Mislearning from Censored Data:

The Gambler's Fallacy in Optimal-Stopping Problems

Kevin He

September 16, 2019

Behavioral Motivation

Examples of the **gambler's fallacy (GaFa)**:

- expect fair coin to land on heads after 3 tails in a row
- believe daughter is "due" after several consecutive sons

In general: behavioral bias of expecting systematic reversals from random sequences

- **Lab evidence**: when asked to produce i.i.d. random sequences, subjects alternate too much between symbols
- **With strict incentives**: Lab experiments and field data find GaFa even in settings where GaFa is strictly costly
- **Chen, Moskowitz, and Shue (2016)**: GaFa also affects experienced decision-makers in high-stakes settings, e.g. asylum court, loan approval, ...

This project: Novel implications of GaFa in optimal-stopping problems (OSP) when agents learn about the underlying distributions from histories.

Illustrative Example

A junior HR manager recruiting for one job opening

- ... by sequentially interviewing candidates
- To decide whether to hire a candidate, must form belief about the **distribution** of talent in future applicants
- She asks senior managers about their beliefs, based on their hiring experience

Suppose all managers have GaFa:

- From a fixed labor pool, exaggerate how unlikely it is to get:
    - **consecutive below-average** applicants
    - **consecutive above-average** applicants
- Same psychology as exaggerating how unlikely it is for a fair coin to land on tails or heads consecutively

Question: What are the implications of GaFa for managers' beliefs and behavior over time?

Histories as Censored Data

Social belief aggregates past managers' **hiring experience** / **histories**

But these histories **truncated** when predecessors decided to stop

- No one observes the additional candidates who "would have" been found with longer searches
- History truncation as a **data censoring mechanism**

Stopping rules impose a **selective** censoring effect on histories

- In recruiting and other natural OSP, stop when early draws deemed "good enough"
- Whether later draw is censored depends on the realization of the earlier draw

Censoring effect, interacted with GaFa, leads to over-pessimism

Censoring Effect + GaFa ⇒ Over-Pessimism

Managers with below-average early interviewees keep searching

Suppose their later candidates are...

- Above-average ⇒ **expected** positive reversal after bad initial outcomes, not strong signals about pool
- Below-average ⇒ **surprised** (because they understate prob of bad streaks), strong and negative signals about pool

Under-infer from good draws, **over**-infer from bad draws

- Predecessors communicate an over-pessimistic impression of the pool to today's manager
- In turn affects her stopping strategy and the kind of (censored) history she gets

Symmetrically, over-optimistic after consecutive good draws

- But, can't have this kind of history due to censoring effect
- An **asymmetric** prediction from the "symmetric" GaFa bias

Summary of the Model

The Environment. Population of short-lived agents (e.g. HR managers in different years) take turns playing the same **OSP**:

- Get a "draw" each period
- Decide between stopping and getting payoff based on current draw, or continuing for a future payoff based on future draw

Domain of Learning. Agents uncertain about underlying **distributions** of draws (e.g. talent distribution in labor pool)

- Infer parameters of these distributions — the "fundamentals" — from histories
- Agents are Bayesians except for the behavioral bias

Main Contributions

1. Novel **channel of mislearning** for behavioral agents
   - **Interaction** between bias (GaFa) and data censoring
   - Agents learn correctly if either obstacle removed
   - This channel drives the main results:
     - Over-pessimistic beliefs in long run, starting from any priors
     - Positive-feedback cycle between distorted beliefs and distorted stopping rules
2. **Dynamics** of misspecified learning from **endogenous data**
   - Actions (stopping thresholds) affect observables (histories)
   - Prove a.s. convergence to unique steady state, from any prior
   - Builds on a literature that has studied this problem in:
     - Binary-states setting: Fudenberg, Romanyuk, Strack (2017)
     - Self-confirming setting: Heidhues, Kőszegi, Strack (2018)
       - Fixing DGP, some feasible belief fits data exactly
       - Thus, learning steady state = self-confirming eqm
   - This paper: non-self-confirming misspec. with rich state space and action space ▶ other literature

