A Multi-Agent Model of Misspecified Learning with Overconfidence

Cuimin Ba & Alice Gindin

University of Pennsylvania

June 24, 2021

Motivation

- Overconfidence is an important psychological bias.
- It is well-documented that overconfident individuals commit fundamental attribution errors in learning (Langer and Roth, 1975; Miller and Ross, 1975; Ross and Sicoly, 1979; Campbell and Sedikides, 1999).
    - Failures attributed to others or the outside environment
    - Successes attributed to own ability or hard work
- **How do *multiple* overconfident agents learn when interacting repeatedly?**

Related Literature

- Psychological evidence on overconfidence:
  - Langer and Roth (1975); Svenson (1981); BenoÃ®t, Dubra, and Moore (2015).
- Economic implications of overconfidence:
  - Camerer and Lovallo (1999); Gervais and Goldstein (2007); Anderson, Brion, Moore, and Kennedy (2012); Compte and Postlewaite (2012).
- Misspecified learning:
  - Esponda and Pouzo (2016), Fudenberg, Romanyuk and Strack (2017), Bohren and Hauser (2019), Heidhues, Koszegi, and Strack (2018a, 2019), He (2019), Esponda, Pouzo, and Yamamoto (2020), Murkuro and Yamamoto (2021).

Setup

- Two agents, $i \in \{1,2\}$; discrete time $t \in \{1,2,3,...\}$.
- Agent i chooses effort $e_t^i \in [\underline{e}, \overline{e}]$ in period t.
- Agent i obtains an individual payoff

$$q_t^i = Q^i(e_t^i, e_t^j, a^i, \phi) + \epsilon_t^i.$$

  - ϵ_t^i is independently drawn from F (density f) with zero mean.
  - a^i is agent i's ability which he is *dogmatic* about.
  - ϕ is a common fundamental he is *uncertain* about.
    - e.g. quality of a joint idea, market demand, etc.
  - Q^i is twice continuously differentiable, strictly concave in e^i, increasing in a^i and ϕ.
    - Assume $Q_a^i, Q_\phi^i > 0$ and the signs of $Q_{e^i a}^i$ and $Q_{e^i \phi}^i$ are different.

Adding Misspecification

- A^i is the *true* ability, $A^i \in [\underline{a}, \overline{a}]$.
- Φ is the *true* fundamental, drawn from full-support continuous density π_0^i with potentially unbounded support $\left[\underline{\phi}, \overline{\phi}\right]$.
- However, agent i has a degenerate belief that his ability is $\tilde{a}^i \in [\underline{a}, \overline{a}]$.
  - If $\tilde{a}^i > A^i$ the agent is overconfident (focus of talk).
  - If $\tilde{a}^i = A^i$ the agent has a correct belief.
  - If $\tilde{a}^i < A^i$ the agent is underconfident (in paper).

Adding Misspecification

- A^i is the *true* ability, $A^i \in [\underline{a}, \overline{a}]$.
- Φ is the *true* fundamental, drawn from full-support continuous density π_0^i with potentially unbounded support $\left[\underline{\phi}, \overline{\phi}\right]$.
- However, agent i has a degenerate belief that his ability is $\tilde{a}^i \in [\underline{a}, \overline{a}]$.
  - If $\tilde{a}^i > A^i$ the agent is overconfident (focus of talk).
  - If $\tilde{a}^i = A^i$ the agent has a correct belief.
  - If $\tilde{a}^i < A^i$ the agent is underconfident (in paper).
- $\tilde{a}^i$ is common knowledge.
- Each agent i has non-degenerate belief μ^i about the other agent's true ability A^j.

Learning about ϕ

- Agent i starts out with the correct prior π_0^i over ϕ and μ_0^i over A^j and performs Bayesian updating.
- After observing q_t^i, q_t^j, e^i, and e^j he updates to π_t^i and μ_t^i.

