Best Response Correspondences

Game Theory · Module 1 — Strategic Form Games

Node 1.4

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.4

Dominance (Node 1.3) prunes strategies that are bad against every opponent profile. The **best-response correspondence** is the positive dual: it tells you exactly which strategies are optimal *against a given opponent profile*. Nash equilibrium will be defined as a fixed point of the joint best-response correspondence (Module 2), so this node is the direct precursor to the central solution concept of the course. It is also the object whose continuity, convex-valuedness, and upper-hemicontinuity power the Nash existence proof via Kakutani's theorem (Node 2.4).

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite game with mixed extension $\Delta(S_i)$ and multilinear expected-payoff U_i.

Best response (pure version). Given $s_{-i} \in S_{-i}$,

$$\mathrm{BR}_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}) \subseteq S_i.$$

Best response (mixed version). Given an opponent mixed profile $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$,

$$\mathrm{BR}_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i, \sigma_{-i}) \subseteq \Delta(S_i).$$

Joint best-response correspondence. Define

$$\mathrm{BR} : \prod_i \Delta(S_i) \rightrightarrows \prod_i \Delta(S_i), \qquad \mathrm{BR}(\sigma) = \mathrm{BR}_1(\sigma_{-1}) \times \cdots \times \mathrm{BR}_n(\sigma_{-n}).$$

A **Nash equilibrium** is a profile σ^* with $\sigma^* \in \mathrm{BR}(\sigma^*)$ — a fixed point of the joint correspondence.

Because BR_i may return multiple optima (ties are common when opponents mix), it is generally a *correspondence* (set-valued map), not a function. Writing $\mathrm{BR}_i : \cdot \rightrightarrows \cdot$ is standard notation for a correspondence.

Intuition

Three ways to think about best response.

As an oracle. If a small demon tells you exactly what your opponents will do, the best response is simply what you should do given that information. It is the single-player decision problem that emerges inside a game once the opponents' play is fixed.

As a reaction curve. In economic applications (e.g., Cournot duopoly), the best response is a function mapping the opponent's quantity to your optimal quantity. Plotting it gives the "reaction function"; the intersection of reaction functions is the Nash equilibrium.

As a fixed-point operator. The joint best response BR is a map from profiles to sets of profiles. An equilibrium is a profile that is returned by the map when fed into the map. This lifts the problem from "solve for equilibrium" to "find a fixed point of an operator," which is a *much* more powerful framing because we have centuries of fixed-point theorems to apply.

The set-valued nature is essential: when two pure strategies tie in expected payoff against σ_{-i}, both are best responses, and so is any convex combination of them. Ignoring this would miss most mixed Nash equilibria.

Structural Role

Best response is the joint pivot between dominance and equilibrium:

- **1.3 ⇒ 1.4.** A strictly dominated strategy is never a best response. The best-response graph lives inside the dominance-reduced game.
- **1.4 ⇒ 2.1.** Nash equilibrium is defined as a fixed point of BR.
- **1.4 ⇒ 2.4.** The existence proof (Kakutani) checks that BR is upper-hemicontinuous with non-empty compact convex values — each of these hypotheses is a property of BR, not of the payoffs directly.
- **1.4 ⇒ 2.5.** Rationalizability is the fixed point of iterated best-response elimination.
- **1.4 ⇒ 9.1.** Fictitious play iterates best responses against empirical frequencies.
- **1.4 ⇒ 9.3.** Counterfactual regret minimization is a weakening of best-response dynamics that guarantees convergence of time averages.

Best response is the functional heart of the course: almost every theorem and almost every algorithm in the remaining modules is phrased as a statement about BR.

Mathematical Development

Characterization of mixed best response

By linearity of U_i in σ_i, a mixed strategy σ_i is in $\mathrm{BR}_i(\sigma_{-i})$ if and only if

$$\mathrm{supp}(\sigma_i) \subseteq \arg\max_{s_i \in S_i} U_i(s_i, \sigma_{-i}).$$

Equivalently, the pure strategies in $\mathrm{supp}(\sigma_i)$ must all tie for first place in expected payoff, and every pure strategy outside the support must do no better. This is the **support characterization** of best response.

