Learning under Misspecification:

"Pathwise Concentration Bounds.." Fudenberg, Lanzani, Strack
"Limit Points of Endogenous Misspecified Learning" FLS
"Which Misperceptions Persist?" FL

Princeton

May 13, 2021

Motivation

- People often have incorrect views of the world despite abundant data.
- Examples:
  - Belief that taxes are linear in income when they are not;
  - Belief in the "law of small numbers" and the gambler's fallacy;
  - "Causation neglect" about the impact of actions on outcomes;
  - Ignoring informative signals in the belief that they don't matter.
- It is important to understand how such agents learn from data, and how they will behave.
- To understand behavior with endogenous data it's not enough to know that beliefs converge, we need some bounds on the rate vis à vis the CLT.

Learning from Exogenous Data

- Berk [1966]: With exogenous i.i.d observations and a possibly misspecified prior, the posterior concentrates around the Kullback-Leibler minimizers with respect to the true data generating process. Various extensions in the statistics literature to other exogenous data generating processes. More recent work by economists extends this to endogenous data.
- Diaconis and Freedman [1990]: Sufficient condition for Bayesian posteriors to converge to the empirical distribution at a uniform exponential rate. This implies that non-myopic agents play myopically once they have enough data, a fact that has been useful in the analysis of non-equilibrium learning in games and related topics.
- "Pathwise Concentration Bounds for Bayesian Beliefs," FLS [2021] extends Berk [1966] by providing a rate of convergence, and extends Diaconis and Freedman [1990] to more general priors.

Misspecified Learning from Endogenous Data

- Esponda and Pouzo [2016] defines Berk-Nash equilibrium, and shows that it is a necessary property for limit points when the payoff function is subject to i.i.d. random shocks.
- Fudenberg, Romanyuk, and Strack [2017] shows the actions and beliefs of a misspecified patient agent can cycle when they would converge if the agent were correctly specified or myopic.
- Bohren and Hauser [2021] characterizes when actions converge with myopic agents and finite-support priors.
- Esponda, Pouzo, and Yamamoto [2021] uses stochastic approximation to characterize asymptotic action frequencies.
- Frick, Iijima, and Ishii [2021] characterizes global convergence.
- Other recent related work: Heidhues, Koszegi, and Strack [2018],[2021], Molavi [2019], He [2021].
- "Limit Points of Endogenous Misspecified Learning," FLS [2021] characterizes exactly which Berk-Nash equilibria can be limit points for general priors and possibly patient agents.

Persistent Misspecification

- He and Libgober [2020] looks at the competition between a misspecified model and the correctly specified model in games.
- Murooka and Yamamoto [2021] study games where all agents have the same misspecification.
- In Gagnon-Bartsch, Rabin, and Schwartzstein [2019] the agent's attention partition determines whether they become correctly specified.
- Montiel-Olea, Ortoleva, Pai, and Prat [2021] shows that misspecified agents with lower dimensional models initially have a higher willingness to pay for an object, while correctly specified agents have a higher long-run willingness to pay.
- "Which Misperceptions Persist," Fudenberg and Lanzani. Purely Bayesian agents can never realize they are misspecified. We use an evolutionary model to see which misperceptions can persist. Mutations that lead some agents to use a better-fitting model can yield lower payoffs and fail to spread.

Learning from Exogenous Data

- Y is a finite set of possible outcomes.
- $P = \Delta(Y)$ is the set of probability measures p over Y, endowed with the total variation norm.
- $\mu_0 \in \Delta(P)$ denotes a prior distribution over distributions of outcomes, and $\Theta = \operatorname{supp} \mu_0$ denotes its support.
- We don't require that all $p \in \Theta$ assign positive probability to all outcomes, or that they are mutually absolutely continuous.
- A data set $y^t = (y_1, y_2, \ldots, y_t) \in Y^t$ is a vector of outcomes. For every data set y^t we let μ_t be the posterior belief, which is required to satisfy Bayes rule on histories where it is defined:

$$\mu_t(C \mid y^t) = \frac{\int_{p \in C} \prod_{\tau=1}^{t} p(y_\tau) d\mu_0(p)}{\int_{p \in P} \prod_{\tau=1}^{t} p(y_\tau) d\mu_0(p)}. \qquad \text{(Bayes Rule)}$$

- The empirical distribution $f_t \in P$ is

$$f_t(z) = \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}_{y_\tau = z} \,.$$

- Let $H : P \times P \to \bar{\mathbb{R}}$ denote the (possibly infinite) KL divergence of q with respect to p:

$$H(q,p) = \sum_{z \in Y} q(z) \log \left(\frac{q(z)}{p(z)} \right) \,.$$

with the convention that $\frac{0}{0} = 0$ and $0 \log 0 = 0$.

