Best-Response Equilibrium: An Equilibrium in Finitely Additive Mixed Strategies

Igal Milchtaich

Bar-Ilan University

Wish-granting game

- The payoff is any real number s the player chooses, $u(s) = s$
- Any choice s^1 is inferior to any $s^2 > s^1$, which is inferior to s^3…
- Does the whole sequence $(s^n)_{n\in\mathbb{N}}$ represent a strategy?
- Does $\lim_{n\to\infty} u(s^n) = \infty$ make it an equilibrium strategy?
- A mixed strategy σ involves assignment of probabilities
- An equilibrium should satisfy $\sigma(\{s^n\}) = 0$ for all n
- Additivity then implies $\sigma(\{s^1, s^2, \dots, s^n\}) = 0$
- Hence, sigma-additivity cannot hold
- Strategy σ is defined as a finitely-additive probability
- For $A \subseteq \mathbb{R}$,
  - $\sigma(A) = 0$ if $s^n \in A^\mathrm{C}$ for almost all n
  - $\sigma(A) = 1$ if $s^n \in A$ for almost all n
- Can be extended to the entire power set of $\mathbb{R}$

Wish-granting game

- Strategy σ describes a rational choice of action
- No single action is optimal, as $\sup_{s \in \mathbb{R}} u(s) = \infty$
- Strategy σ excludes the choice of actions yielding low payoffs, no matter how 'low payoff' is understood
- For every $a < \sup u = \lim_{n\to\infty} u(s^n)$, only finitely many n's satisfy $u(s^n) < a$, and so
$$\sigma(\{ s \in \mathbb{R} \mid u(s) < a \}) = 0$$
- Strategy σ is a best-response equilibrium
- A similar construction works for any one-player game
- Applicable to any action set and payoff function
- Here, specifically, σ formalizes the choice of "infinity": strategy δ_∞
- For every set A bounded from above, $\delta_\infty(A) = 0$ and $\delta_\infty\left(A^C\right) = 1$

Finitely additive probabilities

- The power set $\mathcal{P}(S)$ of a set S is the collection of all its subsets
- $\{\emptyset\} \subseteq \mathcal{A} \subseteq \mathcal{P}(S)$ is an algebra if $A, B \in \mathcal{A}$ implies $A^C, A \cup B \in \mathcal{A}$
- Its elements are the measurable sets
- A finitely additive probability is a function $\mu: \mathcal{A} \longrightarrow [0,1]$ satisfying
  - $\mu(A) + \mu(B) = \mu(A \cup B)$ for all disjoint $A, B \in \mathcal{A}$
  - $\mu(S) = 1$
- It is a probability if for all disjoint $A_1, A_2, \ldots \in \mathcal{A}$ with $\bigcup_{k=1}^{\infty} A_k \in \mathcal{A}$
$$\sum_{k=1}^{\infty} \mu(A_k) = \mu\left(\bigcup_{k=1}^{\infty} A_k\right)$$
- A finitely additive probability $\mu': \mathcal{A}' \longrightarrow [0,1]$ is an extension of μ if $\mathcal{A} \subseteq \mathcal{A}'$ and $\mu = \mu'|_{\mathcal{A}}$, and it is a total extension if $\mathcal{A}' = \mathcal{P}(S)$
- The outer measure of μ is the function $\mu^*: \mathcal{P}(S) \longrightarrow [0,1]$ defined by
$$\mu^*(C) = \inf \{\mu(A) \mid A \supseteq C, A \in \mathcal{A}\}$$
- A set C with $\mu^*(C) = 0$ is μ-null

Integration

- A simple measurable function $f\colon S \longrightarrow \mathbb{R}$ takes only finitely many values and satisfies $f^{-1}(\{x\}) \in \mathcal{A}$ for every value x
- The integral of f with respect to a finitely additive probability μ is
$$\int_S f(s)\, d\mu(s) = \sum_{x\in\mathbb{R}} x\, \mu\big(f^{-1}(\{x\})\big)$$
- More generally, $f\colon S \longrightarrow \mathbb{R}$ is μ-integrable if there are simple measurable functions $(f_n)_{n\in\mathbb{N}}$ such that for every $\epsilon > 0$
$$\lim_{n\to\infty} \mu^*(\{s \in S \mid |f(s) - f_n(s)| > \epsilon\}) = 0,$$
$$\lim_{m,n\to\infty} \int_S |f_m(s) - f_n(s)|\, d\mu(s) = 0$$
- The integral of such f is (well) defined by
$$\int_S f(s)\, d\mu(s) = \lim_{n\to\infty} \int_S f_n(s)\, d\mu(s)$$

