Mental Models and Endogenous Learning

Hakan Özyılmaz

Toulouse School of Economics

EEA
Aug 2023

Motivation

- Large literature on biases in decision making

Motivation

- Large literature on biases in decision making
  - overconfidence (Moore and Healy 2008, Möbius et al. 2014), law of small numbers (Tversky and Kahneman 1971, Rabin 2002), failures in contingent thinking (Martínez-Marquina, Niederle, and Vespa 2019)

Motivation

- Large literature on biases in decision making
  - overconfidence (Moore and Healy 2008, Möbius et al. 2014), law of small numbers (Tversky and Kahneman 1971, Rabin 2002), failures in contingent thinking (Martínez-Marquina, Niederle, and Vespa 2019)
- **Question:** What are the implications of these biases on how people learn and behave?

Motivation

- Large literature on biases in decision making
  - overconfidence (Moore and Healy 2008, Möbius et al. 2014), law of small numbers (Tversky and Kahneman 1971, Rabin 2002), failures in contingent thinking (Martínez-Marquina, Niederle, and Vespa 2019)
- **Question:** What are the implications of these biases on how people learn and behave?
  - Biases might persist and confound learning about other variables

Motivation

- Large literature on biases in decision making
  - overconfidence (Moore and Healy 2008, Möbius et al. 2014), law of small numbers (Tversky and Kahneman 1971, Rabin 2002), failures in contingent thinking (Martínez-Marquina, Niederle, and Vespa 2019)
- **Question:** What are the implications of these biases on how people learn and behave?
  - Biases might persist and confound learning about other variables
  - Learning provides opportunities to correct biases

Motivation

- Large literature on biases in decision making
  - overconfidence (Moore and Healy 2008, Möbius et al. 2014), law of small numbers (Tversky and Kahneman 1971, Rabin 2002), failures in contingent thinking (Martínez-Marquina, Niederle, and Vespa 2019)
- **Question:** What are the implications of these biases on how people learn and behave?
  - Biases might persist and confound learning about other variables
  - Learning provides opportunities to correct biases
- Know little about the persistence of biases and how their persistence may generate mislearning

Motivation

- Growing theoretical literature on how biased individuals learn about payoff-relevant parameters in their environment
  - Take biases given as "misspecified mental models"
  - Otherwise neoclassical: Bayesian and optimal actions
  - Berk 1966, Esponda and Pouzo 2016, Heidhues, Kőszegi, and Strack 2018, Fudenberg, Lanzani, and Strack 2021

Motivation

- Growing theoretical literature on how biased individuals learn about payoff-relevant parameters in their environment
  - Take biases given as "misspecified mental models"
  - Otherwise neoclassical: Bayesian and optimal actions
  - Berk 1966, Esponda and Pouzo 2016, Heidhues, Kőszegi, and Strack 2018, Fudenberg, Lanzani, and Strack 2021
- **This paper:**
  - Conceptualize relative overconfidence as a source of model misspecification as in Heidhues, Kőszegi, and Strack 2018
  - Experimentally study how (relatively) overconfident people learn about a parameter that informs their action

Motivation

- Growing theoretical literature on how biased individuals learn about payoff-relevant parameters in their environment
  - Take biases given as "misspecified mental models"
  - Otherwise neoclassical: Bayesian and optimal actions
  - Berk 1966, Esponda and Pouzo 2016, Heidhues, Kőszegi, and Strack 2018, Fudenberg, Lanzani, and Strack 2021
- **This paper:**
  - Conceptualize relative overconfidence as a source of model misspecification as in Heidhues, Kőszegi, and Strack 2018
  - Experimentally study how (relatively) overconfident people learn about a parameter that informs their action

    More broadly, an experimental test of misspecified model approach to study the persistence of biases