- The KL minimizers for q are $M(q) = \operatorname{argmin}_{p \in \Theta} H(q,p)$.

- The log-likelihood of y^t under distribution p is

$$\log\left(\prod_{\tau=1}^{t} p(y_\tau)\right) = \sum_{z\in Y} t\, f_t(z)\log p(z)$$
$$= -tH(f_t,p) + t\sum_{z\in Y} f_t(z)\log f_t(z)\,.$$

- Thus $M(f_t)$ corresponds to the outcome distributions that maximize the likelihood of y^t.
- And the posterior odds ratio for a set C is

$$\frac{\mu_t(C)}{1-\mu_t(C)} = \frac{\int_{\pi\in C}\exp(-H(f_t,\pi)t)d\mu_0(\pi)}{\int_{\pi\notin C}\exp(-H(f_t,\pi)t)d\mu_0(\pi)}.$$

- $M_\varepsilon(f)$ is the set of distributions that come within ε of the minimum KL divergence:

$$M_\varepsilon(f) = \left\{p'\in\Theta : H(f,p') \le \min_{p\in\Theta} H(f,p) + \varepsilon\right\}\,.$$

Uniform Pathwise Concentration

- The prior μ_0 is ϕ positive if for $\phi : \mathbb{R}_{++} \to \mathbb{R}_{++}$, we have $\mu_0(B_\varepsilon(p)) \geq \phi(\varepsilon)$ for every $p \in P$ and $\varepsilon > 0$.

Theorem (Diaconis and Freedman 1990)

If μ_0 is ϕ positive, for every $\alpha \in (0,1)$ there is a function $\tilde{A}_\alpha : \mathbb{R}_{++} \to \mathbb{R}_{++}$ such that

$$\frac{\mu_t(B_\varepsilon(f_t))}{1 - \mu_t(B_\varepsilon(f_t))} \geq \tilde{A}_\alpha(\varepsilon) \exp(\alpha \varepsilon t),$$

for all $\varepsilon \in (0,1), t \in \mathbb{N}$, and $f_t \in \Delta(Y)$.

- In words, the beliefs concentrate around the empirical distribution at a uniform and exponential rate.
- Choosing larger α yields a better rate but a smaller multiplicative term; in the medium run and for small ϕ functions an intermediate α provides the best bound.

Limitations of ϕ positivity

- ϕ positivity rules out cases where the support of the prior
  - is finite
  - corresponds to a parametric model with dimension less than that of Y, or
  - rules out the true data generating process.
- We relax ϕ positivity: μ_0 is ϕ positive on Θ if for $\phi : \mathbb{R}_{++} \to \mathbb{R}_{++}$ we have $\mu_0(B_\varepsilon(p)) \geq \phi(\varepsilon)$ for every $p \in \Theta \equiv \operatorname{supp} \mu_0$ and every $\varepsilon > 0$.
- Let $\Delta^\Theta(Y)$ be the set of empirical distributions for which Bayes rule is well defined.
- And let $D_\varepsilon(p) = B_\varepsilon(M_\varepsilon(p))$ be the ε-ball around the ε-KL minimizers for p.

Theorem ("Uniform Pathwise Bounds")

If μ_0 is ϕ positive on Θ, then for every $\alpha \in (0,1)$, there is a function $A_\alpha : \mathbb{R}_{++} \to \mathbb{R}_{++}$ such that

$$\frac{\mu_t(D_\varepsilon(f_t))}{1-\mu_t(D_\varepsilon(f_t))} \geq A_\alpha(\varepsilon)\,\exp(\alpha\varepsilon t)\,,$$

for all $\varepsilon \in (0,1)$, $t \in \mathbb{N}$ and $f_t \in \Delta^\Theta(Y)$.

- Exponential convergence rate as in Diaconis Freedman, but now around the ε-KL minimizers $D_\varepsilon(f_t) = B_\varepsilon(M_\varepsilon(f_t))$.
- Beliefs needn't concentrate around the empirical frequency f_t because f_t might not be in the support of the prior.
- If Θ is convex points can strengthen this to the ε ball $B_\varepsilon(M(f_t))$ around the exact minimizers $M(f_t)$.
- But need convexity as otherwise points far from the minimizers can have almost the same divergence.

Proof Sketch

- For any $\alpha \in (0,1)$ there is a wedge of $\alpha\varepsilon$ in the KL divergence between distributions outside $D_\varepsilon(f_t)$ and those inside a sufficiently small ball around the exact KL-minimizers.
- We then use the lower semicontinuity of the KL divergence and a compactness argument to show that there is a bound for the size of this ball that holds uniformly over all empirical distributions f_t.
- This wedge guarantees the odds ratio $\frac{\mu_t(D_\varepsilon(f_t))}{1-\mu_t(D_\varepsilon(f_t))}$ grows at an exponential rate.
- Then the result follows from ϕ positivity.

