Higher-order Learning

September 18, 2021

Motivation

- A number of papers (e.g., Grether, 1980; Holt and Smith, 2009) studied how people update beliefs
- Other papers (e.g., Nagel, 1995) studied how people form beliefs about other people
- Missing gap: **how do people update their beliefs about the beliefs of other people?**

Common learning

	A	B	W
A	$1,1$	$-c,-c$	$-c,0$
B	$-c,-c$	$-c,-c$	$-c,0$
W	$0,-c$	$0,-c$	$0,0$

Parameter θ_A

	A	B	W
A	$-c,-c$	$-c,-c$	$-c,0$
B	$-c,-c$	$1,1$	$-c,0$
W	$0,-c$	$0,-c$	$0,0$

Parameter θ_B

FIGURE 1.—Payoffs from a potential joint opportunity, with actions A, B, or wait (W) available to each agent in each period.

For player 1 to choose A, it needs to be the case that

$$P_1(a_2 = A, \theta = \theta_A) \geq q = \frac{c}{c+1}$$

This implies:

$$P_1(\theta = \theta_A) \geq q$$

But also:

$$P_1(a_2 = A) \geq q \Rightarrow$$
$$P_1[P_2(a_1 = A, \theta = \theta_A) \geq q] \geq q \Rightarrow$$
$$P_1[P_2(\theta = \theta_A) \geq q] \geq q$$

Experimental design

- At the beginning of the experiment, each subject assigned to one of three possible roles: Player 1, Player 2, and Player 3
- Subjects are matched into teams of three and the teams are held fixed for the duration of the session (3 rounds)
- In the beginning of each round, a state $\theta \in \{ORANGE, PURPLE\}$ is drawn and held fixed for 30 periods
- State determines composition of the urn *in the round*:
  - ORANGE: 2 orange balls and 1 purple ball
  - PURPLE: 1 orange ball and 2 purple balls

Experimental design (round)

- In every period of every round, subjects observe a ball drawn from the unknown urn
- Two treatments:
  - **Public:** All players on the same team observe the same ball
  - **Private:** The draws are made with replacement
- After observing his draw,
  - Player 1 reports his belief $\in [0, 1]$ about the color of the urn
  - Player 2 reports his belief $\in [0, 1]$ about the reported belief of Player 1
  - Player 3 reports his belief $\in [0, 1]$ about the reported belief of Player 3
- Composition of the urn revealed at the end of the round
- No feedback regarding other subjects' decisions
- Each subject paid based on ONE randomly chosen period of ONE randomly chosen round
- Experiment for Player 1 is same as our classroom experiment

