Compatible Explanations

Youngji Sohn

University of Warwick

October 2026

Explaining strategic behavior

Players may not fully know the environment in which they interact.

They use models to explain their observations and guide their strategies.

With limited feedback and prior information, these models may be misspecified.

In equilibrium as a steady-state play (Self-confirming, Berk–Nash equilibrium),
each player forms the best statistical explanation of her observations and best respond to that.

A player's explanation also makes claims about what the opponent would observe and the opponent's response.

Could the opponent's own inference and incentives justify that conjecture?

How can equilibrium beliefs incorporate this reasoning?

Approach

Esponda and Pouzo's (2016) framework for misspecified games.

A Berk–Nash equilibrium requires each player to:

1. hold a belief supported on hypotheses that best account for her observations within her model;
2. choose a strategy optimal under that belief.

Propose an equilibrium refinement that are only supported by players' explanations that are compatible with each others' statistical inference and strategic incentives.

Example 1: Investment game

Player 1 invests (I) or takes an outside option (O). Player 2 assists (A) or blocks (B). They choose simultaneously and observe only their own action and payoff.

	A	B
I	$(2,1)$	$(0,0)$
O	$(1,1)$	$(1,0)$

Consider two steady state play:

Investment: (I, A)

1. Player 1's correct conjecture $\widehat{\sigma}_2 = A$ and the observed payoff $u_1 = 2$ match.
2. $\sigma_1 = I$ best responds to the conjecture.

Outside option: (O, A)

1. A wrong conjecture $\widehat{\sigma}_2 = B$ and the observed payoff $u_1 = 1$ match.
2. $\sigma_1 = O$ best responds to the conjecture.

Why expect blocking?

Now suppose player 1 knows player 2's payoffs and regards her as rational.

	A	B
I	$(2, \underline{1})$	$(0, 0)$
O	$(1, \underline{1})$	$(1, 0)$

Player 2 always prefers A, whatever she believes about player 1.

$\sigma_1 = O$ can be optimal only if player 1 expects B with probability $\geq 1/2$.

Respecting player 2's rationality rules out (O, A).

Strategic reasoning under uncertainty

In the investment game, player 2's incentives rule out blocking whatever she believes.

With uncertainty about Nature, an action may be optimal under some beliefs and suboptimal under others.

To justify a conjecture about the opponent's behavior, we therefore need to ask what statistical inference she can make from her observations and model.

Example 2: Coordination game

Two players choose L or R simultaneously without observing Nature's state, α or β.
The true probability $p(\alpha) = 1/2$.

The outcome is high exactly when their actions match and the state is α. Otherwise low.

	Player 1		Player 2	
Own action	High	Low	High	Low
L	3	0	3	0
R	0	2	1	2

Each observes only her own action and her outcome.

Subjective models

- Player 1 considers $\Pi^1 = \{1/2, 1\}$.
- Player 2 considers $\Pi^2 = \{1\}$.

Two pure Berk–Nash equilibria

Own action	Player 1 High	Player 1 Low	Player 2 High	Player 2 Low
L	3	0	3	0
R	0	2	1	2

$\Pi^1 = \{1/2, 1\}$

$\Pi^2 = \{1\}$

Play: (L, L)

	$\pi^i(\alpha)$	$\widehat{\sigma}^j(L)$	$E[u^i(L)]$	$E[u^i(R)]$
Player 1	1	1/2	**3/2**	1
Player 2	1	1/2	**3/2**	3/2

Play: (R, R)

	$\pi^i(\alpha)$	$\widehat{\sigma}^j(L)$	$E[u^i(L)]$	$E[u^i(R)]$
Player 1	1/2	0	0	**1**
Player 2	1	1/2	3/2	**3/2**

Both profiles satisfy each player's own statistical and optimality requirements.

Explanation supporting (L, L)

Player 1 chooses L and observes high half the time.

Exactly two hypotheses in her model reproduce those observations:

$\pi^1(\alpha)$	$\widehat{\sigma}^2(L)$	$E[u^1(L)]$	$E[u^1(R)]$	L optimal?
$1/2$	1	$3/2$	2	No
1	$1/2$	$3/2$	1	Yes

Thus **every belief supporting** L must give positive probability to the second hypothesis: Nature is always α and player 2 mixes equally.

Could player 2 justify that mixed strategy if the hypothesis were true?

What would player 2 observe?

Suppose Nature is always α, player 1 chooses L, and player 2 mixes equally between L and R.

Player 2 would observe after each action:

	High outcome	Payoff
L	Always	3
R	Never	2

In her model, Nature is always α. These observations can therefore be explained only by conjecturing that player 1 always chooses L.

Under that belief, L pays 3 and R pays 2, so the equal mixed strategy is not optimal.

Every belief supporting player 1's L must use this hypothesis, which cannot justify player 2's mixed strategy.

Thus (L, L) is not compelling.

Explanation supporting (R, R)

Player 1 chooses R and observes high half the time.

Exactly two hypotheses in her model reproduce those observations:

$\pi^1(\alpha)$	$\widehat{\sigma}^2(L)$	$E[u^1(L)]$	$E[u^1(R)]$	R optimal?
$1/2$	0	0	1	Yes
1	$1/2$	$3/2$	1	No

Using the **first hypothesis**, player 2 would also observe high half the time after R.