- He observes $q_t^i = Q^i(e_t^i, e_t^j, A^i, \Phi) + \epsilon_t^i,$
  while expecting to see $Q^i\left(e_t^i, e_t^j, \tilde{a}^i, \phi\right) + \epsilon_t^i,$ where $\phi \sim \pi_{t-1}^i$.
- Assume that $\forall \boldsymbol{e}, \tilde{a}^i, \Phi$, there always *exists* a $\phi \in \left(\underline{\phi}, \overline{\phi}\right)$ s.t. the no-gap condition holds,

$$Q^i\left(e^i, e^j, \tilde{a}^i, \phi\right) = Q^i\left(e^i, e^j, A^i, \Phi\right).$$

Choosing Actions

- Agents are *myopic*
- Since both agents observe $\{q_t^i, e_t^i\}_{i=1,2}$, posteriors π_t^i, μ_t^i are common knowledge.
- Players use iterated elimination of dominated strategies to determine their actions.
  ▸ Conditions for DS
- Since Q^i is increasing in a^i and ϕ, overconfidence leads to an underestimation of the fundamental.
  - Output underperforms expectations which players missattribute to a worse than expected fundamental.
- The cross derivative conditions imply that overconfidence and underestimation push effort in the same direction.

Characterizing the Berk-Nash equilibrium

- Agents will converge to a Berk-Nash equilibrium (Esponda & Pouzo, 2016):
  - Optimality: strategies are optimal given equilibrium beliefs.
  - Consistency: equilibrium belief minimizes the relative entropy (Kullback-Leibler divergence) ▸ details.

Characterizing the Berk-Nash equilibrium

- Agents will converge to a Berk-Nash equilibrium (Esponda & Pouzo, 2016):
  - Optimality: strategies are optimal given equilibrium beliefs.
  - Consistency: equilibrium belief minimizes the relative entropy (Kullback-Leibler divergence) ▸ details.

no-gap condition: $Q^i\left(e_\infty^i, e_\infty^j, A^i, \Phi\right) - Q^i\left(e_\infty^i, e_\infty^j, \tilde{a}^i, \phi_\infty^i\right) = 0, \forall i.$

Characterizing the Berk-Nash equilibrium

- Agents will converge to a Berk-Nash equilibrium (Esponda & Pouzo, 2016):
  - Optimality: strategies are optimal given equilibrium beliefs.

$$Q^i_{e^i}\left(e^i_\infty, e^j_\infty, \tilde{a}^i, \phi^i_\infty\right) = 0$$

  - Consistency: equilibrium belief minimizes the relative entropy (Kullback-Leibler divergence) details.

$$\text{no-gap condition: } Q^i\left(e^i_\infty, e^j_\infty, A^i, \Phi\right) - Q^i\left(e^i_\infty, e^j_\infty, \tilde{a}^i, \phi^i_\infty\right) = 0, \forall i.$$

- In our equilibrium, players have correct understanding of the game structure and correct beliefs about the other player's strategy.

Characterizing the Berk-Nash Equilibrium

Lemma 1

There exists at least one equilibrium and each equilibrium takes the form of $(\boldsymbol{e}_\infty, \boldsymbol{\phi}_\infty)$ as described previously.

Assumption 1

There is a unique equilibrium. sufficient conditions

- Denote the equilibrium beliefs given $\tilde{\boldsymbol{a}}$ as $\boldsymbol{\phi}_\infty(\tilde{\boldsymbol{a}})$ and actions as $\boldsymbol{e}_\infty(\tilde{\boldsymbol{a}})$.

Informational Externalities

Definition

1. Agent j creates an *informational externality* for agent i when e^j changes agent i's inference about ϕ

2. The informational externality is
   1. Positive if e^j and e^i are complements and affect i's inference in the same way
   2. Negative if e^j and e^i are substitutes and affect i's inference in opposite ways
   3. *Neither positive or negative* otherwise.

Informational Externalities

Definition

1. Agent j creates an *informational externality* for agent i when e^j changes agent i's inference about ϕ, or equivalently, at least one of $Q^i_{e^j a}, Q^i_{e^j \phi}$, or $Q^i_{e^i e^j}$ is nonzero

2. The informational externality is
   1. Positive if e^j and e^i are complements and affect i's inference in the same way

      $$Q^i_{e^i e^j} \geq 0, \operatorname{sgn}(Q^i_{e^i x}) = \operatorname{sgn}(Q^i_{e^j x}), \forall x = \phi, a.$$

   2. Negative if e^j and e^i are substitutes and affect i's inference in opposite ways

      $$Q^i_{e^i e^j} \leq 0, \operatorname{sgn}(Q^i_{e^i x}) \neq \operatorname{sgn}(Q^i_{e^j x}), \forall x = \phi, a.$$

   3. *Neither positive or negative* otherwise.