Theoretical Framework

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

- His ability is $a_0 \in A = \{20, 40, 60, 80, 100\}$

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

- His ability is $a_0 \in A = \{20, 40, 60, 80, 100\}$

  but he incorrectly believes $a = \tilde{a} > a_0$ with probability 1

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

- His ability is $a_0 \in A = \{20, 40, 60, 80, 100\}$

  but he incorrectly believes $a = \tilde{a} > a_0$ with probability 1

- The project's quality $\phi_0 \in \Phi = \{0, 1, ..., 100\}$ is uncertain

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

- His ability is $a_0 \in A = \{20, 40, 60, 80, 100\}$

  but he incorrectly believes $a = \tilde{a} > a_0$ with probability 1
- The project's quality $\phi_0 \in \Phi = \{0, 1, ..., 100\}$ is uncertain

  and drawn from a discrete uniform dist'n independent from a

Environment

- A manager periodically chooses an investment level $e_t \in [0, 100]$ on a project to maximize profit y_t

$$y(e_t, a, \phi) = (a + e_t)\phi - \frac{e_t^2}{2}, \;\; e^*(\phi) = \phi$$

- His ability is $a_0 \in A = \{20, 40, 60, 80, 100\}$

  but he incorrectly believes $a = \tilde{a} > a_0$ with probability 1
- The project's quality $\phi_0 \in \Phi = \{0, 1, ..., 100\}$ is uncertain

  and drawn from a discrete uniform dist'n independent from a
- The manager receives noisy performance feedback $f_0(e_t) \sim Bernoulli(\mu(e_t, a_0, \phi_0))$

$$\mu(e_t, a_0, \phi_0) = \frac{y(e_t, a_0, \phi_0) - y_{min}}{y_{max} - y_{min}}$$

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?
- **Exogenous Learning**: As a benchmark case, fix the investment at the optimal level

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?
- **Exogenous Learning**: As a benchmark case, fix the investment at the optimal level
  - **Hypothesis 1:** the overconfident manager grows pessimistic beliefs on the project quality

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?
- **Exogenous Learning**: As a benchmark case, fix the investment at the optimal level
  - **Hypothesis 1:** the overconfident manager grows pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?
- **Exogenous Learning**: As a benchmark case, fix the investment at the optimal level
  - **Hypothesis 1:** the overconfident manager grows pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

    $\rightarrow$ accumulates evidence revealing lower than expected profit

Hypotheses — Exogenous Learning

- What is the limit belief and action for the overconfident manager who uses Bayesian updating and myopically optimal actions?
- **Exogenous Learning**: As a benchmark case, fix the investment at the optimal level
  - **Hypothesis 1:** the overconfident manager grows pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

    $\rightarrow$ accumulates evidence revealing lower than expected profit

    $\rightarrow$ attributes lower than expected profit to a lower project quality

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences
  - **Hypothesis 2:** the overconfident manager grows *even more* pessimistic beliefs on the project quality

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences
  - **Hypothesis 2:** the overconfident manager grows *even more* pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences
  - **Hypothesis 2:** the overconfident manager grows *even more* pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

    $\rightarrow$ attributes lower than expected profit to a lower project quality

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences
  - **Hypothesis 2:** the overconfident manager grows *even more* pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

    $\rightarrow$ attributes lower than expected profit to a lower project quality

    $\rightarrow$ *chooses a suboptimal investment level and lowers the profit further*

Hypotheses — Endogenous Learning

- **Endogenous Learning**: Now allow the manager to change his investment level in response to his inferences
  - **Hypothesis 2:** the overconfident manager grows *even more* pessimistic beliefs on the project quality
  - **intuition**: expects an unrealistically high average profit as $\tilde{a} > a_0$ and $y_a \geq 0$

    $\rightarrow$ attributes lower than expected profit to a lower project quality

    $\rightarrow$ *chooses a suboptimal investment level and lowers the profit further*

    $\rightarrow$ rationalizes lower than expected profit with even lower project quality

