Dynamic Concern for Misspecification

Giacomo Lanzani

Harvard

August 2024

Motivation

- Misspecification—beliefs that rule out the true data generating process— is a pervasive issue for economic agents.
  - A central bank may consider different model economies of the data-generating process for output and inflation, but none may be correct.
  - Agents in a social network may find it hard to correctly account for redundant sources of information.
  - Misspecification is even more relevant when dealing with entirely novel issues, such as those related to human impact on climate change.

- Fast-growing literature studies the consequence of Bayesian learning with misspecified models (Esponda and Pouzo, 2016 and the follow-up work).
- Misspecified learning as a unifying explanation for many biases (overconfidence, failures to understand regression to the mean, suboptimal behavior in the face of complex tax schedule, lemon's problems).
- But this literature also assumes that the agents always ignore the possibility of being misspecified.
- Normatively unappealing because misspecification has a strong predicted impact on learning problems.
- Also descriptively unrealistic, as some models of uncertainty/ambiguity-averse preferences embody incomplete trust in a single probability, and we have evidence in favor of it in many settings.

My approach

- I propose a model of agents who:
  - Learn about the data generating process from the consequences of their actions;
  - Hedge against misspecification in a way that generalizes the robust control preferences of Hansen and Sargent (2001);
  - Adaptively adjust their misspecification concern as a function of how well their subjective model explained past data.

Outline of Results

- There exists a unique "statistical" way to adjust misspecification concern with good performance under any possible DGP.

- Characterize the limit actions for different attitudes towards model failures, providing a learning foundation for different decision criteria under uncertainty.

- Applications to misperception of tax schedules (Rees-Jones and Taubinsky, 2020) and cyclical monetary policies (Sargent, 2008,2009).

Static Decision Problem

- Agent who repeatedly chooses from a finite number of actions $a \in A$.
- Consequences $y \in Y$.
- Utility index $u : A \times Y \to \mathbb{R}$ over actions and consequences.
- Each $a \in A$ induces an objective probability distribution $p_a^* \in \Delta(Y)$.
- Agent knows each period consequence only depends on the current action.
- The agent does not know $p^* = (p_a^*)_{a \in A}$ and deals with this uncertainty in a quasi-Bayesian way.

Purely Bayesian Benchmark

- The agent's models are described by parameters Θ, and a prior belief μ_0 with support Θ.
- Each $\theta \in \Theta$ is associated with a distribution $q^\theta = \left(q_a^\theta\right)_{a \in A} \in \Delta\left(Y\right)^A$.
- The agent is *correctly specified* if there exists $\theta \in \Theta$ such that $q^\theta = p^*$.
- The actions that maximize the SEU of an agent with belief μ are:

$$BR^{SEU}\left(\mu\right) = \operatorname{argmax}_{a \in A} \int_\Theta \mathbb{E}_{q_a^\theta}\left[u\left(a, y\right)\right] \mathrm{d}\mu\left(\theta\right).$$

Average Robust Control

- Agent is concerned that none of these models is exact but only an approximation, with concern measured by $\lambda \geq 0$.
- For $p, p' \in \Delta(Y)$, let $R(p||p')$ be the relative entropy between p and p'.
- The agent evaluates action a with the *average robust control* criterion of Cerreia-Vioglio, Hansen, Maccheroni, Marinacci (2022):

$$\int_\Theta \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a}\left[u\left(a,y\right)\right] + \frac{R\left(p_a||q_a^\theta\right)}{\lambda} \right) d\mu\left(\theta\right).$$

- Hansen and Sargent (2001)'s robust control: μ is a Dirac on (the usually correct) θ^*: rational expectations + robustness.
- The nondegenerate average over models captures the fact that the agent hasn't yet discovered the true model.

Reaction to Information

- We want the agent's beliefs and concern for misspecification to adapt to the received information.
- Belief $\mu \in \Delta(\Theta)$ updated to $\mu(\cdot|h_t)$ through standard Bayesian updating.
- Concern for misspecification is a function $\boldsymbol{\lambda}(h_t)$ of how well the agent's model explained the current history h_t.
- The relation between $\boldsymbol{\lambda}(h_t)$ and the average log-likelihood ratio of the model (LLR) turns out to be crucial.

Two Desiderata: Safety and No Regret

- **Safety:** The limit time-average payoff should be at least what the agent can guarantee against *every* possible data generating process.
- Mild but it has a significant bite under misspecification: often a Bayesian SEU agent fails it, a lot of criticism for misspecified learning comes from this.
- **No Regret Under Correct Specification:** If the agent is correctly specified limit time-average payoff should converge to the payoff of the objectively best action.
- Can be interpreted as requiring that the long run probability of Type I error goes to 0.

Theorem (Informal)

1. *For every decision problem and there exists a* $\bar{c} \geq 0$ *such that for all* $c \geq \bar{c}$*, the average robust control with*

$$\boldsymbol{\lambda}(h_t) = c\frac{LLR(h_t, \Theta)}{t}$$

is safe and has no regret under correct specification.