Comparison: Single Agent v.s. Two Agents

- Consider a single-agent benchmark:
  - Fix agent j's action at $e_S^j \equiv e_\infty^j(\boldsymbol{A})$ but allow agent i to optimize his action.
  - Let $\boldsymbol{\phi_S} = \left(\phi_S^1, \phi_S^2\right)$ denote the steady-state beliefs about fundamentals.
  - Let $\boldsymbol{e_S} = \left(e_S^1, e_S^2\right)$ denote the steady-state actions.
- Assume positive informational externalities
  - e^i and e^j are complements.
  - e^i and ϕ are complements.

Single-Agent Learning (Heidhues et al., 2018)

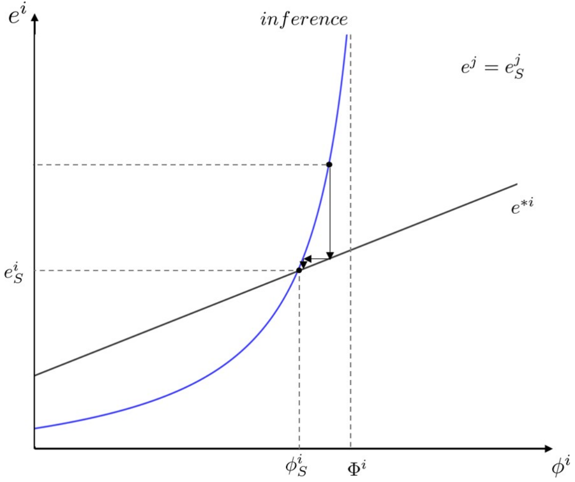

Figure: Single-agent learning.

Two Agents: Mutually-Reinforcing Learning

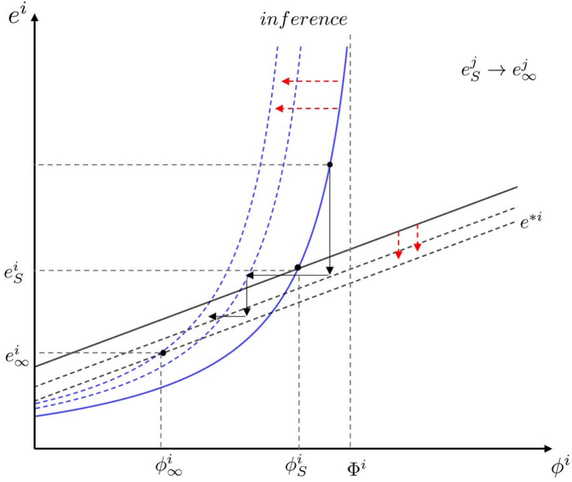

Figure: Mutual-reinforcing learning.

Mutually-Reinforcing Learning

Proposition 1 - Positive Info Ext + Mutual Learning = Worse Inferences

Suppose both agents create positive informational externalities.

1. Then both agents' underestimation of their fundamentals is reinforced when agent j is free to optimize and not reinforced when agent j's action is fixed at the level $e^j_S = e^j_\infty(\boldsymbol{A})$.

   That is, $\boldsymbol{\phi_\infty} < \boldsymbol{\phi_S} < \boldsymbol{\Phi}$. simulation

Mutually-Reinforcing Learning

Proposition 1 - Positive Info Ext + Mutual Learning = Worse Inferences

Suppose both agents create positive informational externalities.

1. Then both agents' underestimation of their fundamentals is reinforced when agent j is free to optimize and not reinforced when agent j's action is fixed at the level $e^j_S = e^j_\infty(\boldsymbol{A})$.

   That is, $\boldsymbol{\phi_\infty} < \boldsymbol{\phi_S} < \boldsymbol{\Phi}$. ▸ simulation

2. When either of the agents are more overconfident, both agents' underestimation of the fundamental is more severe.

   That is, let $\tilde{\boldsymbol{a}} > \hat{\boldsymbol{a}} > \boldsymbol{A}$, then $\boldsymbol{\phi_\infty}(\tilde{\boldsymbol{a}}) < \boldsymbol{\phi_\infty}(\hat{\boldsymbol{a}}) < \boldsymbol{\Phi}$.