Player 2's model has Nature always choosing α.
She can fit these observations by conjecturing that player 1 mixes equally:
both actions pay $3/2$, so R is optimal.

Can player 1 justify this mixed strategy under player 2's explanation?

Justifying player 1's hypothetical mixed strategy

Under player 2's explanation, Nature is always α and player 2 chooses R.

Player 1 would observe **low after L and high after R** and payoff of 0.

Believing that Nature is always α and player 2 always chooses R fits these observations.

And under such a belief, player 1 would expect payoff 0 by deviating to L. So the belief supports player 1 choosing R.

The justification continues.

Main results

I define such a refinement as rationalizable Berk–Nash equilibrium (RBNE).

1. **Compatible explanations.** A Berk–Nash profile is an RBNE exactly when each player has a supporting belief whose explanations belong to a family closed under further justification.

2. **Existence.** Nash play under possibly different, mistaken beliefs survives when the associated hypotheses best account for both actual and hypothetical observations reciprocally.

3. **Under simplifying assumptions.** An informative strategy profile is an RBNE exactly when it remains a Nash equilibrium under every successive pair of projected beliefs about Nature.

Related literature

1. **Equilibrium as steady state play**

   SCE: Fudenberg and Levine (1993). Berk–Nash: Esponda and Pouzo (2016).

   Optimal choices under beliefs sustained by observations, allowing best available statistical fit under misspecification.

2. **Rationalizable conjectures**

   RCE: Rubinstein and Wolinsky (1994), Esponda (2013).
   Rationalizable partition-confirmed equilibrium: Fudenberg and Kamada (2015, 2018).

   Conjectures are disciplined recursively through opponents' rationality and the observations available to them.

3. **Strategic inference under misspecification**

   Frick, Iijima, and Ishii (2022); He and Libgober (2026); (Heidhues, Kőszegi, and Strack (2025)).

   Inference about others' behavior under misspecified models.

This paper: a general refinement of Berk–Nash equilibrium requiring recursive justification of conjectures through opponents' own models, hypothetical observations, and incentives.

Framework

Game

Nature draws a state ω from a finite Ω according to a full-support distribution p.

Two players choose simultaneously without observing the state.

1. Player i chooses x^i from a finite set X^i.
2. She observes her own action and a consequence in a finite set Y^i:

$$y^i = f^i(x^i, x^j, \omega).$$

3. Her payoff is $u^i(x^i, y^i)$.

A strategy is $\sigma^i \in \Sigma^i = \Delta(X^i)$. $\sigma = (\sigma^1, \sigma^2)$.

Under p and σ, the true distribution of player i's consequences is

$$Q^i_\sigma(y^i \mid x^i) = \sum_{\omega \in \Omega} \sum_{x^j \in X^j} p(\omega)\sigma^j(x^j)\mathbf{1}\{f^i(x^i, x^j, \omega) = y^i\}.$$

Subjective models and beliefs

Player i considers state distributions in $\Pi^i \subseteq \Delta(\Omega)$ and every opponent strategy in Σ^j possible.

$$\underbrace{\theta^i = (\pi^i, \widehat{\sigma}^j)}_{\text{a hypothesis}} \in \underbrace{\Theta^i = \Pi^i \times \Sigma^j}_{\text{subjective model}}.$$

- A hypothesis specifies Nature's distribution and the opponent's strategy, independently.
- A belief $\mu^i \in \Delta(\Theta^i)$ assigns probabilities to hypotheses.
- The model excludes the truth when $p \notin \Pi^i$.

Under $\theta^i = (\pi^i, \widehat{\sigma}^j)$, the predicted distribution of player i's consequences is

$$Q^i_{\theta^i}(y^i \mid x^i) = \sum_{\omega \in \Omega} \sum_{x^j \in X^j} \pi^i(\omega) \widehat{\sigma}^j(x^j) \mathbf{1}\{f^i(x^i, x^j, \omega) = y^i\}.$$

Berk–Nash equilibrium

Weighted KL loss compares actual observations with a hypothesis's predictions:

$$K^i(\sigma, \theta^i) = \sum_{x^i \in X^i} \sum_{y^i \in Y^i} \underbrace{\sigma^i(x^i) Q^i_\sigma(y^i \mid x^i)}_{\text{steady-state frequency}} \ln \frac{Q^i_\sigma(y^i \mid x^i)}{Q^i_{\theta^i}(y^i \mid x^i)}.$$

A profile σ is a **Berk–Nash equilibrium** if, for each player i, some belief $\mu^i \in \Delta(\Theta^i)$ satisfies:

1. **Statistical fit:** every hypothesis in the support best accounts for the observations.

$$\operatorname{supp} \mu^i \subseteq \arg\min_{\theta^i \in \Theta^i} K^i(\sigma, \theta^i).$$

2. **Optimality:** her strategy maximizes expected payoff under that belief.

$$\sigma^i \in \mathrm{BR}^i(\mu^i).$$

The opponent's hypothetical problem

The game structure and both subjective model classes are common knowledge.

