Stable Berk-Nash Equilibrium

Yuichi Yamamoto
University of Pennsylvania

Bayesian Learning with Model Misspecification

- One agent. Infinite periods $t = 1, 2, \cdots$.
- In each period t...
  - The agent chooses an action $a \in A$.
  - Observes a signal $y \in Y$, with the dist $q(a) = (q(y|a))_{y \in Y}$.
  - Payoff is $u(a, y)$.
- Subjective model θ: y follows a dist $q^\theta(a) = (q^\theta(y|a))_{y \in Y}$.
- Initial prior μ on the subjective model space Θ.
  - Models are *misspecified* if $q \neq q^\theta$ for all θ.
- Agent chooses myopically optimal action which maximizes

$$g^{\mu^t}(a) = \sum_{\theta \in \Theta} \mu^t(\theta) \sum_{y \in Y} q(y|a) u(a, y).$$

Example: Monopolist with Unknown Demand

- Monopolist chooses a price a.
- The demand y is given by

$$y = f(a) + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2).$$

- In a subjective model θ, the demand is

$$y = b - \theta a + \varepsilon.$$

- Other examples:
  - Overconfidence. (Heidhues-Kosegi-Strack)
  - Gambler's fallacy. (He)

What is known?

- If models are correctly specified, the agent learns the true model.
  - Given any action, the belief moves toward the true model θ on average.
- With model misspecification, the process may not converge.
  - The "drift" of the belief depends on the action.
  - Model θ may fit the true distribution for some a, but not for other $\tilde{a}$.
- Use relative entropy to measure fitness. (Berk)

What is known?

- Esponda-Pouzo introduce *Berk-Nash equilibrium*.
  - Action-belief pair (α, μ) which is a fixed point of the learning process.
- Defined in a static model. Existence.
  - α may be a mixed action, in many cases.
- Necessity.
  - If action converges, it must be a Berk-Nash equilibrium action.
- No sufficiency.
  - Even if a BNE exists, the learning process may not converge.

What is known?

- $|\Theta| = 2$. (General payoff function. Fudenberg-Romanyuk-Strack)
  - If there is a pure BNE, (a, μ) converges almost surely.
  - If not, (a, μ) oscillates forever almost surely.
- Continuous Θ. (Examples only.)
  - HKS, He: Unique pure BNE, global convergence.
  - Nyarko: No pure BNE, action oscillates forever.
- Can mixed BNE be a log-run outcome?
  - Agent almost always chooses a pure action.
  - Action cannot converge to a mixture.

Findings

- Yes, the belief can converge to a mixed equilibrium belief.
  - Example with continuous Θ.
  - Minor modification to a finite model.
  - Action does not converge, but action frequency does.
  - Direct foundation of mixed BNE.
  - Perhaps action convergence is too demanding.

- General analysis for continuous Θ.
  - Log-concavity assumption.
  - If $\Theta = [0, 1]$, the belief converges almost surely.
  - If $\Theta \subseteq \mathbf{R}^k$, the belief may not converge.

Notation

- $h = (a^\tau, y^\tau)_{\tau=1}^{\infty}$: Sample path.
- $h^t = (a^\tau, y^\tau)_{\tau=1}^{t}$: History with length t.
- A strategy is a mapping $s : \bigcup_{t=0}^{\infty} H^t \to \triangle A$.
- P^s: Probability measure over all sample paths h, given s.
- $\mu^{t+1}(h^t)$: the posterior belief in period $t+1$, given h^t.
- $\mu^{t+1}(h) = \mu^{t+1}(h^t)$, where h^t is a truncation.

Stable Beliefs

- s is *myopically optimal* given an initial prior μ if for each h^t,

$$s(h^t) \in \arg\max g^{\mu^{t+1}(h^t)}(\alpha).$$

- Belief μ is *stable* if there is an initial prior with full support and a myopically optimal strategy s such that

$$P^s \left(\lim_{t \to \infty} \mu^t(h) = \mu \right) > 0.$$

- Use a topology of weak convergence.

- Goal: Characterize stable beliefs.