Berk with a Convergence Rate

Theorem (Berk 1966)

If Θ is regular, then for all $\varepsilon \in (0,1)$

$$\lim_{t\to\infty} \mu_t(B_\varepsilon(M(p^*))) = 1 \qquad p^*\text{-a.s.}$$

- Here "regular" is a set of technical conditions that are satisfied in our setting if every $p \in \Theta$ is arbitrarily close to a $p' \in \Theta$ that has full support.

Theorem

If Θ is regular and μ_0 is ϕ positive on Θ, then for every $\varepsilon \in (0,1)$ with probability $1 - O(\exp(-t))$, $\mu_t(\Theta \setminus B_\varepsilon(M(p^*)))$ is $O(\sqrt{t}\exp(-Kt))$.

- This follows from the uniform pathwise bound and the CLT.

Learning from Endogenous Data

- Every period $t \in \mathbb{N}$, the agent chooses an action a from the finite set A.
- Action a has two consequences:
  - Induces objective probability distribution $p_a^* \in \Delta(Y)$;
  - Directly influences the agent's payoff through $u : A \times Y \to \mathbb{R}$.
- Now the agent's beliefs p are over action-dependent outcome distributions with components p_a.
- A (pure) policy $\pi : \bigcup_{t=0}^{\infty} A^t \times Y^t \to A$ specifies an action for every history $(a_\tau, y_\tau)_{\tau=0}^{t} = (a^t, y^t) \in A^t \times Y^t$.
- The policy is chosen to maximize expected discounted utility with discount factor $\beta \in [0, 1)$.
- $BR(\mu) = \arg\max_{a \in A} \int_P \mathbb{E}_{p_a}[u(a, y)]\, d\mu(p)$ is the set of myopic best replies to belief μ.

Berk-Nash Equilibrium

- For each action a, let $\hat{\Theta}(a) = \underset{p \in \Theta}{\operatorname{argmin}} H(p_a^*, p_a)$.

  This is the set of action-contingent outcome distributions in Θ that minimize the KL divergence relative to p_a^* when the agent plays a.
- Action a is a Berk-Nash equilibrium (Esponda and Pouzo [2016]) if there is a belief $\nu \in \Delta(\hat{\Theta}(a))$ such that *a* is myopically optimal given ν.
- Agents need not be very patient, so they may have little or no data about the consequences of some actions.
- Two outcome distributions $p, p' \in \Theta$ are observationally equivalent under action a if $p_a = p'_a$.
- Let $\mathcal{E}_a(p) \subseteq \Theta$ denote the outcome distributions in Θ that are observationally equivalent to p under a.
- While playing a the agent never updates the relative likelihood of the two distributions in the same equivalence class $\mathcal{E}_a(p)$.

Refinements of Berk-Nash Equilibrium

Definition (Uniform and Uniformly Strict Berk-Nash Equilibria)

Action a is a

(i) uniform Berk-Nash equilibrium if for every KL minimizing outcome distribution $p \in \hat{\Theta}(a)$, there is a belief over the observationally equivalent distributions $\nu \in \Delta\left(\mathcal{E}_a(p)\right)$ such that $a \in BR(\nu)$.

(ii) uniformly strict Berk-Nash equilibrium if $\{a\} = BR(\nu)$ for every belief in $\nu \in \Delta(\hat{\Theta}\left(a\right))$.

- When the agent is correctly specified (i.e. $p^* \in \Theta$), p_a^* is the unique KL minimizer for a, so

$$\text{Uniform B-N} = \text{B-N} = \text{Self-Confirming}.$$

- When are there multiple KL Miminimizers?

Technical Assumptions

- The prior μ_0 has subexponential decay: there is $\phi : \mathbb{R}_+ \to \mathbb{R}$ such that for every $p \in \Theta$ and $\varepsilon > 0$, $\mu_0(B_\varepsilon(p)) \geq \phi(\varepsilon)$ with

$$\lim \phi(K/n) \exp(n) = \infty \qquad \forall K > 0.$$

- A strengthening of ϕ-positivity.
- Priors with a density that is bounded away from 0 on their support, priors with finite support, and Dirichlet priors all have subexponential decay. Fudenberg, He, and Imhof [2017] show that Bayesian updating can behave oddly on priors w/o subexponential decay.
- **Simplyfing assumption for the talk:** For all $p \in \Theta$, p and p^* are mutually absolutely continuous. This guarantees that no conceivable distribution is ruled out after a finite number of observations.