1. Learning with a Sequence of Agents

Stage Game: Optimal-Stopping Problem

- At the center of the model is a (single-agent) stage game
- Agents play this stage game, one at a time
- Results hold for all games in a class of OSP, but focus on a simple example for today's talk:
  - Draw X_1 in period 1
  - Can either **stop** and get payoff X_1...
  - ... or **continue** and draw X_2 in period 2
    - » in which case, payoff is X_2

The Domain of Learning

- Society uncertain about the distribution of (X_1, X_2)
- Agents engage in Bayesian learning about this distribution
  - Start with a prior over a class of distributions (**models**)
  - Learning = updating belief over models in light of data
- Will work with a general class of log-concave distributions. Some notation:
  - f_1, f_2 strictly positive densities on $\mathbb{R}$ with finite variance
  - Assume f_1, f_2 are strictly log-concave, symmetric, mean-zero
  - e.g., normal distribution, logistic distribution, ...
  - $f_1(\cdot \mid \tau_1)$ and $f_2(\cdot \mid \tau_2)$ shifted versions of f_1, f_2 centered around τ_1, τ_2
- now describe: what a typical model looks like

A Subjective Model of (X_1, X_2)

Definition

For $\mu_1, \mu_2 \in \mathbb{R}$, $\gamma \geq 0$, $\Psi(\mu_1, \mu_2; \gamma)$ is the distribution where

$$\left\{ \begin{array}{c} X_1 \sim f_1(\cdot \mid \mu_1) \\ (X_2|X_1 = x_1) \sim f_2(\cdot \mid \mu_2 - \gamma(x_1 - \mu_1)) \end{array} \right\}$$

For $\gamma > 0$, $(X_2 \mid X_1 = x_1)$ decreases in FOSD order with x_1

- Reflects a belief in reversal in luck, conditional on fundamentals
- X_1 above average $\Rightarrow$ X_2 likely below average, vice versa

A Subjective Model of (X_1, X_2)

Equivalently, $\Psi(\mu_1, \mu_2; \gamma)$ is the model where:

$$X_1 = \mu_1 + \epsilon_1$$
$$X_2 = \mu_2 + \epsilon_2$$

with ϵ_1, ϵ_2 mean-zero "luck" terms s.t.

- $\epsilon_1 \sim f_1$
- $(\epsilon_2 \mid \epsilon_1) \sim f_2(\cdot \mid -\gamma\epsilon_1)$.

Negative correlation between ϵ_1, ϵ_2 conditional on μ_1, μ_2

$\Rightarrow$ belief in reversal of luck

Feasible Models

- Objectively, $(X_1, X_2) \sim \Psi(\mu_1^\bullet, \mu_2^\bullet; \mathbf{0})$
  - i.e. $X_1 \sim f_1(\cdot \mid \mu_1^\bullet)$ independent from $X_2 \sim f_2(\cdot \mid \mu_2^\bullet)$.
- Prior belief supported on **feasible models**

$$\{\Psi(\mu_1, \mu_2; \gamma) : (\mu_1, \mu_2) \in \mathcal{M}\}$$

for **fixed** $\gamma > 0$, where $\mathcal{M} \subseteq \mathbb{R}^2$ are **feasible fundamentals**.
  - Take easiest case for talk: $\{\Psi(\mu_1^\bullet, \mu_2; \gamma) : \mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]\}$
  - Prior given by density $g_0 : [\underline{\mu}_2, \bar{\mu}_2] \to \mathbb{R}_{>0}$

Misspecified Bayesian Learning

- Agents dogmatic about $\gamma > 0$, a mistake given true $\gamma = 0$
- True distribution not a feasible model, so updating g_0 amounts to **misspecified** Bayesian learning
  - Use misspecification to represent/study behavioral bias
    - Agents are misspecified Bayesians and the nature of their misspecification interpreted as GaFa
  - Take as given agents have GaFa, explore implications
  - Consistent with field evidence that GaFa persists with learning

Stopping Strategies and Censoring of Histories

Definition

A (stage-game) **strategy** is a function $S : \mathbb{R} \to \{\text{Stop, Continue}\}$

Strategy S maps realization of $X_1 = x_1$ into a stopping decision.