Relative Entropy

- $p(\alpha)$: True distribution of (a, y) given a mixed action α.
- $p^\theta(\alpha)$: Distribution of (a, y) induced by a model θ.
- Relative entropy between $p(\alpha)$ and $p^\theta(\alpha)$:

$$d(p(\alpha), p^\theta(\alpha)) = \sum_{(a,y)} p(a, y|\alpha) \log \frac{p(a, y|\alpha)}{p^\theta(a, y|\alpha)}$$

- Expected "surprise."
- A "pseudo" metric.
  - $d(p, \tilde{p}) \geq 0$. Equality iff $p = \tilde{p}$.
  - Does not satisfy symmetry, triangle inequality.

Berk-Nash Equilibrium

- "Best" model given α:

$$\Theta(\alpha) = \arg\min_{\theta \in \Theta} d(p(\alpha), p^{\theta}(\alpha)).$$

- (α^*, μ^*) is a *Berk-Nash equilibrium* if...
  (i) $\alpha^* \in \arg\max_{\alpha \in \triangle A} g^{\mu^*}(\alpha)$.
  (ii) $\text{supp}\mu^* \subseteq \Theta(\alpha^*)$.

- Fixed point of the learning process.
  - If the current belief is μ^*, the agent must choose α^*.
  - If α^* is chosen forever, $\text{supp}\mu^t$ must converge to $\Theta(\alpha^*)$. (Berk)

- Stable beliefs=BNE beliefs?

Necessity

- Subjective models are *regular* if...
  (i) Θ is finite, or a compact convex set in $\mathbf{R}^k$.
  (ii) If Θ is convex, $q^\theta(y|a)$ is continuous in θ.
  (iii) $q^\theta(y|a) > 0$ for all θ, a, and y.
- Stronger than necessary...

Proposition

Under regularity, any stable belief μ is a BNE belief.

- So if the belief converges, it is a BNE belief.
- Esponda-Pouzo: Any stable action α is a BNE action.

Proof Idea

- Pick an arbitrary stable belief μ^*.
- Need to find α^* s.t. (α^*, μ^*) is a BNE.
- Case 1: Optimal action is unique given μ^*.
  - Let a^* denote the optimal action.
  - Optimal action is a^* whenever the belief is close to μ^*.
  - So when the belief converges to μ^*, a^* is taken forever.
  - This implies $\text{supp}\mu \subseteq \Theta(a^*)$. (Berk)
  - (a^*, μ^*) is a Berk-Nash equilibrium.

Proof Idea

- Case 2: Both a^* and $\tilde{a}^*$ are optimal given μ^*.
  - Pick a sample path h such that $\lim_{t\to\infty} \mu^t(h) = \mu^*$.
  - When the belief approaches μ^*, the agent is almost indifferent.
  - But not exactly indifferent, as the belief is different from μ^*.
  - Action switches infinitely often.
  - Let $\alpha^t(h) = \frac{1}{t}\sum_{\tau=1}^{t} s(h^{t-1})$ be the action frequency until period t.
  - $\alpha^t(h) \in \triangle A$, so there is a convergent subsequence.
  - Let $\alpha^*(h)$ be the limit of the subsequence.
  - $\alpha^*(h)$ is optimal given μ, as it is a mixture of a^* and $\tilde{a}^*$.
  - For "almost all" h, $\text{supp}\mu \subseteq \Theta(\alpha^*(h))$.

Action Convergence vs Belief Convergence

- In general, the action cannot converge to a mixed action.
  - Optimal action is almost always unique.
- But the belief may converge to a mixed BNE belief.
  - In this case, action does not converge.
  - But action frequency can converge to a mixed action.
  - More on this later.

Sufficiency: Pure Strict BNE

Proposition

Assume regularity. If (a, μ) is a pure strict BNE with $|\Theta(a)| = 1$, then μ is stable.

- In this case, the action also converges to the equilibrium action.
- Proof idea:
  - Suppose that the initial prior is close to μ.
  - Then the agent chooses a.
  - Then her belief will be even closer to μ, on average.