Experimental Design

Overview

Consists of 4 parts

Overview

Consists of 4 parts

- Part 1. Measurement of ability — "IQ score" using Raven matrices

Overview

Consists of 4 parts

- Part 1. Measurement of ability — “IQ score” using Raven matrices
- Part 2. Eliciting prior beliefs on IQ rank quintiles

Overview

Consists of 4 parts

- Part 1. Measurement of ability — “IQ score” using Raven matrices
- Part 2. Eliciting prior beliefs on IQ rank quintiles
- Part 3 (**Main Part**). Investment decision $\rightarrow$ Feedback

Main Part — Treatments

- **Exogenous:**
  - Subjects make investment decisions but decisions are implemented at the end of the experiment
  - Feedback is based on $e = 100$ in each period

Main Part — Treatments

- **Exogenous:**
  - Subjects make investment decisions but decisions are implemented at the end of the experiment
  - Feedback is based on $e = 100$ in each period
- **Endogenous:**
  - Subjects' investment decisions are implemented immediately
  - Feedback is based on subject's decision for that period

Main Part — Key Design Features

To isolate the role of overconfidence in potentially generating mislearning and suboptimal behavior, I aim to minimize

- errors related to Bayesian updating
- errors related to optimization

Main Part — Minimizing updating errors

To minimize errors related to Bayesian updating, I provide subjects with Bayesian posterior means for each fixed model $\bar{a}$:
$E_{\pi_t^{\bar{a}}}[\phi \mid \mathcal{F}_{t-1}]$

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

Main Part — Minimizing optimization errors

To minimize errors related to optimization, I provide subjects with special calculators

Calculator	The Statistician's Report	
Enter Likelihood (out of 100)	Your IQ Rank Score	Project Quality
	20	50
	40	50
	60	50
50	80	50
50	100	50
CALCULATE	50	CLEAR

Main Part — Decision Screen

Period **1**

Calculator	The Statistician's Report	
Enter Likelihood (out of 100)	**Your IQ Rank Score**	**Project Quality**
	20	50
	40	50
	60	50
50	80	50
50	100	50
CALCULATE	50	CLEAR

Decision Box

Choose Investment Amount

SUBMIT

Main Part — Feedback

Period **1**

The Company's Evaluation: **Satisfactory**

Your profit is higher than the company's expectation.

Main Part — Accelerating Learning

- **Periods 1 to 10**

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods
  - aggregate feedback for 10 periods is provided after each decision

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods
  - aggregate feedback for 10 periods is provided after each decision
- **Periods 101 to 1000**

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods
  - aggregate feedback for 10 periods is provided after each decision
- **Periods 101 to 1000**
  - decision in every 100 periods

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods
  - aggregate feedback for 10 periods is provided after each decision
- **Periods 101 to 1000**
  - decision in every 100 periods
  - each decision is implemented for the consecutive 99 periods

Main Part — Accelerating Learning

- **Periods 1 to 10**
  - decision in each period
  - single feedback after each decision
- **Periods 11 to 100**
  - decision in every 10 periods
  - each decision is implemented for the consecutive 9 periods
  - aggregate feedback for 10 periods is provided after each decision
- **Periods 101 to 1000**
  - decision in every 100 periods
  - each decision is implemented for the consecutive 99 periods
  - aggregate feedback for 100 periods is provided after each decision

Overview

Consists of 4 parts

- Part 1. Measurement of ability — “IQ score” using Raven matrices
- Part 2. Eliciting prior beliefs on IQ rank quintiles
- Part 3 (**Main Part**). Investment decision $\rightarrow$ Feedback
- Part 4. Eliciting posterior beliefs on IQ rank quintiles

Overview

Consists of 4 parts

- Part 1. Measurement of ability — "IQ score" using Raven matrices
- Part 2. Eliciting prior beliefs on IQ rank quintiles
- Part 3 (**Main Part**). Investment decision $\rightarrow$ Feedback
- Part 4. Eliciting posterior beliefs on IQ rank quintiles