Take a Berk–Nash profile σ and a supporting belief μ^i for player i.

Each supported hypothesis $(\pi^i, \widehat{\sigma}^j)$, together with her strategy σ^i, gives an explanation:

$$e^i = \big(\underbrace{\pi^i,\ \sigma^i}_{j\text{'s hypothetical problem}} \ ; \ \underbrace{\widehat{\sigma}^j}_{j\text{'s strategy to justify}} \big).$$

Let $\mathcal{R}^{j,0}(\pi^i, \sigma^i)$ denote the strategies player j can justify in this problem under some belief over her own hypotheses satisfying:

1. **Statistical fit:** every hypothesis in the support best accounts for the observations in the hypothetical problem.
2. **Optimality:** the strategy maximizes expected payoff under that belief.

The conjectured strategy is justified when

$$\widehat{\sigma}^j \in \mathcal{R}^{j,0}(\pi^i, \sigma^i).$$

Recursive justification

Apply the same requirement recursively.

For $n \geq 0$, define

$$\tau^i \in \mathcal{R}^{i,n+1}(\rho, \tau^j) \iff \exists \nu^i \in \Delta(\Theta^i) \quad \text{such that}$$
$$\tau^i \in \mathrm{BR}^i(\nu^i), \qquad \nu^i\big(\{(q, \widehat{\tau}^j) \in \Theta^i_\rho(\tau) : \widehat{\tau}^j \in \mathcal{R}^{j,n}(q, \tau^i)\}\big) = 1.$$

Return to player i's actual problem (p, σ^j).

Definition

A profile $\sigma \in \Sigma$ is a **rationalizable Berk–Nash equilibrium (RBNE)** if

$$\sigma^i \in \mathcal{R}^{i,n}(p, \sigma^j) \qquad \text{for every } i \text{ and every integer } n \geq 0.$$

What does RBNE add?

Berk–Nash equilibrium explains each player's choice by solving her own statistical problem.

But a hypothesis supporting her choice also describes an environment for the opponent.
In that environment, the opponent's own inference and incentives may contradict the behavior being conjectured.

RBNE requires supporting beliefs whose hypotheses can justify that behavior recursively.

Strategic reasoning can therefore narrow equilibrium predictions without additional observations.

Which explanations survive?

Pure profile	Own fit and optimality	Recursive justification
(L, L)	Satisfied	Fails
(R, R)	Satisfied	Satisfied

At (L, L), every supporting belief for player 1 needs a hypothesis that cannot justify player 2's conjectured mixed strategy.

At (R, R), the required explanations justify one another, even though player 2 remains mistaken about Nature.

RBNE distinguishes mistaken beliefs that can be justified through opponents' inference and incentives from those that cannot.

Existence

Characterization by compatible explanations

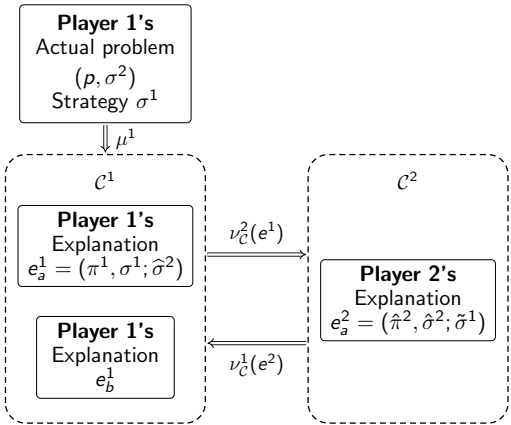
A family of explanations is **compatible** if each explanation can be justified through the opponent's own model, hypothetical observations, and incentives, using only further explanations within the family.

Characterization

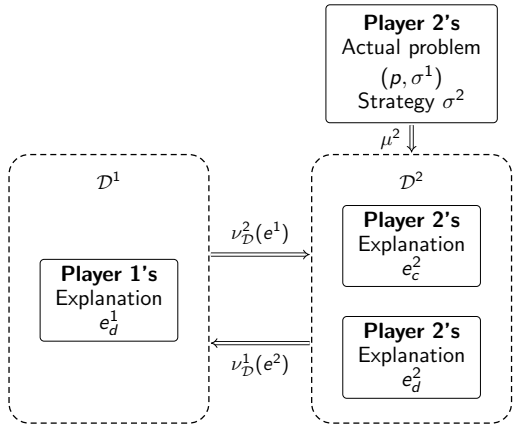
Under the maintained regularity conditions, a profile is an RBNE if and only if each player has a Berk–Nash supporting belief whose every generated explanation belongs to a compatible family.

Families may branch or be infinite, and the two players may use different families.

When do subjective models and feedback allow such supporting beliefs?