Only Uniform-Berk Nash Equilibria are Limit Actions

Theorem (Limit Actions are Uniform Berk-Nash Equilibria)

If actions converge to a $\in A$ *with positive probability, a is a uniform Berk-Nash equilibrium.*

Proof Sketch

- The agent's belief concentrates around the KL minimizers from the empirical frequency at an exponential rate that is uniform over the sample realizations.
- While playing a, the empirical frequency converges to p_a^*.
- The difference between the empirical frequency and p_a^* is a random walk that oscillates towards the minimizers.
- By the Central Limit Theorem these oscillations die out at rate $\frac{1}{\sqrt{t}}$, which is slower than the exponential concentration of beliefs.
- So we can show that the beliefs concentrate around each minimizer infinitely often.
- If a is not a uniform B-N, this induces the agent to switch to another action.

Possible Non-convergence

- Nyarko [1991] shows by example that misspecified learning may not converge.

- Our theorem shows that if no equilibrium is uniform, actions cannot converge; this may be easier to check than directly verifying non-convergence.

- We show by example that uniform B-N equilibria need not exist.

- One case where they do exist is if the agent is correctly specified.

Uniform Stability

Definition (Stability)

A Berk-Nash equilibrium a is uniformly stable if for every $\kappa \in (0,1)$, there is an $\epsilon > 0$ such that for all initial beliefs $\nu \in \Delta(\Theta)$ such that $\nu(\hat{\Theta}(a)) > 1 - \epsilon$, the action prescribed by *any* optimal policy converges to $a \in A$ with probability greater than $1 - \kappa$.

Theorem (Characterization theorem)

Action $a \in A$ is uniformly stable if and only if it is a uniformly strict Berk-Nash equilibrium.

Note that this covers the case where the agent perceives an information value from experimentation, as in the example of Fudenberg, Romanyuk, and Strack [2017].

Proof Sketch

- If a is not a uniformly strict B-N there is *some* belief over minimizers such that a is not strictly optimal, there is an optimal policy where a is not the limit outcome.
- Conversely, if a uniformly strict, it is the unique myopic best reply to every action- contingent outcome distribution p in a ball around the KL minimizers $\hat{\mathcal{P}}(a)$.
- And for any discount factor $\beta \in (0,1)$ a is the unique optimal choice for beliefs in some smaller ball, as the probability of learning from other actions becomes negligible.
- We then transform the odds ratio between the non-KL minimizers and the minimizers to make it a positive supermartingale, as in Frick, Iijima, and Ishii [2021].
- Then we generalize the "active supermartingales" of Fudenberg and Levine [1993] to show that if this odds ratio starts sufficiently low, it is unlikely to increase enough for the agent to change their action.

Positive Attractiveness

Definition (Positively Attracting)

Action $a \in A$ is positively attracting if for every optimal policy π

$$\mathbb{P}_\pi \left[\lim_{t \to \infty} a_t = a \right] > 0 \,.$$

- Benaim and Hirsch [1999] show that linearly stable Nash equilibria are positively attractive under stochastic fictitious play.
- We show by example that uniformly strict BNE need not be positively attractive.
- Uniformly strict BNE *are* positively attracting in some cases of interest, such as when the agent believes that the distribution over outcomes is the same for all actions ("causation neglect") or with one-dimensional priors and a supermodularity property.

Which Misperceptions Persist?

- When will agents realize they are misspecified?
- Bayesian updating can't lead to a positive probability on a data generating process that lies outside the support of the prior.
- We propose an evolutionary criterion to evaluate the stability of misspecified Bayesian models.
- Each agent observes the actions and outcomes of the previous generation, and uses them to update beliefs within their subjective model.
- The agents then choose an action that is a best reply to these beliefs.
- Different agents may employ different subjective models, and the relative frequency of the models that induce better actions increases over time.
- The steady states of this process coincide with Berk-Nash equilibria.

- "Mutations" lead some agents to adopt an expanded subjective model with a larger support.
- They use the expanded model to make their inferences and choose their actions.
- We then ask whether the use of the expanded model will spread, or whether the existing model "resists mutations."
- Not all equilibria are unstable, because the share of mutants only increases if they do better than agents using the prevailing paradigm.
- We consider two ways that subjective models can expanded: "local" expansions to nearby subjective models, and "one-hypothesis relaxations" that drop one of the hypotheses that characterize the subjective model.
- We characterize stability in both cases.
- We apply the results to several common misspecifications.