Properties of BR_i

Non-empty: $\Delta(S_{-i})$ is compact and U_i is continuous, so the argmax is non-empty. **Compact:** the argmax is a closed subset of the compact simplex. **Convex:** any convex combination of two best-response mixtures has support inside the same argmax set, hence is also a best response. (This uses multilinearity crucially.) **Upper hemicontinuous:** if $\sigma_{-i}^{(k)} \to \sigma_{-i}$ and $\sigma_i^{(k)} \in \mathrm{BR}_i(\sigma_{-i}^{(k)})$ with $\sigma_i^{(k)} \to \sigma_i$, then $\sigma_i \in \mathrm{BR}_i(\sigma_{-i})$. This is the Berge Maximum Theorem applied to a continuous objective on a compact domain.

These four properties — non-empty, compact, convex, upper-hemicontinuous — are exactly the hypotheses of Kakutani's fixed-point theorem.

Reaction functions in continuous games

For games with continuous action spaces and strictly concave payoff in own action (e.g., Cournot with quadratic costs and linear demand), the argmax is a singleton, and the best-response correspondence reduces to a genuine function $r_i(s_{-i})$. The reaction function is differentiable under standard regularity conditions, and its slope at equilibrium carries information about stability of best-response dynamics.

Cournot example. Two firms produce quantities $q_1, q_2 \geq 0$; price is $P(Q) = a - Q$ with $Q = q_1 + q_2$; marginal cost c. Firm 1's payoff is $\pi_1 = (a - q_1 - q_2 - c)q_1$. First-order condition:

$$\frac{\partial \pi_1}{\partial q_1} = a - 2q_1 - q_2 - c = 0 \quad \Longrightarrow \quad r_1(q_2) = \frac{a - c - q_2}{2}.$$

Symmetric for firm 2. The Nash equilibrium is the intersection $q^* = (a - c)/3$ each.

Fixed-point formulation

A profile $\sigma^* \in \prod_i \Delta(S_i)$ is a Nash equilibrium iff σ^* is a fixed point of the self-map $f(\sigma) := \mathrm{BR}(\sigma)$ — meaning $\sigma^* \in \mathrm{BR}(\sigma^*)$ in the set-valued sense. All of Module 2 builds on this formulation.

Worked Example

Cournot with three firms

Three symmetric firms produce $q_i \geq 0$; demand $P = 12 - Q$ with $Q = q_1 + q_2 + q_3$; marginal cost $c = 0$. Firm i's profit:

$$\pi_i(q_i, q_{-i}) = (12 - q_i - q_{-i,\mathrm{sum}})q_i, \qquad q_{-i,\mathrm{sum}} = \sum_{j \neq i} q_j.$$

FOC: $12 - 2q_i - q_{-i,\mathrm{sum}} = 0 \Rightarrow r_i(q_{-i,\mathrm{sum}}) = (12 - q_{-i,\mathrm{sum}})/2$.

At a symmetric equilibrium $q_1 = q_2 = q_3 = q^*$, the opponent sum is $2q^*$. The fixed-point equation becomes

$$q^* = (12 - 2q^*)/2 = 6 - q^* \quad \Longrightarrow \quad q^* = 3.$$

Each firm produces 3 units, total output is 9, price is 3, and each firm earns 9. This is the unique symmetric Nash equilibrium, obtained entirely by setting $\sigma^* = \mathrm{BR}(\sigma^*)$.

Discrete analogue. Consider Matching Pennies. Row's best response to Column playing $(q, 1 - q)$:

$$\mathrm{BR}_1(q) = \begin{cases} \{\text{Heads}\} & q > 1/2 \\ \Delta(\{H, T\}) & q = 1/2 \\ \{\text{Tails}\} & q < 1/2. \end{cases}$$

The fixed point is $q = 1/2$, where both players mix 50/50. Notice BR_1 is convex-valued at $q = 1/2$ — the entire simplex is the best-response set — which is what allows the fixed point to land there.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
$\mathrm{BR}_i(\sigma_{-i})$	The maximally exploitative strategy against a fixed opponent strategy
$\mathrm{BR}(\sigma^*) = \sigma^*$	GTO — every player is already playing maximally against every other
Best response as exploit	"Nodelocking" a solver to opponent tendencies and reading the exploitative strategy off the solution
Upper hemicontinuity of BR	Small changes in opponent frequencies lead to small changes in the exploit — until a threshold is crossed, then the exploit jumps
Convex-valued	Any mixture of two tied best responses is also a best response — "mix at indifferences"

Concrete situation

100 BB deep, BTN vs BB. BTN opens 2.5x, BB calls. Flop $K\spadesuit 9\heartsuit 4\diamondsuit$. BTN bets 33% pot. BB's raise frequency in equilibrium is around 8% on this board with a specific range.