(a) Player 1's reasoning



(b) Player 2's reasoning

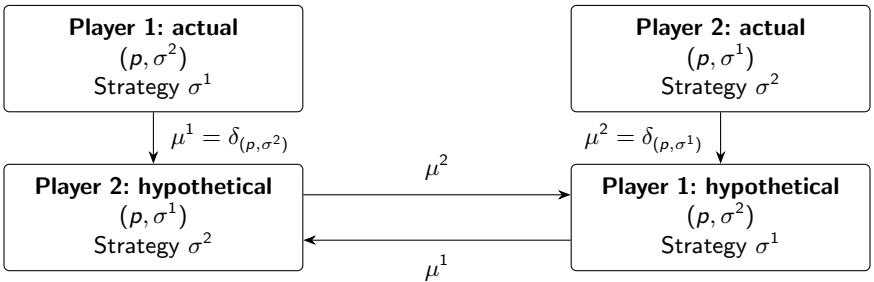
1 Nash cycles

Suppose $p \in \Pi^1 \cap \Pi^2$. At $\sigma \in NE(p)$, player i can use the correct belief $\delta_{(p,\sigma^j)}$.

Existence when both players consider the true state distribution

Since the objective game has a Nash equilibrium,

$$p \in \Pi^1 \cap \Pi^2 \quad \Longrightarrow \quad NE(p) \subseteq \mathcal{E}^\infty, \qquad \mathcal{E}^\infty \neq \varnothing.$$

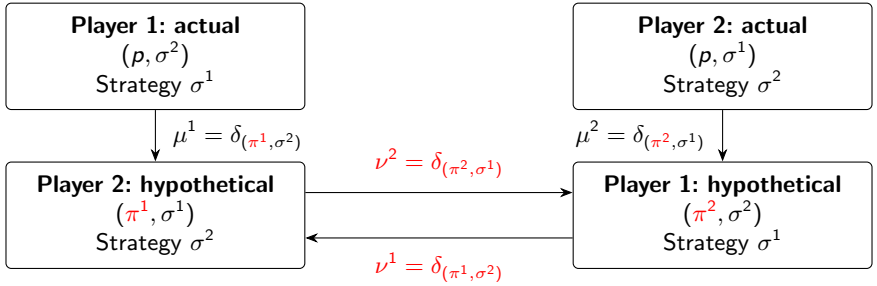


2 Nash cycles with diverse priors

Take $\pi^i \in \Pi^i$ and $\sigma \in NE(\pi^1, \pi^2)$.

Sufficient condition

If each hypothesis (π^i, σ^j) best accounts for observations under σ in both her actual and hypothetical problems, then $\sigma \in \mathcal{E}^\infty$.



3 All-or-nothing feedback

Feedback is **all-or-nothing** in two respects:

1. **Opponent's action.** After each own action, feedback either identifies the opponent's action across states, or is independent of that action when the state is fixed.
2. **Nature.** At each fixed action profile, feedback either identifies the state or is independent of it.

The wKL loss therefore separates inference about the opponent from inference about Nature.

For a Nature distribution ρ, define its **KL projections** onto player i's model by

$$\mathrm{P}_i(\rho) = \arg\min_{q \in \Pi^i} D(\rho \| q), \qquad D \text{ is KL divergence.}$$

A correct opponent conjecture combined with a KL projection of Nature is sufficient for statistical fit.

3 Existence through reciprocal projections

Suppose both players' feedback is all-or-nothing. Choose priors satisfying, for both players $i \neq j$,

$$\pi^i \in \underbrace{\mathrm{P}_i(p)}_{\text{actual problem}} \cap \underbrace{\mathrm{P}_i(\pi^j)}_{\text{hypothetical problem}}.$$

At any $\sigma \in NE(\pi^1, \pi^2)$, the point belief $\delta_{(\pi^i, \sigma^j)}$ best accounts for both problems and makes σ^i optimal. These beliefs justify one another.

Nash play with diverse priors survives

$$NE(\pi^1, \pi^2) \subseteq \mathcal{E}^\infty, \qquad \mathcal{E}^\infty \neq \varnothing.$$

Common model. If $\Pi^1 = \Pi^2 = \Pi$, write $\mathrm{P} = \mathrm{P}_1 = \mathrm{P}_2$. Any $\pi \in \mathrm{P}(p)$ gives $NE(\pi, \pi) \subseteq \mathcal{E}^\infty$, even when $p \notin \Pi$.

Convexity and informative play are not required.

3 A characterization through successive projections

Maintain regularity and full support of p. Suppose both feedback functions are all-or-nothing, both models are convex, and σ is informative.

For full-support ρ, write $\mathrm{P}_i(\rho) = \{\mathrm{p}_i(\rho)\}$ for the unique KL projection. Start with the projections of p and project each player's distribution onto the opponent's model:

$$\pi_0^i = \mathrm{p}_i(p), \qquad \pi_{k+1}^i = \mathrm{p}_i(\pi_k^j), \qquad i \neq j, \quad k \geq 0.$$

Characterization

$$\sigma \in \mathcal{E}^\infty \quad \Longleftrightarrow \quad \sigma \in NE(\pi_k^1, \pi_k^2) \quad \text{for every } k \geq 0.$$

A profile is *informative* if each player uses an action identifying opponent behavior and some profile in the joint support has an injective state map for that player. These opportunities need not coincide.

The same profile must remain optimal under every pair of projected distributions. The index k records reasoning depth, not time. Strategies stay fixed, and no convergence of the projections is required.