Mixed BNE for Finite Θ

- Example: Two actions a_1, a_2. Two signals y_1, y_2.
- In the true model, two signals are equally likely, given any action.
- Two subjective models, θ_1 and θ_2.
- $g^{\theta_1}(a_1) = g^{\theta_2}(a_2) = 1$, and $g^{\theta_1}(a_2) = g^{\theta_2}(a_1) = 0$.
- Given a_1, the signal dist is 0.3-0.7 at θ_1, 0.6-0.4 at θ_2.
- Given a_2, the signal dist is 0.4-0.6 at θ_1, 0.7-0.3 at θ_2.

Mixed BNE for Finite Θ

- Simpler version of FRS example.
- If the model θ_1 is likely, the agent chooses a_1.
- But then she gets more convinced that the model is likely to be θ_2.
- Then she starts to choose a_2.
- Now she gets more convinced that the model is likely to be θ_1.
- Does this process converge?

Mixed BNE for Finite Θ

- No pure-strategy BNE.
- Unique mixed BNE (α, μ):
  - α mixes a_1 and a_2 equally.
  - μ puts prob 0.5 on both models.
- In this equilibrium...
  - Two models are equally close to the true world, given α.
  - Two actions are indifferent, given μ.
- If the belief converges, it must be the equilibrium belief.

Mixed BNE for Finite Θ

- But this equilibrium belief μ cannot be stable.
  - Suppose that the current belief is μ.
  - If the current signal is y_1, the posterior on θ_1 becomes

$$\mu'(\theta_1) = \frac{0.5 \cdot 0.3}{0.5 \cdot 0.3 + 0.5 \cdot 0.6} = \frac{1}{3}.$$

  - So the belief cannot stay around μ.

- Hence the belief oscillates.

- In general, mixed BNE is unstable when $|\Theta| < \infty$.

Mixed BNE for Convex Θ

- Now modify the subjective model space: $\Theta = [0, 1]$.
  - Given a_1, the prob of y_1 is $0.3(1-\theta) + 0.6\theta$.
  - Given a_2, the prob of y_1 is $0.4(1-\theta) + 0.7\theta$.
- $g^\theta(a_1) = 1 - \theta$ and $g^\theta(a_2) = \theta$.
- $\theta = 0, 1$ are identical with the models θ_1, θ_2 in the previous example.
- But now we have many "intermediate" models.
- Does it change the result?

Mixed BNE for Convex Θ

- There is a unique mixed BNE (α, μ):
  - α mixes a_1 and a_2 equally.
  - μ puts prob one on the model $\theta = 0.5$.
- Equilibrium action is the same, but the belief is not.
  - Relative entropy $d(p, \tilde{p})$ is convex in $\tilde{p}$.
  - "Moving to the middle" improves the fitness.
- This new belief is more "robust."
  - Observing one more signal does not change the posterior much.
- So it can be stable.

Mixed BNE for Convex Θ

Proposition

For any initial prior with full support and for any myopically optimal strategy s,

$$P^s\left(\lim_{t\to\infty}\mu^t(h)=\mu\right)=1.$$

- Global convergence to the mixed BNE belief.
- Action frequency also converges to the equilibrium action.
- First direct foundation of mixed BNE.
  - Esponda-Pouzo: I.i.d. payoff shock, so BNE is purified.
- New interpretation of Nyarko.
  - The belief can converge, even when the action does not.

Proof Idea

- Step 1: Belief concentration.

- Likelihood maximizer in period t:

$$\theta^t_{\max}(h^{t-1}) \in \arg\max_\theta \mu^t(h^{t-1})[\theta].$$

- After a long time, the posterior becomes approximately $\delta_{\theta^t_{\max}(h^{t-1})}$.

Lemma

For any smooth initial prior with full support, there is $T > 0$ and $K > 0$ such that for any $t \geq T$, for any history h^t, and for any θ,

$$\frac{\mu^{t+1}(\theta^{t+1}_{\max})}{\mu^{t+1}(\theta)} \geq \exp\left\{\frac{tK}{2}(\theta^{t+1}_{\max} - \theta)^2\right\}.$$

- Key: Convexity of relative entropy.

Proof Idea

- Step 2: Show $P^s(\lim_{t\to\infty} \theta^t_{\max}(h) = 0.5) = 1$.
- In the long run, BR is determined by $\theta^t_{\max}$.

