensive-form game with perfect recall has a sequential equilibrium; every sequential equilibrium is subgame perfect; trembling-hand perfection in the *agent normal form* implies sequentiality.

Remark: The refinements hierarchy

For finite extensive-form games with perfect recall:

$$\text{THP (agent form)} \subseteq \text{sequential} \subseteq \text{PBE} \subseteq \text{subgame perfect} \subseteq \text{Nash}.$$

Further refinements—properness (Myerson), stability (Kohlberg–Mertens), forward induction, the intuitive criterion (Theorem 11.1)—address signaling and forward-looking inferences. The refinements program of the 1980s never produced a single universally accepted concept, but its ideas (credibility, consistency of beliefs, forward induction) permanently shaped applied modeling.

13.5 Epistemic foundations

What assumptions on knowledge and belief justify each solution concept? A sample of landmark results from *epistemic game theory*:

- **Rationalizability**: common knowledge of rationality characterizes (correlated) rationalizability, i.e. IESDS survivors (Theorem 3.2).
- **Nash equilibrium** (Aumann–Brandenburger 1995): in two-player games, mutual knowledge of rationality and of the players' *conjectures* about each other implies the conjectures form a Nash equilibrium—common knowledge is not needed, but knowing the opponent's belief is a strong requirement.
- **Correlated equilibrium** (Aumann 1987): if players share a common prior and are "Bayes rational" given their information, their behavior is a correlated equilibrium. This gives correlated equilibrium a cleaner epistemic pedigree than Nash.
- **Agreeing to disagree** (Aumann 1976): with a common prior, if two players' posteriors about an event are common knowledge, they must be *equal*—rational agents cannot agree to disagree. Corollary (no-trade theorems, Milgrom–Stokey 1982): rational risk-

averse traders with a common prior cannot trade on differences of information alone—if someone wants to bet with you, that itself is information (compare the winner's curse, Example 10.5).

Exercises

Exercise 13.1: Existence via Kakutani, hands-on

Consider a two-player game where each player picks $x_i \in [0,1]$ and $u_i = -(x_i - \frac{1}{2}x_j)^2$. Compute best responses, verify the conditions of Theorem 13.4, and find the unique pure equilibrium.

Exercise 13.2: Perfection in the normal form

	L	C	R
T	$2,2$	$1,1$	$0,0$
M	$1,1$	$1,1$	$1,1$
B	$0,0$	$1,1$	$2,2$

Find all pure Nash equilibria and determine which are trembling-hand perfect.

Exercise 13.3: Sequential vs. PBE

Construct (or recall from the literature) a three-player extensive-form game with a PBE whose off-path beliefs cannot arise as limits of trembles, so the assessment is not sequential. Explain exactly where consistency bites.

Exercise 13.4: No-trade

Two risk-neutral agents with a common prior each receive a private signal about an asset's value and are offered the chance to bet against each other; betting is voluntary. Prove that in any equilibrium, the probability of a bet actually occurring is zero. Where exactly does the common prior enter, and what happens without it?