Appendix

Objective game and regularity

The objective game is $\mathcal{O} = \langle I, \Omega, p, X, Y, f, u \rangle$, with $I = \{1, 2\}$. State, action, and consequence sets are finite, and p has full support.

$$X = X^1 \times X^2, \qquad Y = Y^1 \times Y^2, \qquad \Sigma = \Delta(X^1) \times \Delta(X^2).$$

Each $f^i : X^i \times X^j \times \Omega \to Y^i$ specifies feedback, and $u^i : X^i \times Y^i \to \mathbb{R}$ specifies payoff. Neither player receives information about the state before choosing.

Regularity. For every player i:

1. Π^i is nonempty and compact.
2. Every $\theta^i \in \Theta^i$ is the limit of a sequence $(\theta^i_m)_m \subseteq \Theta^i$ such that $Q^i_{\theta^i_m}(y^i \mid x^i) > 0$ for every m, own action x^i, and consequence y^i feasible after that action.

The underlying game structure and both subjective models are common knowledge. The true distribution and actual beliefs need not be known to the opponent.

Statistical fit in a hypothetical problem

In player j's problem (ρ, τ^i), test τ^j; write $\tau = (\tau^1, \tau^2)$ in player order. Let Q^j_{ρ,τ^i} denote her consequence distribution under Nature ρ and opponent strategy τ^i.

$$K^j_\rho(\tau; \theta^j) = \sum_{x^j \in X^j} \tau^j(x^j) \sum_{y^j \in Y^j} Q^j_{\rho,\tau^i}(y^j \mid x^j) \ln \frac{Q^j_{\rho,\tau^i}(y^j \mid x^j)}{Q^j_{\theta^j}(y^j \mid x^j)}.$$

$$\Theta^j_\rho(\tau) = \arg\min_{\theta^j \in \Theta^j} K^j_\rho(\tau; \theta^j).$$

- The candidate's own action frequencies τ^j weight the fit; only played actions receive weight.
- Minimize over the full original model Θ^j. Actual fitting is the case $\rho = p$ and $\tau = \sigma$.

Within each hypothesis, Nature and the opponent's action are independent. A belief mixing hypotheses may produce dependence.

Use $0 \ln(0/b) = 0$ for $b \geq 0$, $a \ln(a/0) = +\infty$ for $a > 0$, and $0 \cdot (+\infty) = 0$.

Supporting beliefs at each depth

Start with the original full-model minimizers, and retain those whose opponent conjectures pass the depth-$(n-1)$ test:

$$\Theta^{i,0}(\sigma) = \Theta^i_p(\sigma), \qquad \Theta^{i,n}(\sigma) = \left\{(\pi^i, \widehat{\sigma}^j) \in \Theta^{i,0}(\sigma) : \widehat{\sigma}^j \in \mathcal{R}^{j,n-1}(\pi^i, \sigma^i)\right\} \quad (n \geq 1).$$

For every $n \geq 0$, the supporting beliefs are

$$\mathcal{B}^{i,n}(\sigma) = \left\{\mu^i \in \Delta(\Theta^i) : \mu^i\left(\Theta^{i,n}(\sigma)\right) = 1, \quad \sigma^i \in \mathrm{BR}^i(\mu^i)\right\}.$$

$$\mathcal{E}^n = \left\{\sigma \in \Sigma : \mathcal{B}^{i,n}(\sigma) \neq \varnothing \text{ for every } i\right\}.$$

- One belief must support the entire mixed strategy. Its individual hypotheses need not do so separately. Optimality compares every available action, including unplayed actions.
- Every hypothesis used by that belief must pass. Excluding a minimizer does not make a worse fit admissible.

Compatible families: formal definition

Seek a family containing all further explanations generated by its chosen justifying beliefs. Player i's explanations form the set

$$E^i = \left\{ (\rho^i, \tau^i; \widehat{\tau}^j) : (\rho^i, \widehat{\tau}^j) \in \Theta^i,\ \tau^i \in \Sigma^i \right\}.$$

Collect them in $\mathcal{C} = (\mathcal{C}^1, \mathcal{C}^2)$, where each $\mathcal{C}^i \subseteq E^i$ is nonempty and closed.

The family is **compatible** if, for every $e^i = (\rho^i, \tau^i; \widehat{\tau}^j) \in \mathcal{C}^i$, there exists a belief $\nu^j \in \Delta(\Theta^j)$ such that:

- ν^j is supported on player j's full-model KL minimizers in the problem (ρ^i, τ^i) when she uses $\widehat{\tau}^j$.
- $\widehat{\tau}^j \in \mathrm{BR}^j(\nu^j)$.
- Every supported hypothesis generates an explanation in the other player's set:

$$(\rho^j, \widehat{\tau}^i) \in \operatorname{supp} \nu^j \quad \Longrightarrow \quad (\rho^j, \widehat{\tau}^j; \widehat{\tau}^i) \in \mathcal{C}^j.$$

The choice of ν^j may depend on e^i. Families may branch and need not be finite or form a cycle.