Parametric Models and Inference

- Objective contingent outcome distributions $p^*(\cdot|\cdot) \in \Delta(Y)^A$.
- A subjective model for an agent is the set of parameters $\Theta \subseteq \mathbb{R}^k$ they consider possible, where each $\theta \in \Theta$ is associated with family of probability distributions $p_\theta(\cdot|a)$.
- Let $\psi \in \Delta(A)$ be a distribution over actions in the population.
- Given a distribution ψ let

$$H_\psi(p^*, p_\theta) = \sum_{a \in A} \psi(a) H\left(p^*(\cdot|a), p_\theta(\cdot|a)\right).$$

- Let $\Theta(\psi)$ denote KL minimizers for ψ:

$$\Theta(\psi) : = \operatorname*{argmin}_{\theta \in \Theta} H_\psi(p^*, p_\theta).$$

- These are the parameters that best fit the data generated by ψ.

State of the System

- There is a continuum of agents.
- The state of the system is a finite-support joint distribution $\pi \in \Delta(\mathcal{K} \times A) =: \Pi$ over the subjective models and actions of the agents.
- Each agent's posterior beliefs are supported on the KL-minimizing parameters in their Θ with respect to the distribution over actions in the last period.
- Agents play a best response to their beliefs.
- $\pi^{t+1}(\cdot|\Theta) \in \Delta(A)$ denotes the distribution over actions played at time $t+1$ by the agents with subjective model Θ when the previous state is π^t.

Evolutionary Dynamics

- We assume that the share of agents with a particular subjective model evolves according to a payoff monotone (Samuelson and Zhang [1992]) dynamic $T : \Pi \to \Delta(\mathcal{K})$: better-performing models increase their shares.
- In our model, the beliefs themselves are not inherited; what is inherited is how to use observables to reach a conclusion.
- The offspring then learn from the data of the previous period.
- A solution of the system is a sequence of states where the shares of the subjective models evolve according to T, and agents play best replies to the KL-minimizers given the previous period's data.
- A steady state is a fixed point of that process, i.e. a constant solution.

Steady States and Mutations

Lemma

For all $\Theta \in \mathcal{K}$ *and* $\psi \in \Delta(A)$, $(\delta_\Theta \times \psi)$ *is a steady state if and only if* (Θ, ψ) *is a Berk-Nash equilibrium.*

Now we consider "mutations" that lead some agents to change their subjective model.

Definition

$\bar{\pi}$ is an ε mutation of a steady state $\delta_\Theta \times \psi$ to $\Theta' \supseteq \Theta$ if

(a) its marginal over subjective models is $(1-\varepsilon)\delta_\Theta + \varepsilon\delta_{\Theta'}$ and

(b) $\bar{\pi}(\cdot|\tilde{\Theta}) \in \Delta(BR(\Delta(\tilde{\Theta}(\psi))))$

Resistance to Mutations

Definition

A Berk-Nash equilibrium (Θ, ψ) resists invasion by Θ' if, after a sufficiently small mutation, the aggregate behavior of the population converges back to ψ.

Definition (Improving Mutations)

The ε mutation of a steady state $\delta_\Theta \times \psi$ to Θ' is improving if Θ' allows a lower KL divergence w.r.t. ψ than Θ does.

Proposition

A Berk-Nash equilibrium resists every mutation that is not explanation improving: A more open-minded model can destabilize an equilibrium only if it better explains the equilibrium distribution.

Local Mutations

Definition (Local Mutations)

The ε expansion of Θ is $\Theta_\varepsilon = \{\theta' \in \mathbb{R}^k : \exists \theta \in \Theta, ||\theta - \theta'|| \leq \varepsilon\}$.

Definition

A Berk-Nash equilibrium (Θ, ψ) resists local mutations if it resists invasion by every sufficiently small ε expansion of Θ.

- In a uniformly strict equilibria, the equilibrium action a is a strict best response to every KL-minimizing parameter.
- So after a small mutation beliefs are concentrated on a neighborhood where the unique best reply is still a.

Proposition

Every uniformly strict Berk-Nash equilibrium resists local mutations.

Most Improving Parameters

- Uniformly strict Berk-Nash equilibria need not exist.
- At other Berk-Nash equilibria, multiple strategies are a best reply to some belief over KL-minimizers.
- Only some of these strategies are best replies to the KL minimizers of the expanded model.
- The relative performance of these best replies and the equilibrium is what determines resistance to local mutations.
- The most improving parameters at a steady state are

$$\mathcal{M}_{\Theta,\psi}(\varepsilon) = \underset{\theta \in \Theta_\varepsilon}{\operatorname{argmin}} \, H_\psi \left(p^*, p_\theta\right).$$

  These parameters make the greatest local improvement in the explanation of the equilibrium data.
- Following an ε expansion, the mutants' beliefs will be concentrated on these parameters.