Product of finitely additive probabilities

- $(\mu_i)_{i=1}^n$ defined on algebras $(\mathcal{A}_i)_{i=1}^n$ of subsets of sets $(S_i)_{i=1}^n$
- The product algebra $\mathcal{A} = \prod_i \mathcal{A}_i$ consists of all finite unions of sets $A \subseteq S = \prod_i S_i$ of the form $A = \prod_i A_i$ with $A_i \in \mathcal{A}_i$ for all i
- The product $\mu = \prod_i \mu_i$ is a finitely additive probability defined on $\mathcal{A}$
- For a rectangular set A as above, $\mu(A) = \prod_i \mu_i(A_i)$

Lemma. For a bounded function $f: S \longrightarrow \mathbb{R}$,

$$\int_S f(s)\, d\mu(s) = \int_{S_n} \cdots \int_{S_1} f(s_1, s_2, \ldots, s_n)\, d\mu_1(s_1) \cdots d\mu_n(s_n),$$

provided that the “multiple” and iterated integral both exist.

- In particular, the latter does not depend on the order of integration
- However, Fubini’s theorem does not hold here
- It is possible that only the multiple or only the iterated integral exists

Best-response equilibrium

- Each player i has an action set S_i and a payoff function $u_i: S \longrightarrow \mathbb{R}$
- A strategy for i is any finitely additive probability $\sigma_i: \mathcal{A}_i \longrightarrow [0,1]$
- A strategy profile $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ may be identified with $\sigma = \prod_i \sigma_i$
- For any i, it may also be written as (σ_i, σ_{-i}), where $\sigma_{-i} = \prod_{j \neq i} \sigma_j$
- A strategy profile σ is a best-response equilibrium if for every i
  - the following integral exists for every $s_i \in S_i$
$$v_i(s_i) \coloneqq \int_{S_{-i}} u_i(s_i, s_{-i})\, d\sigma_{-i}(s_{-i})$$
  - the function $v_i: S_i \longrightarrow \mathbb{R}$ satisfies for every $a < \sup_{s_i \in S_i} v_i(s_i)$
$$\sigma_i^*(\{\, s_i \in S_i \mid v_i(s_i) < a \,\}) = 0$$
- Thus, actions yield well-defined expected payoffs, and any set of low-payoff actions is σ_i-null (the best-response requirement)

Best-response equilibrium

Proposition 1. If $\sup v_i < \infty$, the best-response requirement holds if and only if v_i is σ_i-integrable and

$$\int_{S_i} v_i(s_i)\, d\sigma_i(s_i) = \sup v_i\,.$$

- Player i's equilibrium payoff is $\int_S u_i(s)\, d\sigma(s)$ – if the integral exists
- If u_i is not σ-integrable, the equilibrium payoff is not well defined
- A best-response equilibrium excludes the choice of low-payoff actions, without necessarily identifying expected payoffs

Proposition 2. Every strategy profile $\tilde{\sigma}$ that extends a best-response equilibrium σ is also a best-response equilibrium.
At least one such $\tilde{\sigma}$ is total (in the sense that $\mathcal{A}_i = \mathcal{P}(S_i)$ for all i).

Bilateral trade

- An item's worth is 0 to the seller and 1 to the buyer
- The buyer has to offer a price $0 \leq p \leq 1$
- The seller has to select the interval of acceptable prices
- Accepting any $p > 0$ is a weakly dominant strategy
- But there is no mixed equilibrium of which it is a part
- Intuitively, the buyer should offer "very little", or "an ϵ"
- A best-response equilibrium does exist: the seller's strategy is δ_{0^+}
- For $A \subseteq [0,1]$ that includes a right neighborhood of 0, $\delta_{0^+}(A) = 1$
- The equilibrium payoffs are 1 to the buyer and 0 to the seller

Price competition

- The n identical firms with cost function C set prices $p_1, p_2, \dots, p_n$
- Those tied for the lowest price p equally share the demand $D(p)$
- Competition may be expected to drive the price down
- A ("normal") mixed equilibrium may not exist, even for $n = 2$
- Example: $D(p) = 1 - p$ and quasi-fixed cost $C(q) = 0.16 \cdot 1_{q>0}$
- For a monopoly, $p = 0.5$ is profit maximizing, 0.2 gives zero profit
- For any $0.2 \le p \le 0.5$, $\left(\delta_{p^-}, \delta_{p^-}\right)$ is a best-response equilibrium
- For $A \subseteq [0, \infty)$ that includes a left neighborhood of p, $\delta_{p^-}(A) = 1$
- No well-defined equilibrium profits
- More generally, $\left(\delta_{p^-}, \delta_{p^-}, \dots, \delta_{p^-}\right)$ is a best-response equilibrium if
  - $\pi_M(p) = pD(p) - C(D(p))$ is nondecreasing in $(0, p)$, and
  - its supremum there is nonnegative