Definition

A **history** of the stage game is $h \in \mathbb{H} := \mathbb{R} \times (\mathbb{R} \cup \{\varnothing\})$.

- If agent stops after $X_1 = x_1$, then $h = (x_1, \varnothing)$
  - $\varnothing =$ missing data indicator. If agent stops in period 1, history does not contain **counterfactual** second draw
- If agent continues and draws $X_2 = x_2$, then $h = (x_1, x_2)$
- Data **endogeneity**: S determines how histories are censored
- Censoring effect depends on having a dynamic stage-game

Cutoff Strategies

Special subclass: **cutoff strategies** S_c for $c \in \mathbb{R}$, where $S_c(x_1) = \text{Stop}$ iff $x_1 \geq c$.

Lemma (Optimality of cutoff strategies)

The best stopping-strategy under any belief over the feasible models is a cutoff strategy.

Intuition: Better x_1 draw increases stopping payoff and, under GaFa, predicts worse draws in next period

Timeline

A sequence of short-lived agents, one per round $t = 1, 2, 3, \ldots$

Agent 1 starts with some prior g_0

In round t, an agent arrives and...

1. **Adopts** final belief $\tilde{g}_{t-1}$ of agent $t-1$ as own prior
2. **Plays** stage game using cutoff $\tilde{C}_t$, maximizing expected payoff
3. **Generates** stage-game history $\tilde{H}_t$, collects payoff
4. **Updates** $\tilde{g}_{t-1} \to \tilde{g}_t$ by applying Bayes' rule to $\tilde{H}_t$

(Equivalent to a model where each agent starts with g_0, updates it using all predecessors' histories, then plays the stage game.)

Two key stochastic processes:

- $(\tilde{C}_t)_{t \geq 1}$, cutoffs of different agents
- $(\tilde{g}_t)_{t \geq 1}$, beliefs of different agents
- Probability space given by random draws in different rounds

Global Convergence of $(\tilde{C}_t), (\tilde{g}_t)$

$(\tilde{C}_t)$ and $(\tilde{g}_t)$ co-evolve with each other

- Will $(\tilde{C}_t)$ settle down, or can behavior cycle forever? (e.g. Nyarko (1991))
- Different long-run beliefs with different g_0?

Theorem 1

There exist some $c^\infty, \mu_2^\infty \in \mathbb{R}$ not depending on g_0 s.t.

- *If $\mu_2^\infty \in supp(g_0)$, a.s. $\tilde{C}_t \to c^\infty$ and $\tilde{g}_t \to \mu_2^\infty$ in L^1*
- *$\mu_2^\infty < \mu_2^\bullet$ and $c^\infty < c^\bullet$ ($c^\bullet$ = the objectively optimal cutoff)*

In the Long Run, Society Exhibits...

Over-pessimism ($\mu_2^\infty < \mu_2^\bullet$)

- Formalizes the result in the opening example
- In feasible models Ψ, two sources affect realization of X_2:
  - fundamental μ_2
  - reversal effect from X_1 realization
- Due to the **censoring effect**, on average expect positive reversal **when X_2 observed**
- "Two wrongs making a right": Pessimistic μ_2^∞ counteracts false expectation of positive reversal to best fit X_2 data

Early stopping ($c^\infty < c^\bullet$)

- Not implied by over-pessimism — GaFa can lead to late stopping even under (slightly) pessimistic beliefs
  - Intuition: GaFa agents perceive a better continuation value after bad initial draws
- μ_2^∞ sufficiently pessimistic in long-run to overcome this opposing force

Another Learning Environment

It remains to show global convergence.

As a first step, consider another learning environment replacing the single agent in round t with a **large generation** of agents moving simultaneously.