Characterization

Proposition

- *If all best replies to the most improving parameters at an equilibrium induce a higher payoff, the equilibrium does not resist local mutations.*
- *If the equilibrium is quasi-strict and some best reply to the the most improving parameters give a lower payoff, the equilibrium resists local mutations.*

- An equilibrium is quasi-strict if all the best replies are played with positive probability.
- If the equilibrium is not quasi-strict, the feedback gathered from a tiny fraction of mutated agents playing a more revealing action that is not used in equilibrium may change the behavior of the old population, even if they were performing better than the mutants.

Monopoly Pricing and Linear Demand

- A monopolist faces demand function $y = i^* - l^* a + \omega$.
- a is the price chosen by the monopolist and ω is a standard normal shock.
- The monopolist's payoff is $u(a, y) = ay$.
- It is uncertain about the the intercept $i \in \mathbb{R}$ and slope $l \in \mathbb{R}$ of the demand function.
- The true values of the parameters are $(i^*, l^*) = (42, 4)$.
- The monopolist has two actions, $A = \{2, 10\}$.
- Objectively optimal to use action 2.

Nyarko [1994]

- The quasi-strict equilibrium $\psi(2) = \frac{1}{5}$ is unstable to local mutations because the most improving parameter relaxes the unique binding constraint, allowing for a larger slope, which induces the (optimal) low price.

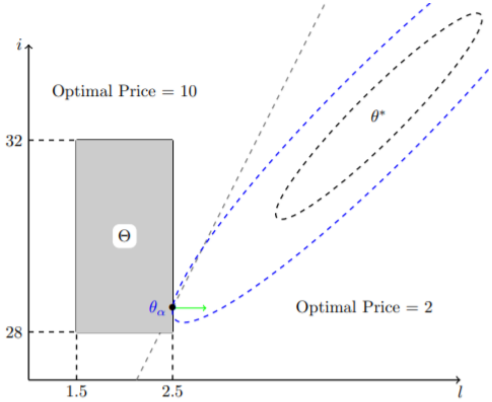

(Green arrow points towards greatest KL improvement.)

Esponda and Pouzo [2016]

- The quasi-strict equilibrium $\psi(2) = \frac{35}{36}$ resists local mutations because the most improving parameter has a larger intercept, which induces the (suboptimal) high price.

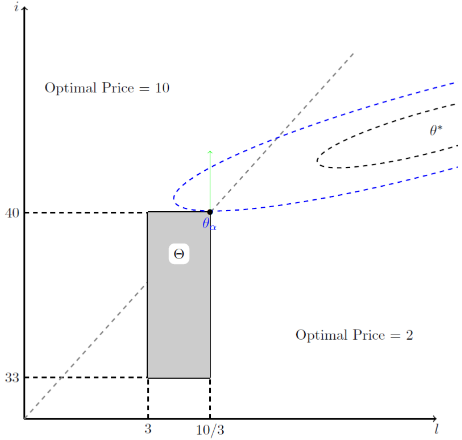

(Green arrow points towards greatest KL improvement.)

One-hypothesis Mutations

- We also consider agents whose subjective model is described by a finite collection of hypotheses about the underlying parameter.
- Quantitative statements like:
  - Restrictions on the possible values of one of the dimensions of the parameter, e.g. an overconfident agent who is sure that their skill is higher than a threshold.
  - Joint restrictions on the parameters, as independence between two variables.
- The hypotheses describe the parts of the agent's model that can be separately relaxed by a mutation.
- Formally, there is a finite collection of continuous functions $(f_i)_{i=1}^m$, $f_i : \mathbb{R}^k \to \mathbb{R}$ such that $\Theta = \{\theta : f_i(\theta) \geq 0, \forall i \in \{1, ..., m\}\}$.

Definition

- The subjective model Θ^l is a one-hypothesis relaxation of

$$\Theta = \{\theta \in \mathbb{R}^k : f_i(\theta) \geq 0, \forall i \in \{1, ..., m\}\}$$

if

$$\Theta' = \{\theta \in \mathbb{R}^k : f_i(\theta) \geq 0, \forall i \in \{1, ..., m\} \setminus \{l\}\}.$$

- A Berk-Nash equilibrium (Θ, ψ) resists one-hypothesis mutations if it resists invasion by every one-hypothesis relaxation.

- Several collections of hypotheses can describe the same Θ, and they can have different sets of one-hypothesis relaxations.
- This is natural, as the hypotheses are part of the agents' model of the world.

Taxation and Overshooting

- An agent chooses effort $a \in A = \{3, 4, 5\}$ at cost $c(a) = 2a/3$, and obtains income $z = a + \omega$, where $\omega \sim N(0, 1)$.
- The agent pays taxes $x = \tau^*(z)$, where τ^* has two income brackets, and the higher one is heavily taxed:

$$\tau^*(z) = \begin{cases} z/6, & \text{if } z \leq 16/3 \\ \frac{11}{12}z - 4, & \text{if } z \geq 16/3. \end{cases}$$

.