We randomly select one part for payment

Procedural Details

- UCSB EBEL
- 210 participants
- Average pay \$27.6 including \$10 show-up fee

Results

Overconfidence as a Source of Model Misspecification

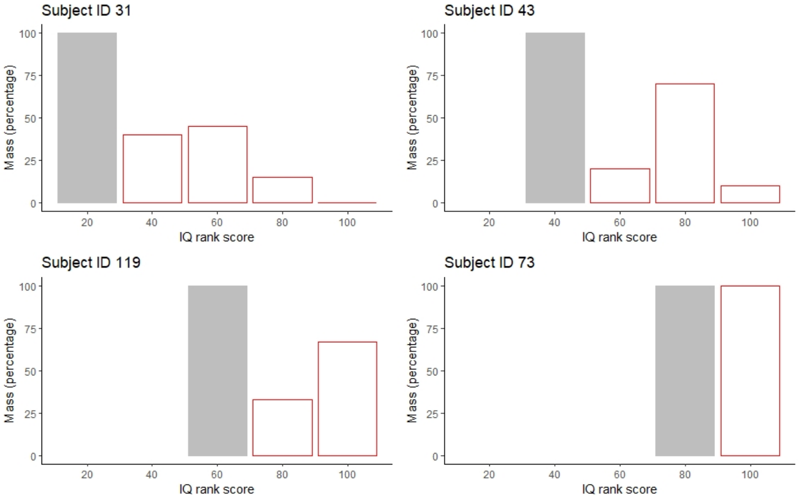

Priors on IQ rank score

30% of subjects assign 100% probability to strictly higher IQ rank scores

Correctly Specified Subjects

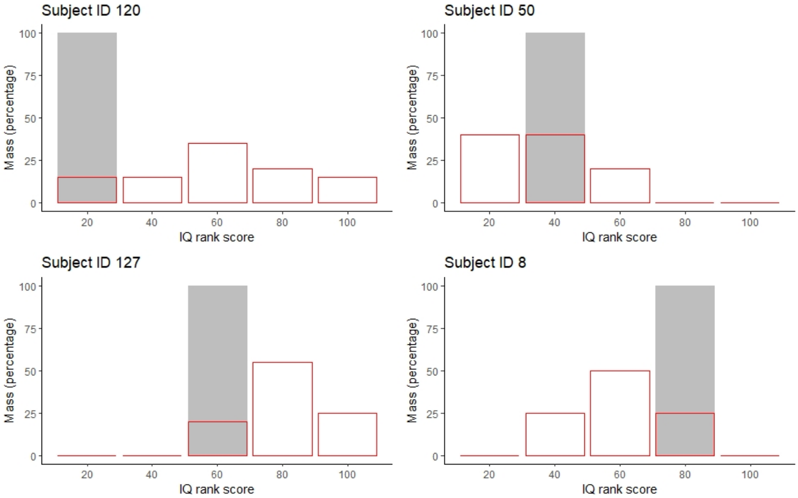

Priors on IQ rank score

70% of subjects assign *some* probability to their correct IQ rank scores

Do Misspecified Models Generate Suboptimal Behavior?

We define a displacement measure $\Delta_{OPT} = e - e^*(\phi_0)$ to measure "distance" from the optimal action

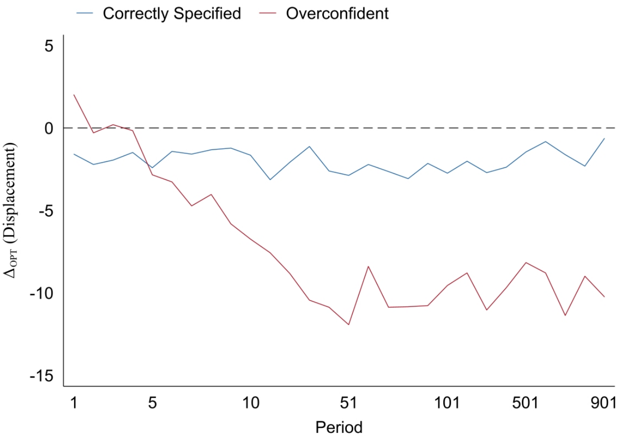

▸ Statistical Evidence

Result#1

- Overconfident subjects mislearn the fundamental and take actions that are *lower* than the first-best action by the end of the experiment