Mutually-Limiting Learning

Proposition 2 - Negative Info Ext + Mutual Learning = Better Inferences

Suppose agent j has a negative informational externality on agent i.

1. Agent i's underestimation of the fundamental is less severe when agent j is free to optimize than when agent j's action is fixed at e_S^j, i.e. $\phi_S^i < \phi_\infty^i(\tilde{\boldsymbol{a}}) < \Phi$.
2. As agent j becomes more overconfident, agent i's underestimation of the fundamental is less severe. That is, for any $\tilde{\boldsymbol{a}} > \hat{\boldsymbol{a}} > \boldsymbol{A}$ such that $\tilde{a}^j > \hat{a}^j$ and $\tilde{a}^i = \hat{a}^i$, we have $\phi_\infty^i(\hat{\boldsymbol{a}}) < \phi_\infty^i(\tilde{\boldsymbol{a}}) < \Phi$.

Welfare Analysis

Overconfidence does not necessarily result in welfare loss.

- Overconfidence may correct the inefficiencies arising from free-riding problems.

In our paper we characterize conditions under which both agents are better off when they are overconfident.

Welfare Examples

Ex 1 - Positive Information Externalities - Agents are learning about a common idea ϕ, while contributing to different parts of a project:

$$Q^i\left(e^i, e^j, a^i, \phi\right) = \left(e^i + e^j\right)\phi + a^i + e^i e^j - c(e^i)^2.$$

Ex 2 - Negative Information Externalities - Venture Capitalists are investing in different start-ups in the same industry:

$$Q^i\left(e^i, e^j, a^i, \phi\right) = e^i(\phi - \lambda e^j) + a^i - c(e^i)^2$$

Welfare Changes

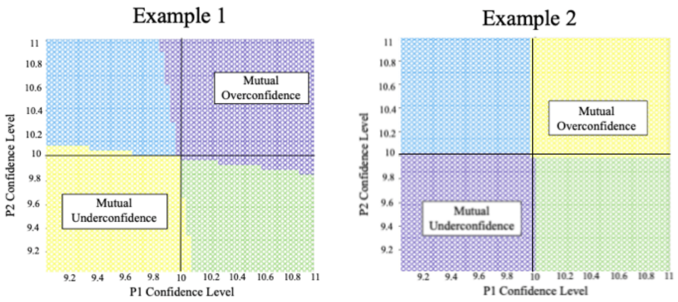

Figure: Purple: both worse off. Yellow: both better off. Green: Player 1 is better off and player 2 is worse off. Blue: Player 1 is worse off and player 2 is better off.

Belief Convergence

Proposition 3

Suppose the informational externalities are either both positive or both negative. Then the agents' actions almost surely converge to $\left(e^1_\infty, e^2_\infty\right)$ and their beliefs almost surely converge in distribution to the Dirac measure at $\left(\phi^1_\infty, \phi^2_\infty\right)$.

- We build our proof on Heidhues et al. (2018)'s contraction argument.
- Key difference: in the case of two agents, informational externalities play a role.

Extensions

- Underconfidence: Mutual learning patterns are reversed.
- Multiple equilibria
- More than two agents

Final Remarks

- Takeaways:
  - Informational externalities $\Rightarrow$ mutually-reinforcing/limiting learning.
  - Which mutual learning pattern occurs depends on the specific game.
  - In contrast to single-agent environment, overconfidence is not always bad.
- Question for future research:
  - What if agents start to exploit their informational externalities?

Thank you!

Sufficient Conditions for a Unique Berk-Nash Equilibrium

Lemma

There exists a unique pure-strategy Berk-Nash equilibrium if for all $i \in I$, derivatives Q^i_ϕ, Q^i_a are bounded, and $|\tilde{a}^i - A^i|$ is sufficiently small.

Definition of KL divergence

- The KL divergence is defined as

$$K^i\left(\boldsymbol{e},\hat{\phi}^i\right)=\mathbb{E}^\epsilon\left[\log\frac{f\left(\epsilon^i\right)}{f\left(Q^i\left(\boldsymbol{e},A^i,\Phi\right)-Q^i\left(\boldsymbol{e},\tilde{a}^i,\phi\right)+\epsilon^i\right)}\right].$$

Simulation: Mutually-Reinforcing Learning