- The agent's payoff is $u(a, (z, x)) = z - x - c(a)$, so their optimal action is 4.
- Their subjective model of the tax schedule is quadratic with random coefficients:

$$\tau_\theta(z) = (\theta_1 + \eta)z + (\theta_2 + \eta)z^2,$$

where η is a standard normal.

- The agent observes $y = (z, x)$ at the end of each period.
- The original paradigm is that the tax schedule is linear.
- Given any action a, the KL-minimizing parameter treats the expected marginal rate as the actual average rate.
- The unique pure Berk-Nash equilibrium is uniformly strict and has too much effort.

- The agent observes $y = (z, x)$ at the end of each period.
- The original paradigm is that the tax schedule is linear.
- Given any action a, the KL-minimizing parameter treats the expected marginal rate as the actual average rate.
- The unique pure Berk-Nash equilibrium is uniformly strict and has too much effort.
- An agent who drops the linearity assumption estimates a very high quadratic term, because most realized income levels will be near the shift point between brackets.
- They extrapolate this progressivity as a global feature which leads them to choose the minimal action 3.
- The mutated agent overshoots the optimum, and the equilibrium resists to "one-hypothesis" mutations.

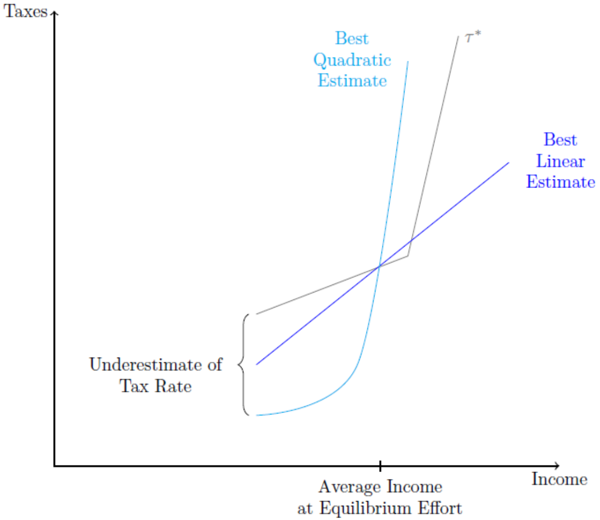

Figure: Misspecified Taxation Schedule

Additive Lemons and Cursed Equilibrium, Esponda [2008]

- We extend the model to allow the agent to observe a signal at the start of each period.
- We use this extension to explore the persistence of cursed equilibria in a lemons problem where the buyer believes that the seller's ask price and value are independent.
- The agent is a buyer with valuation $v = \omega + 5 + s$.
- Seller who owns the object and values it at ω.
- The signal s is a mean-zero shock independent of ω.
- The mechanism used is double action with price at the buyer's bid, so the seller sets their bid x equal to their value.
- The value ω is $\omega = 3$, with probability $1/2$, $\omega = 2$ with probability $1/4$ and $\omega = 1$ with probability $1/4$.
- The value is observed only if a sale occurs.

- A parameter is a probability distribution over prices, and a conditional distribution over values given the price.

- The true distribution of values conditional on an ask price depends on the price.

- However, the agent believes that the price and the value are independent.

Aligned Preferences

- Suppose the distribution of s is a point mass on 0.
- Then the objectively optimal strategy is to bid 3.
- Bidding 2 is a Berk-Nash equilibrium.
- The KL-minimizing parameter is an independent joint probability distribution that is correct about the distribution of seller bids.
- A mutation that drops the restriction that high values have the same probability after each price has no effect: Because 3 is not accepted, the mutated agents cannot infer that there is correlation at the high price level, and so they do not increase their bid.

- Suppose s instead that is uniform over $\{-1, 1\}$.
- The optimal strategy is still to always bid 3 after every signal.
- The strategy $a = 2 + s$ is a Berk-Nash equilibrium— now the agent sometimes bids 3.
- The KL-minimizer is again an independent joint probability distribution that is correct about the distribution of seller bids.
- Since the value is observed only when a transaction occurs, the observed distribution over values is too pessimistic, which leads the agent to bids 1 after signal -1.
- The equilibrium does not resist one-hypothesis mutations: a mutation that drops the restriction that high values have the same probability after each price lets the agent realize the high price is correlated with high value.
- Realizing this leads the agent to always bid high $(a = 3)$.
- This did not happen without the noise since in that case the agent never makes the high bid.