Two goals:

1. Develop tools needed to prove a.s. convergence in Thm 1, when agents move one by one
2. Clarify how censoring effect leads to positive-feedback cycle between distorted beliefs and distorted stopping strategies

2. Learning with Large Generations

Timeline (Large Generations)

A sequence of generations, $t \in \{0, 1, 2, ...\}$

Each with a **continuum** of short-lived GaFa agents $n \in [0, 1]$

- Think of each generation as large but finite
- Limit of what happens as finite cohort size $\to \infty$

When a new agent arrives in gen t:

1. **Observes** a large dataset containing histories of each predecessor in each previous gen, $(h_{s,n})_{n\in[0,1]}$ for $0 \leq s < t$
2. **Infers** fundamentals from histories
   - Agents now uncertain about both μ_1 and μ_2.
   - Feasible models $\{\Psi(\mu_1, \mu_2; \gamma); (\mu_1, \mu_2) \in \mathbb{R}^2\}$, fixed $\gamma > 0$
3. **Plays** stage game
4. **Generates** stage-game history

Inference from One Generation of Histories

Here is the key lemma for analyzing dynamics.

Suppose agent observes many histories censored by S_c. Can show belief eventually concentrates on $\mu_1^*(c), \mu_2^*(c)$.

Lemma

For all $c \in \mathbb{R}$, $\mu_1^*(c) = \mu_1^\bullet$ and $\mu_2^*(c) < \mu_2^\bullet$. Also, $\mu_2^*(c)$ is strictly increasing with $\lim_{c\to\infty} \mu_2^*(c) = \mu_2^\bullet$.

- μ_1 correctly estimated, μ_2 estimate too low for any $c \in \mathbb{R}$
- **Censoring effect**: distortion in belief about μ_2 depends on how data is censored
  - Less distortion with less censoring (i.e., larger c)
  - As $c \to \infty$ (i.e., no data censoring), distortion tends to 0

Can find closed-form expression of $\mu_2^*(c)$ in the Gaussian case:

$$\mu_2^*(c) = \mu_2^\bullet - \gamma(\mu_1^\bullet - \mathbb{E}[X_1 | X_1 \leq c])$$

Positive Feedback

Optimal cutoff: $C(\mu_1, \mu_2; \gamma) = \frac{1}{1+\gamma}\mu_2 + \frac{\gamma}{1+\gamma}\mu_1$

The pair of derivatives, $\frac{d\mu_2^*(c)}{dc} > 0$ and $\frac{\partial C(\mu_1,\mu_2;\gamma)}{\partial \mu_2} > 0$ imply that:

Proposition

Starting with gen 0 playing any cutoff strategy, beliefs $(\mu^*_{2,[t]})$ *and cutoff thresholds* $(c_{[t]})$ *form monotonic sequences across gens.*

$\mu_1^\bullet = \mu_2^\bullet = 0$, bias $\gamma = 0.5$, starting at $c_{[0]} = 0$

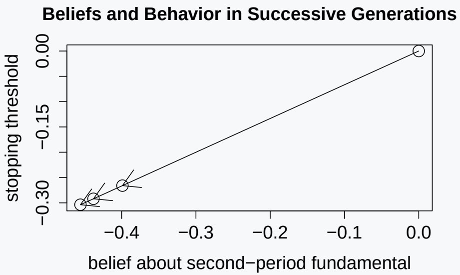

Steady State: Existence and Uniqueness

What are the long-run behavior of $c_{[t]}, \mu^*_{2,[t]}$?

- Tend to infinity?
- Different limits for different initial conditions (different $c_{[0]}$)?

$C(\mu_1, \mu_2; \gamma) = \frac{1}{1+\gamma}\mu_2 + \frac{\gamma}{1+\gamma}\mu_1$ Lipschitz continuous in μ_2 with constant $\frac{1}{1+\gamma} < \frac{1}{\gamma}$.

Next result holds for all stage games where this Lipschitz continuity applies for some constant $\ell < \frac{1}{\gamma}$.

Theorem 2 (monotonic convergence to unique steady state)

There exist $\mu_2^\infty, c^\infty \in \mathbb{R}$ so that for any $c_{[0]}$, $c_{[t]} \to c^\infty$ and $\mu^*_{2,[t]} \to \mu_2^\infty$ monotonically. Also, $(\mu_1^\bullet, \mu_2^\infty, c^\infty)$ is the LR outcome of Theorem 1.