Characterization: the proof

For $e^i = (\rho^i, \tau^i; \widehat{\tau}^j)$, write $h(e^i) = (\rho^i, \tau^i)$ and $c(e^i) = \widehat{\tau}^j$. Define $\mathcal{R}^{i,\infty} = \bigcap_{n \geq 0} \mathcal{R}^{i,n}$.

Compactness. Regularity gives compact response graphs and compact, nested supporting-belief sets. At $\sigma \in \mathcal{E}^\infty$, choose

$$\mu^i \in \bigcap_{n \geq 0} \mathcal{B}^{i,n}(\sigma).$$

The same argument in each hypothetical problem gives a single justifying belief whose generated conjectures survive all depths.

Necessity. Define the nonempty closed sets

$$\mathcal{C}^i = \left\{ e^i \in E^i : c(e^i) \in \mathcal{R}^{j,\infty}\left(h(e^i)\right) \right\}.$$

Each $e^i \in \mathcal{C}^i$ has a justifying belief generating only explanations in $\mathcal{C}^j$. The actual beliefs μ^i generate explanations in $\mathcal{C}^i$.

Sufficiency. Induction shows that every conjecture in a compatible family passes every finite-depth test. Actual supporting beliefs generating explanations in such families therefore support $\sigma \in \mathcal{E}^\infty$.

The sufficiency argument permits different compatible families for the two actual players.

Existence constructions and robustness of nonexistence

Reciprocal projections under all-or-nothing feedback. Define

$$\mathrm{P}_i(\rho) = \arg\min_{q \in \Pi^i} D(\rho \| q).$$

If $\pi^i \in \mathrm{P}_i(p) \cap \mathrm{P}_i(\pi^j)$ for both players, then

$$NE(\pi^1, \pi^2) \subseteq \mathcal{E}^\infty, \qquad \mathcal{E}^\infty \neq \varnothing.$$

Each point belief fits its actual and reciprocal hypothetical observations and supports the same Nash strategy. This sufficient condition requires neither convex models nor an informative profile.

With a common model $\Pi^1 = \Pi^2 = \Pi$, write $\mathrm{P} = \mathrm{P}_1 = \mathrm{P}_2$ and $NE(\pi) = NE(\pi, \pi)$. Every $\pi \in \mathrm{P}(p)$ gives $NE(\pi) \subseteq \mathcal{E}^\infty$, including when $p \notin \Pi$.

Finite detection and local persistence. Under maintained regularity, nestedness and compactness imply

$$\mathcal{E}^\infty(p, u) = \varnothing \quad \Longrightarrow \quad \mathcal{E}^N(p, u) = \varnothing \quad \text{for some finite } N.$$

Closedness of the finite-depth graph then gives $\mathcal{E}^N(p', u') = \varnothing$ for all sufficiently close (p', u') with p' of full support.

State, action and consequence spaces, feedback functions, and subjective models are held fixed; only p and the utilities assigned to observed consequences vary.

Nash cycles: formal sufficient condition

Let $\pi^i \in \Pi^i$ for $i = 1, 2$, and let $\sigma \in NE(\pi^1, \pi^2)$.

For a source distribution ρ, write $\Theta^i_\rho(\sigma) = \arg\min_{\theta^i \in \Theta^i} K^i_\rho(\sigma; \theta^i)$, with $\Theta^{i,0}(\sigma) = \Theta^i_\rho(\sigma)$.

Nash-cycle proposition

If, for each $i \neq j$,

$$(\pi^i, \sigma^j) \in \Theta^{i,0}(\sigma) \cap \Theta^i_{\pi^j}(\sigma),$$

then $\sigma \in \mathcal{E}^\infty$.

The point belief $\delta_{(\pi^i, \sigma^j)}$ fits player i's actual and reciprocal hypothetical observations and makes σ^i optimal. The same two beliefs justify one another.

This is a sufficient condition for the chosen profile. General compatible families can use different beliefs across hypothetical problems.

All-or-nothing feedback

Player i's feedback is *all-or-nothing* if both conditions hold.

1. Opponent's action

For each own action x^i, either different opponent actions always give different consequences, even across states,

$$f^i(x^i, b, \omega) \neq f^i(x^i, b', \omega') \quad \text{for all distinct } b, b' \in X^j \text{ and all } \omega, \omega' \in \Omega,$$

or feedback is independent of the opponent's action when the state is fixed,

$$f^i(x^i, b, \omega) = f^i(x^i, b', \omega) \quad \text{for all } b, b' \in X^j \text{ and } \omega \in \Omega.$$

2. State

For every action profile (x^i, x^j), the map $\omega \mapsto f^i(x^i, x^j, \omega)$ is either injective or constant.

A profile σ is *informative* if, for each player, some own action in $\operatorname{supp} \sigma^i$ reveals the opponent's action, and some profile in $\operatorname{supp} \sigma^i \times \operatorname{supp} \sigma^j$ has an injective state map.

KL projections and convex models

Under all-or-nothing feedback, KL loss separates into opponent-strategy fit and state fit. For $\rho, q \in \Delta(\Omega)$, define

$$D(\rho\|q) = \sum_{\omega\in\Omega} \rho(\omega) \log \frac{\rho(\omega)}{q(\omega)}, \qquad \mathrm{P}_i(\rho) = \arg\min_{q\in\Pi^i} D(\rho\|q).$$

Use $0\log(0/b) = 0$ for $b \geq 0$ and $a\log(a/0) = +\infty$ for $a > 0$.