Result#1

- Overconfident subjects mislearn the fundamental and take actions that are *lower* than the first-best action by the end of the experiment
- Correctly specified subjects' learning process does not lead them away from the first-best action

Does Endogenous Learning Exacerbate Suboptimal Behavior for Overconfident Subjects?

Displacement measure: $\Delta_{OPT} = e - e^*(\phi_0)$

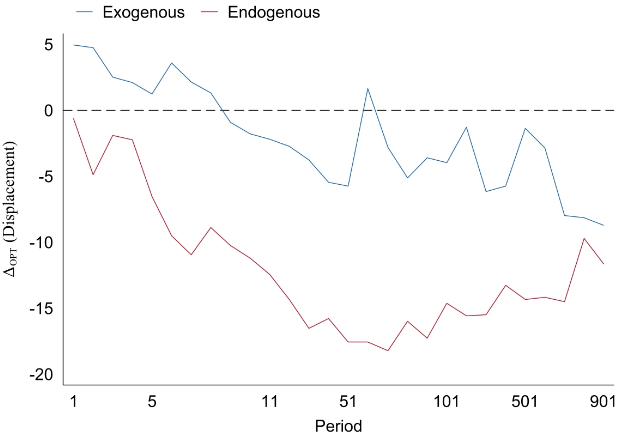

Statistical Evidence

Result#2

- Inconsistent with the theoretical prediction, endogenous learning does not exacerbate the extent of suboptimal behavior for overconfident subjects

Result#2

- Inconsistent with the theoretical prediction, endogenous learning does not exacerbate the extent of suboptimal behavior for overconfident subjects
- **Why?** Any deviation from the theory hinges on
  - Optimization errors: Over 80% of the subjects follow the calculator in both treatments
  - Errors related to Bayesianism: deviations can only occur through learning on own ability by design

Result#2

- Inconsistent with the theoretical prediction, endogenous learning does not exacerbate the extent of suboptimal behavior for overconfident subjects
- **Why?** Any deviation from the theory hinges on
  - Optimization errors: Over 80% of the subjects follow the calculator in both treatments
  - Errors related to Bayesianism: deviations can only occur through learning on own ability by design
- Comparison of elicited priors and posteriors on own ability confirms this suggestion
  - Endogenous feedback leads subjects to reduce their overconfidence further

Conclusion

- Clearly generate misspecified mental models through overconfidence and investigate their influence on long-run behavior

Conclusion

- Clearly generate misspecified mental models through overconfidence and investigate their influence on long-run behavior
- Document relative overconfidence mostly persists, generates mislearning and suboptimal behavior

Conclusion

- Clearly generate misspecified mental models through overconfidence and investigate their influence on long-run behavior
- Document relative overconfidence mostly persists, generates mislearning and suboptimal behavior
- On the bright side, I show endogenous learning may not exacerbate suboptimal behavior

  — abundant feedback "weakens" misspecified mental models at the extensive and intensive margin

  — provide evidence that "weakening" is stronger with endogenous feedback

Conclusion

- Clearly generate misspecified mental models through overconfidence and investigate their influence on long-run behavior
- Document relative overconfidence mostly persists, generates mislearning and suboptimal behavior
- On the bright side, I show endogenous learning may not exacerbate suboptimal behavior

  — abundant feedback "weakens" misspecified mental models at the extensive and intensive margin

  — provide evidence that "weakening" is stronger with endogenous feedback
- Investigating how endogenous learning affects information processing seems a crucial future direction

Thanks!

hakan.ozyilmaz@tse-fr.eu

Do overconfident subjects act consistent with Bayesianism?