Proof Idea: An Auxiliary Learning Environment

Consider an auxiliary large-gen learning environment

- Each gen only infers from histories of immediate predecessor gen
- Gen t believes in $(\mu_1^\bullet, \mu_2) \Rightarrow$ gen $t+1$ believes $\mu_2^*(\ C(\mu_1^\bullet, \mu_2; \gamma)\ )$
- We get a one-generation forward belief map,

$$\mathcal{I}(\mu_2) := \mu_2^*(C(\mu_2^\bullet, \mu_2; \gamma)).$$

$\mathcal{I}$ is a contraction! To see this:

Suppose $|\mu_2^{'} - \mu_2^{''}| = 1$. Then,

$$|C(\mu_1^\bullet, \mu_2^{'}; \gamma) - C(\mu_1^\bullet, \mu_2^{''}; \gamma)| \leq \ell < 1/\gamma$$

Also, can show $\mu_2^*(c)$ Lipschitz continuous w/ constant γ.

In Gaussian case,

$$\mu_2^*(c+1) - \mu_2^*(c) = \gamma\,(\mathbb{E}[X_1 \mid X_1 \leq c+1] - \mathbb{E}[X_1 \mid X_1 \leq c]) \leq \gamma$$

by property of Gaussian tail.

So, $|\mathcal{I}(\mu_2^{'}) - \mathcal{I}(\mu_2^{''})| \leq \gamma\ell < \gamma \cdot (1/\gamma) = 1$

Proof Idea: Dynamics in the Auxiliary Environment

Now use properties of contraction mapping:

- $\mathcal{I}$ has a unique fixed point $\mu_2^\infty \in \mathbb{R}$
- Belief dynamics in aux environment are iterates of $\mathcal{I} \Rightarrow$ converge globally to μ_2^∞

To conclude, compare baseline large-gen environment with auxiliary environment

- Observing datasets from many previous gens, almost all "close" to $\mathcal{H}^\bullet(C(\mu_1^\bullet, \mu_2^\dagger; \gamma)) \approx$ observing one dataset of $\mathcal{H}^\bullet(C(\mu_1^\bullet, \mu_2^\dagger; \gamma))$

3. Returning to a Sequence of Agents

Convergence Result for Sequence of Agents

Recall: Sequence of short-lived agents, play stage game one by one

- Stochastic processes of cutoffs $(\tilde{C}_t)$ and posterior beliefs $(\tilde{g}_t)$
- Feasible models: $\{\Psi(\mu_1^\bullet, \mu_2; \gamma) : \mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]\}$, $\gamma > 0$

Theorem 1

In games where $\mu_2 \mapsto C(\mu_1, \mu_2; \gamma)$ increasing and Lipschitz continuous with constant $\ell < 1/\gamma$, there exist some $c^\infty, \mu_2^\infty \in \mathbb{R}$ not depending on g_0 s.t.

- *If $\mu_2^\infty \in supp(g_0)$, a.s. $\tilde{C}_t \to c^\infty$ and $\tilde{g}_t \to \mu_2^\infty$ in L^1*
- *$\mu_2^\infty < \mu_2^\bullet$ and $c^\infty < c^\bullet$ ($c^\bullet =$ the objectively optimal cutoff)*

▶ stage games where this holds

Now give proof sketch of a.s. convergence

Iteratively Bounding Asymptotic Cutoffs and Beliefs

1. **Using beliefs to bound cutoffs**. $\text{supp}(g_0) = [\underline{\mu}_2, \bar{\mu}_2] \Rightarrow$ cutoff bounded in $[\underline{c}, \bar{c}]$, $\underline{c} = C(\mu_1^\bullet, \underline{\mu}_2; \gamma)$ and $\bar{c} = C(\mu_1^\bullet, \bar{\mu}_2; \gamma)$.
2. **Using cutoffs to bound beliefs**. $\tilde{C}_t$ asymptotically bounded in $[\underline{c}, \bar{c}] \Rightarrow$ belief asymptotically bounded in $[\mu_2^*(\underline{c}), \mu_2^*(\bar{c})]$.
   - Intuition: many histories censored by $S_c \Rightarrow$ belief $\to \mu_2^*(c)$
   - Issue: thresholds $(\tilde{C}_t)$ = stochastic process, can depend on previous histories
   - Solution: simplify log-likelihood using a version of LLN from Heidhues, Kőszegi, Strack (2018) ▶ details

Combining these two bounds, LR belief bounded in $[\mathcal{I}(\underline{\mu}_2), \mathcal{I}(\bar{\mu}_2)]$.