A model Π^i is *convex* if $\lambda q + (1-\lambda)q' \in \Pi^i$ for every $q, q' \in \Pi^i$ and $\lambda \in [0,1]$.

Under regularity, all-or-nothing feedback, and an informative profile, every full-support ρ has a finite KL projection with full support. If Π^i is convex, strict convexity makes it unique. Write $\mathrm{P}_i(\rho) = \{\mathrm{p}_i(\rho)\}$.

Nash equilibrium with different state distributions

For $q^1, q^2 \in \Delta(\Omega)$, let $NE(q^1, q^2)$ be the set of profiles σ such that, for each player i, every action in $\operatorname{supp}\sigma^i$ maximizes expected payoff against σ^j under q^i. Expected payoffs use the consequence distributions $Q^i_{q^i,\sigma^j}(\cdot \mid x^i)$.

Reproducing observations

Exact observational reproduction

Player i's model reproduces her observations at σ if some $\theta^i \in \Theta^i$ satisfies

$$Q^i_{\theta^i}(\cdot \mid x^i) = Q^i_{\sigma}(\cdot \mid x^i) \qquad \text{for every } x^i \in \operatorname{supp} \sigma^i.$$

This is equivalent to $\min_{\theta^i \in \Theta^i} K^i(\sigma, \theta^i) = 0$.

Example. Two players choose whether to play P or stay out O. There are two states, H and L.

Staying out pays zero. Playing pays one at (H, P) or (L, O), where the second coordinate is the opponent's action, and zero otherwise.

Each player observes only her own payoff. The objective distribution and the common subjective model are

$$p(H) = \frac{1}{2}, \qquad \Pi^1 = \Pi^2 = \{q\}, \qquad q(H) = \frac{3}{4}.$$

Exact fit and the opponent's choice

Let r be the probability of P in an opponent conjecture. Under (q, r),

$$U^i(P; (q, r)) = \frac{3}{4}r + \frac{1}{4}(1 - r) = \frac{1}{4} + \frac{1}{2}r > 0 = U^i(O; (q, r)).$$

Thus P is strictly optimal under every belief over subjective parameters.

Under $p(H) = 1/2$, playing pays one with probability $1/2$, whatever the opponent's strategy. The parameter $r = 1/2$ reproduces both actions' consequence distributions at every profile.

At (P, P), this is the unique minimizing parameter, and it supports P. Therefore,

$$\mathcal{E}^0 = \{(P, P)\}.$$

The opponent's choice

The fitted parameter conjectures $\frac{1}{2}P + \frac{1}{2}O$. In the opponent's problem (q, P), O is still strictly suboptimal under every belief. The conjectured mixed strategy cannot be justified.

$$\mathcal{E}^1 = \mathcal{E}^\infty = \varnothing.$$

Relation to rationalizable conjectural equilibrium

Represent rationalizable conjectural equilibrium (RCE) in Esponda (2013) using a common space $\Pi \ni p$ of state distributions as fundamentals and complete mixed strategies. Players know their own strategy and the consequence distribution after each action in its support.

Exact consistency and recursive optimality

At a problem (ρ, τ), a supporting belief $\mu^i \in \Delta(\Pi \times \Sigma^j)$ must satisfy

$$\mu^i\big(\{(q,b) : K^i_\rho(\tau;(q,b)) = 0\}\big) = 1, \qquad \tau^i \in \mathrm{BR}^i(\mu^i).$$

Apply these conditions at every depth. Let $\mathcal{E}^{\mathrm{RCE}}_\Pi(p)$ denote the set of profiles with such recursive support.

With $\Pi^1 = \Pi^2 = \Pi \ni p$, every actual or generated ρ belongs to Π. The parameter (ρ, τ^j) attains zero wKLD, so minimizing wKLD requires exact consistency. Hence

$$\mathcal{E}^\infty = \mathcal{E}^{\mathrm{RCE}}_\Pi(p).$$

RCE already has recursive justification. The present framework allows different models and positive fitting loss.

Cournot: actual fit and optimal choices

Both pure profiles $(2,1)$ and $(1,2)$ generate prices 6 and 10 with probabilities $1/2$ each. Given own quantity, prices identify both the state and the rival's quantity.

Each firm's unique best-fitting hypothesis correctly conjectures the rival quantity. Firm 1 uses high-demand probability $2/3$, while firm 2 uses $1/2$.

	$(q_1, q_2) = (2,1)$	$(q_1, q_2) = (1,2)$
Firm 1: profit from $q_1 = 1, 2$	$26/3$, $\mathbf{34/3}$	$\mathbf{17/3}$, $16/3$
Firm 2: profit from $q_2 = 1, 2$	$\mathbf{5}$, 4	8, $\mathbf{10}$

These point beliefs support the two displayed profiles as Berk–Nash equilibria. They do not exhaust the Berk–Nash equilibrium set.