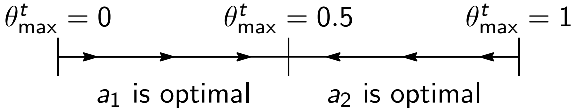

Proof Idea

- Intuition is simple, but the actual proof is not.
- Reason: θ_{max} is not a sufficient statistic.
  - $\theta_{\text{max}}^{t+1}$ depends not only on θ_{max}^{t}, but on μ^{t}.
- Key step: For any $\varepsilon > 0$, if a_1 is chosen forever, θ_{max} cannot drop by ε infinitely often.

Unstable Mixed BNE

- Now assume that...
  - Given a_1, the prob of y_1 is $0.3\theta + 0.4(1-\theta)$.
  - Given a_2, the prob of y_1 is $0.6\theta + 0.7(1-\theta)$.
- $g^\theta(a_1) = 1-\theta$ and $g^\theta(a_2) = \theta$.
- Given a_1, the agent gets more convinced that $\theta = 0$ is more likely.
- Given a_2, the agent gets more convinced that $\theta = 1$ is more likely.

Unstable Mixed BNE

- Two pure-strategy BNE.
  - $a = a_1$ and a degenerate belief on $\theta = 0$.
  - $a = a_2$ and a degenerate belief on θ_1.
- Also a mixed BNE.
  - α mixes two actions equally. A degenerate belief on $\theta = 0.5$.
- This mixed BNE is unstable.

Unstable Mixed BNE

Proposition

For any initial prior with full support, for any myopically optimal strategy s, and for any open set U containing $\theta = 0$ and $\theta = 1$,

$$P^s\left(\lim_{t\to\infty} \mu^t(h)[U \cap \Theta] = 1\right) = 1.$$

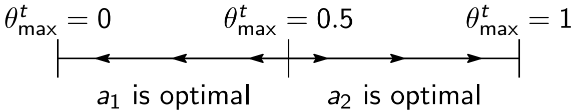

General Analysis for Continuous Θ

- Previous two examples: $q^\theta(y|a)$ is linear w.r.t. θ.
- Belief concentration under a more general condition.
- Signal distributions are *log-concave* if
  (i) Θ is convex and
  (ii) $q^\theta(y|a)$ is log-concave w.r.t. θ.
- Many exiting models satisfy this assumption.
  - Nyarko, He, Section 5 of HKS.
  - HKS's general model does not satisfy this.

General Analysis for Continuous Θ

- Strong identifiability: $q^{\theta}(a) \neq q^{\tilde{\theta}}(a) \ \forall a \forall \theta \forall \tilde{\theta}$.

Proposition

Assume log concavity and strong identifiability. Then $d(f, p^{\theta}(\alpha))$ is convex with respect to θ for any α and $f \in \triangle(A \times Y)$.

- Given an observation f, there is a unique "best" model θ.
- Hence belief concentration.
- Also, any equilibrium belief is δ_{θ} for some θ.

General Analysis for Continuous Θ

- Assume $\Theta = [0, 1]$.
- θ is *absorbing* if there is $\varepsilon > 0$, a_1, and a_2 (possibly $a_1 = a_2$) s.t.
  - Unique optimal action is a_1 for any $\tilde{\theta} \in (\theta - \varepsilon, \theta)$.
  - $\Theta(a_1) > \theta$.
  - Unique optimal action is a_2 for any $\tilde{\theta} \in (\theta, \theta + \varepsilon)$.
  - $\Theta(a_2) < \theta$.
- μ is *absorbing* if $\mu = \delta_\theta$ for some absorbing θ.

General Analysis for Continuous Θ

Proposition

Assume...

- $\Theta = [0, 1]$
- Regularity, log concavity, strong identifiability.
- Optimal action is unique except for finite θ.

Then the belief converges almost surely to an equilibrium belief.

Proposition

In addition, if

- Optimal action is unique for $\Theta(a)$

then the belief converges to an absorbing belief.

General Analysis for Continuous Θ

- When $\Theta \subset \mathbf{R}^k$, $\theta^t_{\max}$ can be cyclic.
- So the belief may not converge, even if there is a unique pure BNE.