- This is because $\mathcal{I}$ is the composition $\mu_2 \mapsto \mu_2^*(C(\mu_1^\bullet, \mu_2; \gamma))$
- $\mathcal{I}$ is a contraction, and its iterates are monotonic (from Theorem 2)
- Repeat this argument $\Rightarrow$ LR belief bounded in $[\mathcal{I}^2(\underline{\mu}_2), \mathcal{I}^2(\bar{\mu}_2)]$, $[\mathcal{I}^3(\underline{\mu}_2), \mathcal{I}^3(\bar{\mu}_2)]$,...
- Iterate to deduce LR belief is the unique fixed point of $\mathcal{I}$

Uncertainty about μ_1

$(\mu_1, \mu_2), (\mu_1', \mu_2')$ on the same line with slope $-\gamma$, with $\mu_1, \mu_1' < \mu_1^\bullet$.

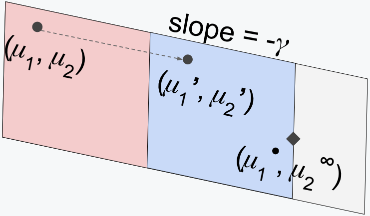

For any $x_1, x_2 \in \mathbb{R}$,

$$f_2(x_2; \mu_2 - \gamma(x_1 - \mu_1)) = f_2(x_2; \mu_2' - \gamma(x_1 - \mu_1'))$$

so $(\mu_1, \mu_2), (\mu_1', \mu_2')$ always explain second-period data equally well.

But, (μ_1', μ_2') explains first-period data strictly better.

Bound **directional deriv** of llh along $(1, -\gamma)$ for **any** stopping rule

- Asymptotic prob. of red region $= 0$
- Repeat $\Rightarrow$ restrict belief to vertical strip around $(\mu_1^\bullet, \mu_2^\infty)$

Comparative Statics and Policy Implications

How do changes in the stage game affect learning outcomes?

- Rational learners, endogenous data: no change because learn correctly
- Biased learners, exogenous data: no change because same data $\Rightarrow$ same inference, decision problem irrelevant

Both misspecification and endogenous data needed for comparative statics

Comparative Statics and Policy Implications

In simple search game, agent utility either x_1 or x_2

In general, stage game described by $u_1(x_1), u_2(x_1, x_2)$

- e.g. search with probability of recall has u_2 depending on x_1

Definition

For two games $(u_1, u_2^H), (u_1, u_2^L)$, u_2^H payoff dominates u_2^L ($u_2^H \succ u_2^L$) if for all x_1, $u_2^H(x_1, x_2) \geq u_2^L(x_1, x_2)$ for all x_2, with strict inequality for a positive-measure set.

Examples:

- $u_2^L = u_2^H - \kappa_{\text{wait}}$, $\kappa_{\text{wait}} > 0$ waiting cost
- $(u_1, u_2^H), (u_1, u_2^L)$ are search with probability of recall, and recall prob. higher in u_2^H

Proposition

$u_2^H \succ u_2^L \Rightarrow$ *steady state for* (u_1, u_2^H) *has strictly more optimistic belief about* μ_2 *and strictly higher cutoff threshold than* (u_1, u_2^L).

Comparative Statics and Policy Implications

Consider workers searching for jobs and inferring wage distribution from offer histories

- (Mueller, Spinnewijn, and Topa (2018) find evidence of GaFa in this setting using longitudinal surveys)
- This project: a mechanism for society of biased learners becoming over-pessimistic about wage distribution in LR

Policies that **subsidize longer search** (e.g., extended unemployment benefits) partially correct LR beliefs

- Such policies corrects beliefs **without** knowing true $\mu_1^\bullet, \mu_2^\bullet$
- Also, a test of misspecification – no effect on steady-state beliefs of correctly specified agents

Conclusion

- A model of GaFa agents learning the draw-generating **distributions** underlying an optimal-stopping problem, using others' **histories**.
- Key takeaway: novel channel of misinference for behavioral agents — **interaction** between bias and data censoring
  - Histories are **endogenously censored** by predecessors' stopping strategies, which depend on their beliefs.
  - **Interaction** between data censoring and GaFa drives main results and policy implications

Thank you!