Spatial competition with three firms

- Uniformly-distributed consumers on $[0,1]$ choose the closest firm
- A firm's profit is the total mass of its consumers
- With three firms, no pure strategy equilibrium exists
- Symmetric equilibrium with uniform distribution on $[1/4,3/4]$
- Unique equilibrium with a mixture of pure and mixed strategies
- One firm at $1/2$, the other two mix with support $[5/24,19/24]$
- Cannot be replaced by any two-point randomization
- Can be replaced by $1/2\,\delta_{x^-} + 1/2\,\delta_{(1-x)^+}$, with $1/4 \leq x \leq 1/3$
- The replacement gives a best-response equilibrium
- Only the player choosing $1/2$ has a well-defined equilibrium payoff

Zero-sum game without a value

- Two-player zero-sum game (Sion and Wolfe 1957)
- Both players' action set is $[0,1]$, and u_1 is
- Maxmim value is $1/3$, and maxmin strategy
$$\sigma_1 = 1/3\,\delta_0 + 2/3\,\delta_1$$
- Minmax value is $3/7$, and minmax strategy
$$\sigma_2 = 1/7\,\delta_{1/4} + 2/7\,\delta_{1/2} + 4/7\,\delta_1$$
- There exists no "normal" mixed equilibrium
- But there is a strategy of player 2 lowering 1's maximum payoff to $1/3$ (Vasquez 2017)
$$\sigma_2' = 1/3\,\delta_{1/2^-} + 2/3\,\delta_1$$
- A best-response equilibrium (σ_1, σ_2')
- Well-defined equilibrium payoffs $(1/3, -1/3)$

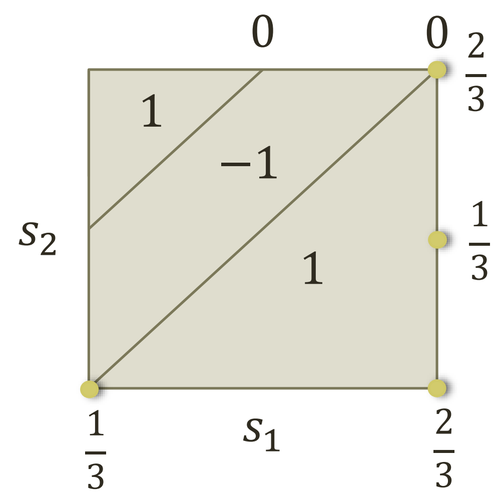

Game without best-response equilibrium

- Three players have the same action set $(0,1)$
- Payoff functions $u_1(s) = -s_1, u_2(s) = -s_2, u_3(s) = \min(s_2/s_1, 1)$
- Best-response requirement for $i = 1,2$

$$\int (-s_i)\, d\sigma_i(s_i) = 0$$

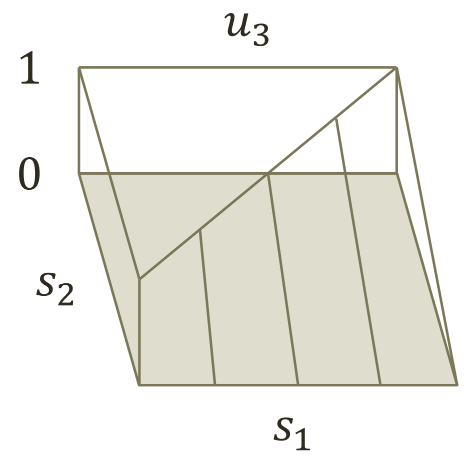

- But $v_3 = \int u_3\, d(\sigma_1 \times \sigma_2)$ does not exist
- The iterated integrals are not equal

$$\int \int \min\left(\frac{s_2}{s_1}, 1\right) d\sigma_2(s_2)\, d\sigma_1(s_1) \leq \int \left(-\frac{1}{s_1}\right) \int (-s_2)\, d\sigma_2(s_2)\, d\sigma_1 = 0$$

$$\int \int \min\left(\frac{s_2}{s_1}, 1\right) d\sigma_1(s_1)\, d\sigma_2(s_2) \geq \int \int \left(1 - \frac{s_1}{s_2}\right) d\sigma_1(s_1)\, d\sigma_2(s_2) = 1$$