Cournot: hypothetical fit and choices

Under firm 1's explanation, the high-demand frequency is $2/3$. Firm 2 uniquely fits $4/5$ and the correct rival quantity:

$$D(2/3\|1/2) - D(2/3\|4/5) = \frac{1}{3}\ln\frac{128}{125} > 0.$$

Profile being tested	$(2,1)$	$(1,2)$
Firm 2: profit from $q_2 = 1, 2$	$31/5$, **$32/5$**	$46/5$, **$62/5$**
Quantity attributed to firm 2	1	2
Firm 1's explanation	Fails	Passes

At $(1,2)$, firm 2's actual explanation uses $1/2$ and attributes $q_1 = 1$. In the resulting problem, firm 1 fits its only allowed demand probability, $2/3$, and prefers quantity 1: $17/3 > 16/3$.

Both firms therefore pass the first test at $(1,2)$. The following compatible family verifies all-depth survival.

Cournot: a continuing justification

At $(q_1, q_2) = (1, 2)$, write each explanation as (high-demand probability, own quantity; rival quantity).

$$\mathcal{C}^1 = \{(2/3, 1; 2)\}, \qquad \mathcal{C}^2 = \{(4/5, 2; 1),\ (1/2, 2; 1)\}.$$

Each row uses a belief concentrated on the opponent's unique best-fitting hypothesis.

Starting explanation	Generated explanation	Conjectured choice is optimal
Firm 1: $(2/3, 1; 2)$	Firm 2: $(4/5, 2; 1)$	$q_2 = 2$: $62/5 > 46/5$
Firm 2: $(4/5, 2; 1)$	Firm 1: $(2/3, 1; 2)$	$q_1 = 1$: $17/3 > 16/3$
Firm 2: $(1/2, 2; 1)$	Firm 1: $(2/3, 1; 2)$	$q_1 = 1$: $17/3 > 16/3$

Every generated explanation stays in the family, so it is compatible. The actual supporting beliefs generate $(2/3, 1; 2)$ and $(1/2, 2; 1)$, both in the family. Hence $(1, 2) \in \mathcal{E}^\infty$.

Firm 2's hypothetical belief at $4/5$ differs from its actual supporting belief at $1/2$.

Cournot: two possible steady states

Two firms choose $q_i \in \{1, 2\}$ before observing demand. With $Q = q_1 + q_2$,

$$P_\ell(Q) = 15 - 3Q, \qquad P_h(Q) = 19 - 3Q, \qquad u_i = q_i(P - 3).$$

Each observes its own quantity and profit. High demand occurs with $1/2$.

Firm 1 allows only high-demand probability $2/3$. Firm 2 allows $1/2$ or $4/5$.

Compare $(q_1, q_2) = (2, 1)$ and $(1, 2)$.
Both generate prices 6 and 10 with probability $1/2$, hence identifying the rival's quantity.
Their unique best-fitting hypotheses use $2/3$ for firm 1 and $1/2$ for firm 2.

Optimal quantity under those beliefs	$(2, 1)$	$(1, 2)$
Firm 1	2	1
Firm 2	1	2

Both profiles are Berk–Nash equilibria: each firm's fitted belief supports its choice.

If firm 1's explanation were right

Firm 1's fitted explanation uses high-demand probability $2/3$ and the correct rival quantity.

What would firm 2 observe if that explanation were true?

Source of observations	Price 6	Price 10	Firm 2's best fit
Actual play	$1/2$	$1/2$	$1/2$
Firm 1's explanation	$1/3$	$2/3$	$4/5$

Prices still identify firm 1's quantity. Firm 2 interprets the implied frequencies within its own model: $4/5$ fits better than its other hypothesis, $1/2$.

The explanation changes the observations to be explained, and hence firm 2's fitted belief.

Does this fit justify firm 2's choice?

At $(2,1)$, firm 1's explanation conjectures $q_2 = 1$. Keep firm 1's quantity fixed at $q_1 = 2$ and compare firm 2's incentives:

Firm 2's belief	Profit: $q_2 = 1$	Profit: $q_2 = 2$	Best q_2
Actual: $1/2$	**5**	4	1
Under firm 1's explanation: $4/5$	$31/5$	**32/5**	2

Firm 1's only best-fitting hypothesis cannot justify the quantity it conjectures. Thus $(2,1)$ is excluded.

At $(1,2)$, the same fitted probability $4/5$ makes firm 2 prefer $q_2 = 2$ against $q_1 = 1$. Firm 1's conjecture passes this test.

A justification creates another explanation

At $(1, 2)$, the belief justifying firm 2's quantity also conjectures that firm 1 chooses $q_1 = 1$.

$$\text{Firm 1: high demand } 2/3, \quad q_1 = 1, \quad \text{conjectures } q_2 = 2$$

$$\downarrow$$

$$\text{Firm 2: high demand } 4/5, \quad q_2 = 2, \quad \text{conjectures } q_1 = 1$$

This new explanation poses the same question in the reverse direction.

If firm 2's explanation were true, would firm 1's own model and incentives justify $q_1 = 1$?

One successful test creates further explanations to justify. The framework specifies this requirement for every explanation used in a supporting belief, at every depth.