Displacement measure: $\Delta_{BAYES} = e - e^*(\Pi_{sim}[\phi])$

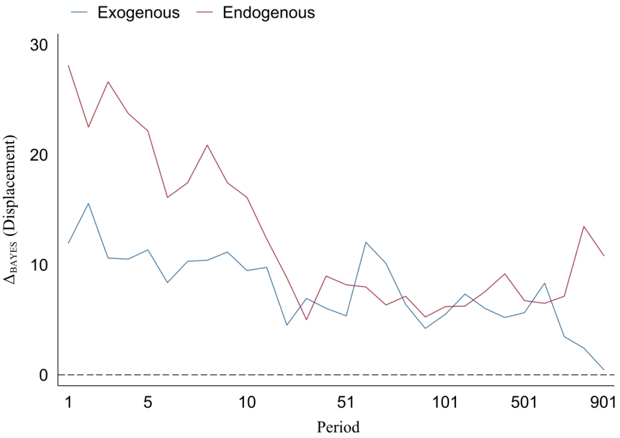

▸ Statistical Evidence

Result#4

- Overconfident subjects in **Exogenous** take actions that are indistinguishable from the Bayesian action in the last period

Result#4

- Overconfident subjects in **Exogenous** take actions that are indistinguishable from the Bayesian action in the last period
- Overconfident subjects in **Endogenous** take significantly higher actions than the Bayesian action in the last period

Result#4

- Overconfident subjects in **Exogenous** take actions that are indistinguishable from the Bayesian action in the last period
- Overconfident subjects in **Endogenous** take significantly higher actions than the Bayesian action in the last period
- The direction of deviation in **Endogenous** is consistent with "weakened" overconfidence

Do overconfident subjects learn about their ability?

"Distance" between expected and true ability: $\Delta_a = E_p[a] - a_0$

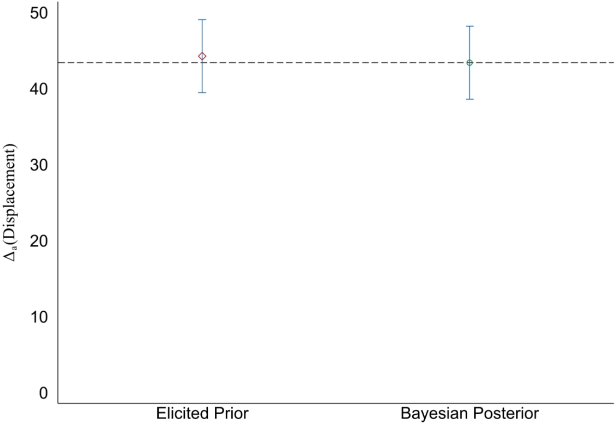

Elicited Prior: 44.3, Bayesian Posterior: 43.4

Overconfident Subjects: Priors vs Posteriors on Ability

"Distance" between expected and true ability: $\Delta_a = E_p[a] - a_0$

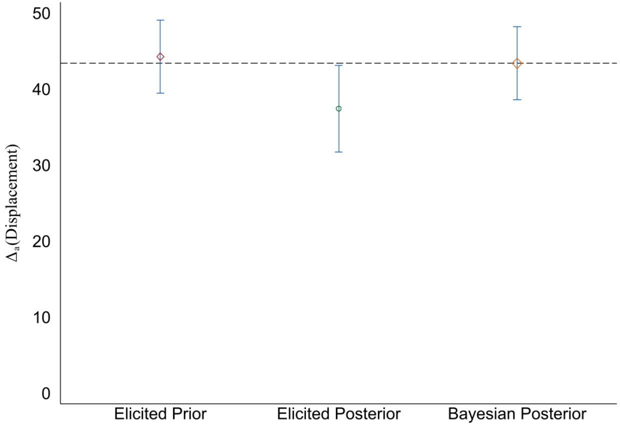

Elicited Prior: 44.3, Elicited Posterior: 37.4 ($p = 0.02$)