2. *There are decision problems for which there are no safe and no regret under correct specification average robust control such that*

$$o\left(\boldsymbol{\lambda}(h_t)\right) = \frac{LLR(h_t, \Theta)}{t} \quad \text{or} \quad \boldsymbol{\lambda}(h_t) = o\left(\frac{LLR(h_t, \Theta)}{t}\right)$$

- It is always safe and no regret under correct specification to keep the concern proportional to the LLR.
- Any rule that is globally more demanding or lenient in the evaluation of the model performance fails one of these conditions in some decision problem.

- I refer to $\boldsymbol{\lambda}(h_t) = c\frac{LLR(h_t,\Theta)}{t}$ as "statistically sophisticated."

- $o\left(\boldsymbol{\lambda}(h_t)\right) = \frac{LLR(h_t,\Theta)}{t}$, is a "demanding type", sort of a belief in the Law of Small Numbers, if the empirical frequency doesn't quickly match the theoretical distribution they grow suspicious of the model.

- $\boldsymbol{\lambda}(h_t) = o\left(\frac{LLR(h_t,\Theta)}{t}\right)$ is a "lenient type" that attributes too much of the unexplained evidence to sampling variability.

Equilibrium Concepts

- Let $\Theta(a) = \operatorname{argmin}_{\theta\in\Theta} R\left(p^*_a || q^\theta_a\right)$ be the set of distributions that best fit the true data generating process while action a is played.

Definition

Action a^* is a:

1. *Berk-Nash equilibrium* if there exists $\nu \in \Delta(\Theta)$ with

$$a^* \in BR^{SEU}(\nu), \qquad \operatorname{supp}\nu \subseteq \Theta(a^*).$$

2. *Maxmin equilibrium* if a^* is the maxmin best reply to the distributions that are absolutely continuous wrt some $\{q_\theta : \theta \in \Theta(a^*)\}$.
3. *c-robust equilibrium* if there exists $\nu \in \Delta(\Theta)$ with

$$a^* \in BR^{\lambda}(\nu), \quad \operatorname{supp}\nu \subseteq \Theta(a^*), \quad \lambda = c \min_{\theta\in\Theta} R\left(p^*_{a^*} || q^\theta_{a^*}\right).$$

Taxonomy Under Misspecification

- Action a is a $\boldsymbol{\lambda}$-*limit action* if there is a $\boldsymbol{\lambda}$-optimal policy Π such that $\mathbb{P}_{\Pi}\left[\sup\{t\colon a_t \neq a\} < \infty\right] > 0$.

Theorem

Suppose that the agent is misspecified and a^* is a $\boldsymbol{\lambda}$-limit action. We have:

1. *If $\boldsymbol{\lambda}$ is lenient, then a^* is a Berk-Nash equilibrium (**SEU**).*
2. *If $\boldsymbol{\lambda}$ is demanding, then a^* is a **maxmin** equilibrium.*
3. *If $\boldsymbol{\lambda}(h_t) = c\frac{LLR(h_t,\Theta)}{t}$ then a^* is a c-**robust** equilibrium.*

Taking Stock

- Learning foundation for different uncertainty attitudes.
- Positive correlation between misspecification and uncertainty aversion.
- Positive correlation between uncertainty aversion and belief in the Law of Small Numbers.
- These relations are causal: misspecification and belief in the Law of Small Numbers induce more uncertainty aversion.
- Long-run uncertainty aversion is higher under convergence to an action with consequences that are less well predicted by the models.

Conclusion

- Model endogenous misspecification concern and establish a normative benchmark.
- Characterize the limit actions and give a learning foundation for different decision criteria under uncertainty.
- Paper shows that the model is consistent with response to tax schedules and monetary policy cycles.

Proof Sketch

- **Step 1**: Show that on almost all histories if the empirical action frequency converges to α then

$$\frac{LLR(h_t,\Theta)}{t} - \min_{\theta\in\Theta}\sum_{a\in A}\alpha\left(a\right)R\left(f_a^t||q_a^\theta\right) \to 0,$$

  where f_a^t is the empirical outcome frequency after action a.
- **Step 2**: Prove that although the $\left(R\left(\cdot|q_a^\theta\right)\right)_{\theta\in\Theta}$ are not continuous (so cannot apply maximum theorem)

$$\min_{\theta\in\Theta}\sum_{a\in A}\alpha_t\left(a\right)R\left(f_a^t||q_a^\theta\right) \to \min_{\theta\in\Theta}\sum_{a\in A}\alpha(a)R\left(p_a^*||q_a^\theta\right) \qquad a.s.$$

- So $\boldsymbol{\lambda}$ converges to 0 (lenient) to a finite positive number (statistically sophisticated), infinity (demanding).

- **Step 3**: Generalization of Berk (1966, exogenous action) and Esponda and Pouzo (2016, finite Y) to show beliefs must concentrate on $\Theta(a)$.

- **Step 4**: Extend the result of Maccheroni, Marinacci, and Rustichini (2006) on the limits for λ to 0 or ∞ to allow for:
    - Evaluation of continuous rather than finite range utility;
    - (Infinite) average of robust control evaluations instead of single one;
    - Convergent but time-changing weights in the average.