Exploit-by-best-response. Suppose you (BTN) are facing a fish whose raise frequency on this board texture is empirically around 2% — they almost never raise. What is your best response? Given the lower raise frequency, your c-bet EV is higher because you face less resistance; in particular, you can c-bet **wider** than solver recommends because the punishment for c-betting a weak range (getting raised off it) is reduced. Concretely, solver says c-bet 62% of range; exploit says c-bet $\approx$ 80%.

Similarly, against a fish whose c-bet defense is too folding-heavy, your best response is to c-bet 100% of range at a smaller size. Against a fish whose c-bet defense is too calling-heavy, your best response is to reduce bluff frequency and value-bet thinner.

These are literally best-response computations in the sense of this node. You fix σ_{-i} (the fish's strategy) and compute $\mathrm{BR}_i(\sigma_{-i})$ — the strategy that maximizes your EV against the observed opponent. Solvers support this via "nodelocking": you freeze the opponent's strategy at some node and re-solve your own play.

What this understanding buys you

Exploit = deviate from GTO toward the best response against observed opponent. GTO is the Nash equilibrium — the fixed point of the joint best-response map. If the opponent is not at their own GTO (they are off-equilibrium), then your GTO is not a best response to them — it is only guaranteed not to *lose*. The best response to their specific mistake is strictly more profitable than GTO.

The art of live poker is **when** to deviate toward the best response and **how far**. Deviate too little, and you leave money on the table; deviate too much, and you become exploitable in return if the opponent adjusts. The best-response correspondence is the map you are walking along when you move from GTO toward pure exploit — and the upper-hemicontinuity of BR is what guarantees small reads translate into small exploits, not wild swings.

Cross-Links

- **1.1 Players, strategies, payoffs** — inputs to the BR operator.
- **1.2 Mixed strategies** — the domain on which BR lives.
- **1.3 Dominance** — dominated strategies lie outside the image of BR.
- **2.1 Nash equilibrium** — fixed point of BR.
- **2.4 Nash existence** — BR's properties power the Kakutani proof.
- **9.1 Fictitious play** — iterated best response to empirical frequencies.

- **Economics** — **reaction functions** — the differentiable case.
- **Systems thinking** — **fixed points** — the unifying structural concept.

Structural Compression

Invariant: For every opponent profile, there is a non-empty, convex, compact set of own-strategies that maximize expected payoff, and this set varies upper-hemicontinuously with the opponent profile.

Mechanism: The Berge Maximum Theorem applied to the multilinear objective U_i on the compact convex domain $\Delta(S_i)$ — continuity plus compactness of the parameter space gives the well-behavedness of the correspondence.

Cross-System Echo: Best-response correspondences are the multi-agent analogue of the gradient in single-agent optimization: they tell you which direction to move given the current state of the environment. The fixed-point formulation of equilibrium is the multi-agent analogue of the first-order condition in calculus.

Exercises

1. Compute $\mathrm{BR}_1(q)$ for Matching Pennies as a function of $q = \Pr[\text{Column plays H}]$. Verify it is convex-valued only at $q = 1/2$.
2. For symmetric Cournot with n firms, demand $P = a - Q$, marginal cost c, derive the best-response function of firm i and the symmetric Nash equilibrium $q^* = (a-c)/(n+1)$.
3. Prove that $\mathrm{BR}_i(\sigma_{-i})$ is always convex: if $\sigma_i, \sigma_i' \in \mathrm{BR}_i(\sigma_{-i})$, then any convex combination is also in $\mathrm{BR}_i(\sigma_{-i})$.
4. Prove that BR_i is upper-hemicontinuous. (You may use the Berge Maximum Theorem.)
5. Give an example of a game in which BR_i is not *lower*-hemicontinuous at some point. (Lower-hemicontinuity is the non-trivial condition — it does not follow from continuity of payoffs.)
6. Verify the support characterization of best response: $\sigma_i \in \mathrm{BR}_i(\sigma_{-i})$ iff every pure strategy in $\mathrm{supp}(\sigma_i)$ yields the maximal expected payoff.
7. **Argue, structurally,** why Nash equilibrium is naturally defined as a fixed point of the joint best-response correspondence BR rather than as a solution of some system of payoff equations directly. What is the mathematical advantage of phrasing equilibrium as a fixed-point problem on a correspondence?

Game Theory · Module 1 — Strategic Form Games · Node 1.4
Concepts: fixed-point, nash-equilibrium, convexity
Prerequisites: 1.1, 1.2, 1.3
Enables: 1.5, 1.6, 2.1, 2.2, 2.3, 2.4

hbar.university — own.your.cognition