Overconfident Subjects: Exogenous and Endogenous

“Distance” between expected and true ability: $\Delta_a = E_p[a] - a_0$

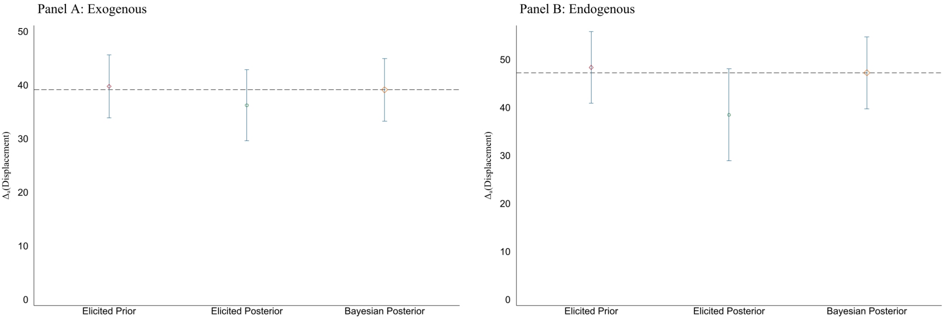

Result#5

- Inconsistent with Bayesian Learning, overconfident subjects have more accurate expectations of their own abilities by the end of the experiment

Result#5

- Inconsistent with Bayesian Learning, overconfident subjects have more accurate expectations of their own abilities by the end of the experiment
- A significant portion of overconfident subjects assign some probability to their true abilities at the end of the experiment

Result#5

- Inconsistent with Bayesian Learning, overconfident subjects have more accurate expectations of their own abilities by the end of the experiment
- A significant portion of overconfident subjects assign some probability to their true abilities at the end of the experiment
- Self-learning is more pronounced with endogenous feedback

"Distance" from the First-Best Optimal

Using $\Delta_{OPT} = e - e^*(\phi_0)$, we estimate a "distance-from-benchmark" regression $\Delta_{OPT} = \alpha + \beta M + \varepsilon$ in each period

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
β	3.133	-6.763*	-9.947***	-9.674***
	(5.088)	(3.168)	(2.256)	(2.430)
α	-1.363	-1.435	-1.541	-0.605
	(2.677)	(1.579)	(1.511)	(1.469)
Observations	210	210	210	210
Period	1	501	701	901

"Distance" from the First-Best Optimal

Using $\Delta_{OPT} = e - e^*(\phi_0)$, we estimate a "distance-from-benchmark" regression $\Delta_{OPT} = \alpha + \theta M + \varepsilon$ in each period

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
θ	-5.603	-12.97*	-6.529	-2.934
	(9.018)	(5.680)	(3.452)	(4.131)
α	4.967	-1.376	-7.963**	-8.711*
	(6.272)	(5.024)	(2.429)	(3.295)
Observations	63	63	63	63
Period	1	501	701	901

Switching Mental Models

		Posterior Models		
		Overconfident	Correctly Specified	Underconfident
Prior	Overconfident	78%	22%	0%
Models	Correctly Specified	13%	85%	2%

Switching Mental Models

		Posterior Models		
		Overconfident	Correctly Specified	Underconfident
Prior	Overconfident	78%	22%	0%
Models	Correctly Specified	13%	85%	2%

- The deviations are significant for both overconfident and correctly specified subjects ($p = 0.002$, $p < 0.001$)

Switching Mental Models

		Posterior Models		
		Overconfident	Correctly Specified	Underconfident
Prior Models	Overconfident	78%	22%	0%
	Correctly Specified	13%	85%	2%

- The deviations are significant for both overconfident and correctly specified subjects ($p = 0.002$, $p < 0.001$)
- The reduction in Δ_a for overconfident subjects mainly comes from subjects who switch from overconfident to correctly specified