Other Related Literature

Empirical evidence of GaFa:

- Producing/recognizing/predicting i.i.d. sequences: Budescu (1987), Bar-Hillel and Wagenaar (1991), Benjamin, Moore, Rabin (2017 WP)
  - State-run lotteries: Terrell (1994), Suetens, Galbo-Jørgensen, and Tyran (2016)
  - Casino gambling: Narayanan and Manchanda (2012)
  - Experienced, high-stakes: Chen, Moskowitz, and Shue (2016)

Inference by GaFa agents:

Rabin (2002), Rabin and Vayanos (2010)

Focus on learning from exogenous data. This paper: learners' actions affect statistical properties of the dataset.

Equilibrium concept under misspecification:

Esponda and Pouzo (2016) — not about global convergence dynamics

In settings with **finite actions**, Esponda, Pouzo, and Yamamoto have in-progress work on convergence of empirical distribution of actions to BNE

◀ Return

Regularity Assumptions on the Stage Game

◄ Return

1. [u_1 increases in x_1, u_2 increases in x_2]

$$x_1^{'} > x_1^{''}, x_2^{'} > x_2^{''} \text{ implies}$$
$$u_1(x_1^{'}) > u_1(x_1^{''}) \text{ and } u_2(x_1^{'}, x_2^{'}) > u_2(x_1^{'}, x_2^{''}).$$

2. [Higher 1st period draw affects u_1 more than u_2]

For $x_1^{'} > x_1^{''}$ and any $\bar{x}_2$, $u_1(x_1^{'}) - u_1(x_1^{''}) > |u_2(x_1^{'}, \bar{x}_2) - u_2(x_1^{''}, \bar{x}_2)|$.

3. [Non-degeneracy]

There exist $x_1^g, x_2^b, x_1^b, x_2^2 > 0$ s.t.

$$u_1(x_1^g) - u_2(x_1^g, x_2^b) > 0, \qquad u_1(x_1^b) - u_2(x_1^b, x_2^g) < 0.$$

4. [Continuity and integrability]

u_1, u_2 continuous and u_2 absolutely integrable against any $f_2(\cdot \mid \tau_2)$.

5. [ℓ-Lipschitz continuity in cutoffs]

Either: There exists $0 < \ell < \frac{1}{\gamma}$ so that for every $x_1, x_2 \in \mathbb{R}$ and $d > 0$,

$$u_1(x_1 + \ell d) - u_1(x_1) \geq u_2(x_1 + \ell d, x_2 + (1 - \gamma\ell)d) - u_2(x_1, x_2)$$

Or: $u_2(x_2)$ does not depend on x_1 and $u_2^{'}$ bounded below by some $\epsilon > 0$

The Log-Likelihood Process

Log-likelihood of data by end of round t under parameter $\mu_2 \in [\underline{\mu}_2, \bar{\mu}_2]$:

$$\sum_{s=1}^{t} \underbrace{\mathbf{1}\{X_{1,s} \leq \tilde{C}_s\} \cdot \ln(f_2(X_{2,s}; \mu_2 - \gamma(X_{1,s} - \mu_1^\bullet)))}_{\varphi_s(\mu_2)}$$

where $(X_{1,s}, X_{2,s})$ is the hypothetical pair of draws in round s.

$\bar{\varphi}_s(\mu_2) := \mathbb{E}[\varphi_s(\mu_2) \mid \mathcal{F}_{s-1}]$ — expected round s contribution to llh, given info up to round $s-1$

Apply a version of LLN for martingales with bounded quadratic variation from Heidhues, Kőszegi, Strack (2018) to show

$$\frac{1}{T}\sum_{s=1}^{T}(\varphi_s(\mu_2) - \bar{\varphi}_s(\mu_2)) \to 0 \quad \text{a.s.}$$

Simplifies analysis to just looking at $\sum_{s=1}^{t} \bar{\varphi}_s(\mu_2)$ ◄ Return