Two-player counterexample

- Two players have the same action set $\mathbb{N}$, and the payoff matrix is

$$
\begin{array}{c c}
 & \begin{array}{cccccc} \mathbf{1} & \mathbf{2} & \mathbf{3} & \cdots & \boldsymbol{n} & \cdots \end{array} \\
\begin{array}{c} \mathbf{1} \\ \mathbf{2} \\ \mathbf{3} \\ \vdots \\ \boldsymbol{n} \\ \vdots \end{array} &
\begin{pmatrix}
1,1 & 1,2 & 1,3 & \cdots & 1,n & \cdots \\
0,1 & 2,2 & 2,3 & \cdots & 2,n & \cdots \\
0,1 & 0,0 & 3,3 & \cdots & 3,n & \cdots \\
\vdots & \vdots & & \ddots & \vdots & \\
0,1 & 0,0 & 0,0 & \cdots & n,n & \\
\vdots & \vdots & \vdots & & & \ddots
\end{pmatrix}
\end{array}
$$

- Strategy σ_1 is "diffuse" ($\sigma_1(\{n\}) = 0$ for all n) $\Longrightarrow \sigma_2(\{1\}) = 1$
- But $\sigma_2(\{1\}) = 1 \Longrightarrow \sigma_1(\{1\}) = 1 \Longrightarrow \sigma_1$ is not diffuse
- Strategy σ_1 is not diffuse $\Longrightarrow \lim_{n\to\infty} v_2(n) = \infty \Longrightarrow \sigma_2$ is diffuse
- But σ_2 is diffuse $\Longrightarrow \lim_{n\to\infty} v_1(n) = \lim_{n\to\infty} n = \infty \Longrightarrow \sigma_1$ is diffuse
- The contradictions prove that no best-response equilibrium exists

Similar solution concepts

- The basic problem with finitely additive probabilities – non-integrability of payoff functions – has been addressed by others
- Finitely additive extension of a zero-sum game (Yanovskaya 1970)
- Optimistic equilibrium (Vasquez 2017: price-competition example)
- Justifiable equilibrium (Flesch et al. 2018): a strategy profile σ such that for every player i and alternative strategy τ_i

$$\overline{\int_S} u_i(s)\, d\sigma(s) \coloneqq \inf\{ \textstyle\int_S g(s)\, d\sigma(s) \mid g \text{ simple measurable}, g \geq u_i \}$$

$$\geq \underline{\int_S} u_i(s)\, d(\tau_i, \sigma_{-i})(s)$$

$$\coloneqq \sup\{ \textstyle\int_S g(s)\, d(\tau_i, \sigma_{-i})(s) \mid g \text{ simple measurable}, g \leq u_i \}$$

- The payoff function u_i is assumed bounded

Similar solution concepts

Theorem. Every best-response equilibrium is a justifiable equilibrium but not conversely.

- A single player has action set $[0,1]$ and payoff function $u = 1_{\mathbb{Q}}$
- The algebra $\mathcal{J}$ is all finite unions of subintervals of $[0,1]$
- A simple measurable function $0 \leq g \leq 1$ satisfies
  - $g \leq 1_{\mathbb{Q}}$ if and only if $g = 0$ outside some finite subset of $\mathbb{Q}$
  - $g \geq 1_{\mathbb{Q}}$ if and only if $g = 1$ outside some finite subset of $\mathbb{Q}^{\mathsf{C}}$
- The first fact gives $\underline{\int} 1_{\mathbb{Q}}\, d\tau = 1$ for $\tau = \delta_0$
- The condition for justifiable equilibrium is $\sigma(\{s\}) = 0$ for all $s \notin \mathbb{Q}$
- Every nonatomic probability on $[0,1]$ (e.g., Lebesgue measure) is one
- The condition for best-response equilibrium is stronger, $\sigma^*\left(\mathbb{Q}^{\mathsf{C}}\right) = 0$, which is equivalent to $\sum_{s \in \mathbb{Q} \cap [0,1]} \sigma(\{s\}) = 1$

Conceptual foundations

- In a mixed equilibrium, the choice of suboptimal actions is excluded
- The condition is both necessary and sufficient
- Only the supports of the players' strategies need to be examined
- Not alternative mixed strategies – the mixed extension is irrelevant
- Furthermore, here mixed strategies are not randomized strategies
- They are not played, and may not even be playable in any sense
- Represent others' assessment of the players' choices of actions
- The equilibrium condition is rational, best-response choice
- Actions yielding low expected payoffs are excluded
- The expectation is with respect to the other players' strategies
- A player has no use for the integral wrt the product probability
- The existence of this integral – the expected payoff – is optional
- It is not a requisite for a meaningful notion of mixed equilibrium