Ongoing Research: Selective Memory

- Here we suppose that the agent's memory is distorted through a memory function $m_s : (Y \times A \times S) \to [0, 1]$, which gives the probability that the agent remembers the outcome, action, signal triplet (y, a, s) when they observe signal s'.
- We assume the agent is unaware of their selective memory and updates beliefs naively using Bayes rule.
- Past work (e.g. Mullinaithan [2002], Kahana [2012]. Bordalo et al [2017]) focuses on the 1-step-ahead implication of selective memory.
- We define a notion of "selective memory equilibrium" and extend our concentration results to show that is is a necessary condition for long-run outcomes.
- Now we are looking at various types of selective memory such as associativeness and confirmation bias to see what we can say about them.

Thanks!

When are there Multiple KL Miminimizers?

- For "generic priors" there is a unique KL minimizer for any distribution p_a^*.
- But symmetry or parametric restrictions are not generic.
- Example: suppose that y is the color of the ball drawn from an urn which is known to contain 6 balls.
- The agent correctly believes their action doesn't affect y.
- Outcome distributions correspond to the urn composition.
- The agent is certain that at most half of the balls have the same color, i.e., that $p(y) \leq 1/2$ for every y.
- In reality the urn has 4 white balls, 1 red, and 1 blue.
- So the two KL minimizers are (3 white, 2 blue, 1 red) and (3 white, 1 blue, 2 red). Text

Infinitely Many Oscillations

- Fix a KL-minimizer q_a for p_a^*. Let E_t be the event in which q_a is a KL minimizers for f_t.
- We can show that the correlation between being a minimizer at time t, and being a minimizer at time s, is not too large in the sense

$$\liminf_{t \to \infty} \frac{\sum_{s=1}^{t} \sum_{r=1}^{t} \mathbb{P}[E_s \text{ and } E_t]}{\left(\sum_{s=1}^{t} \mathbb{P}[E_s]\right)^2} > 0$$

- It then follows from the Kochen-Stone lemma that the probability that infinitely many of the E_t will realize is strictly positive.
- The event "infinitely many E_t realize" is invariant under finite permutations so the Hewitt–Savage 0-1 law implies that the probability must equal one. Text

Payoff Monotonicity

- $\pi_{\mathcal{K}}^{t+1}(\Theta) = T(\pi^t)(\Theta)$, where T is continuous and such that the dynamic is payoff monotone, meaning that

$$\frac{U^*(\pi^t(\cdot|\Theta))}{U^*(\pi^t(\cdot|\Theta'))} > (=)1 \implies \frac{T(\pi)(\Theta)}{T(\pi)(\Theta')} = \frac{\pi_{\mathcal{K}}^{t+1}(\Theta)}{\pi_{\mathcal{K}}^{t+1}(\Theta')} > (=)\frac{\pi_{\mathcal{K}}^{t}(\Theta)}{\pi_{\mathcal{K}}^{t}(\Theta')}. \tag{1}$$

Back to slides

Solution

Definition

A sequence $(\pi^t)_{t\in\mathbb{N}_0} \in \Pi^{\mathbb{N}_0}$ is a solution if satisfies equations

$$\pi^{t+1}(\cdot|\Theta) \in \Delta(BR(\Delta(\Theta(\pi^t_\Pi))))$$

and

$$\frac{U^*(\pi^t(\cdot|\Theta))}{U^*(\pi^t(\cdot|\Theta'))} > (=)1 \implies \frac{\pi^{t+1}_{\mathcal{K}}(\Theta)}{\pi^{t+1}_{\mathcal{K}}(\Theta')} > (=)\frac{\pi^t_{\mathcal{K}}(\Theta)}{\pi^t_{\mathcal{K}}(\Theta')}.$$

for all $t \in \mathbb{N}_0$.

Back to slides

Steady State

Definition

A *steady state* is a $\hat{\pi} \in \Pi$ such that $(\hat{\pi})_{t \in N_0}$ is a solution and $\hat{\pi}_{\mathcal{K}} = \delta_{\Theta}$ for some $\Theta \in \mathcal{K}$. A steady state is *unitary* if $\hat{\pi}(\cdot|\Theta) \in \Delta(BR(\mu))$ for some $\mu \in \Delta(\Theta(\hat{\pi}_A))$.

Back to slides

Resistance to Mutations

Definition

A Berk-Nash equilibrium (Θ, ψ) *resists mutation to* Θ' if there is a collection of solutions $(\pi_\varepsilon^t)_{t\in\mathbb{N}_0,\varepsilon\in(0,1)}$, such that π_ε^0 is the ε mutation of $\delta_\Theta \times \psi$ to Θ', and

$$\lim_{\varepsilon\to 0}\lim_{t\to\infty}(\pi_\varepsilon^t)_A = \psi.$$

Back to slides