Binarized scoring rule

Bayesian Predictions

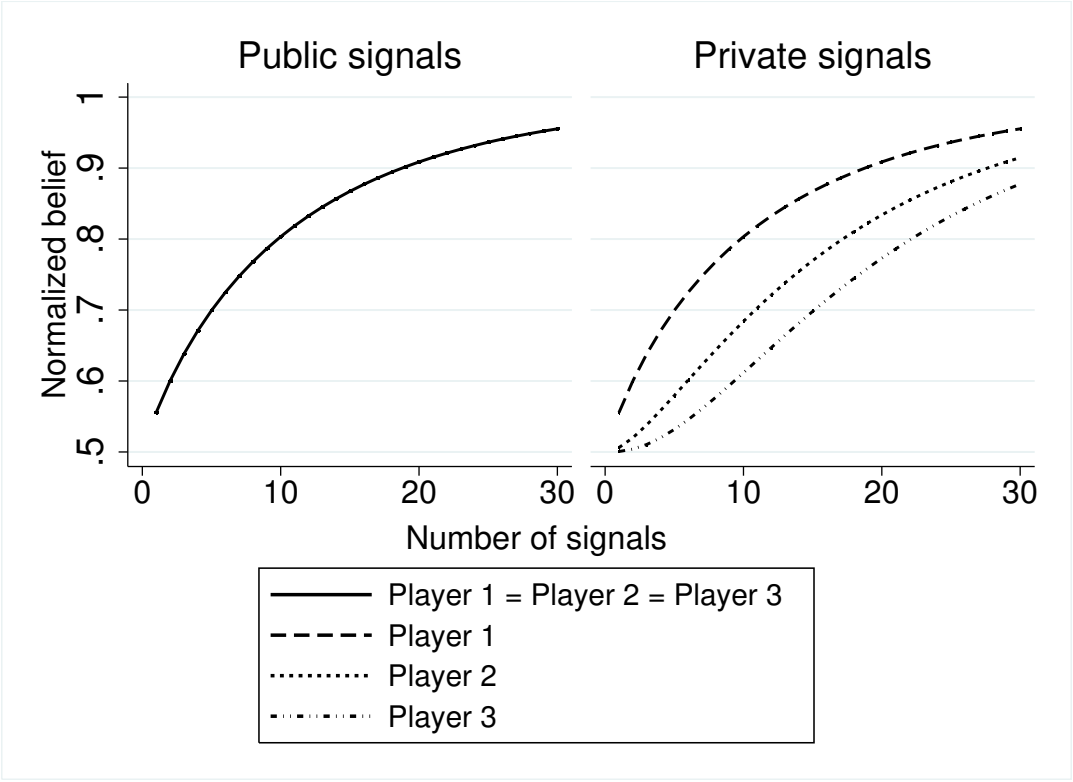

Prediction 1: Higher-order beliefs are closer to the prior with private information.

Bayesian Predictions

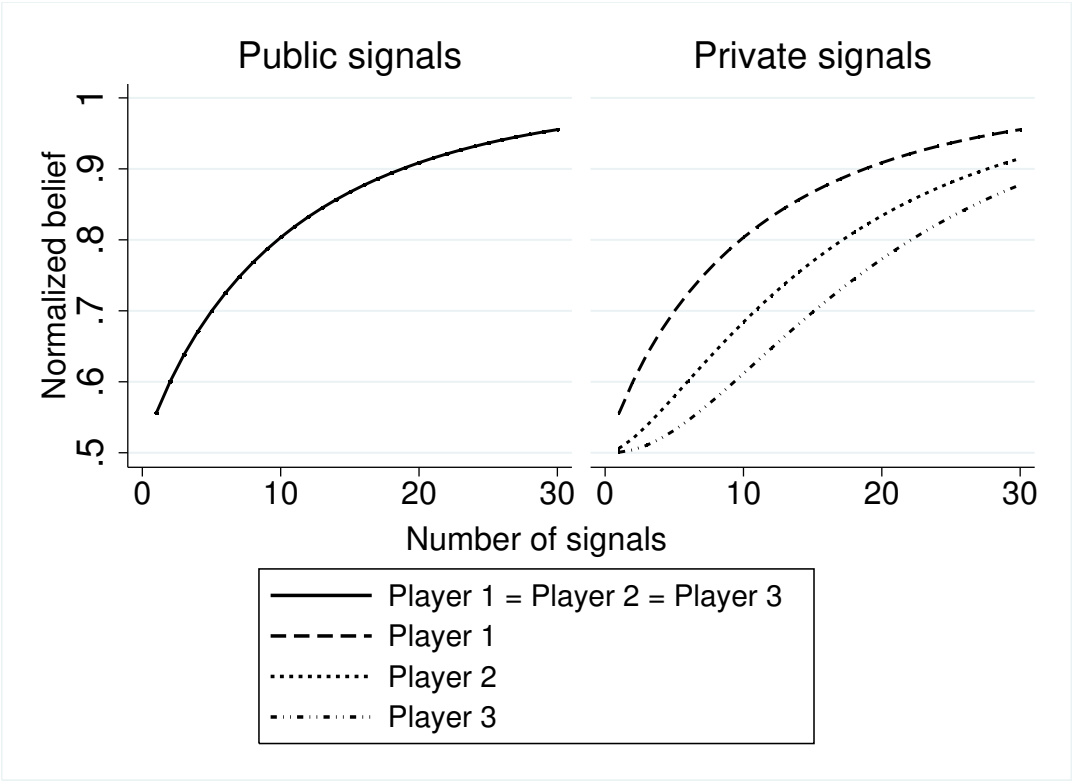

Prediction 2: On average, higher-order beliefs are more accurate in the public treatment, regardless of the number of signals observed.

Higher-order Learning

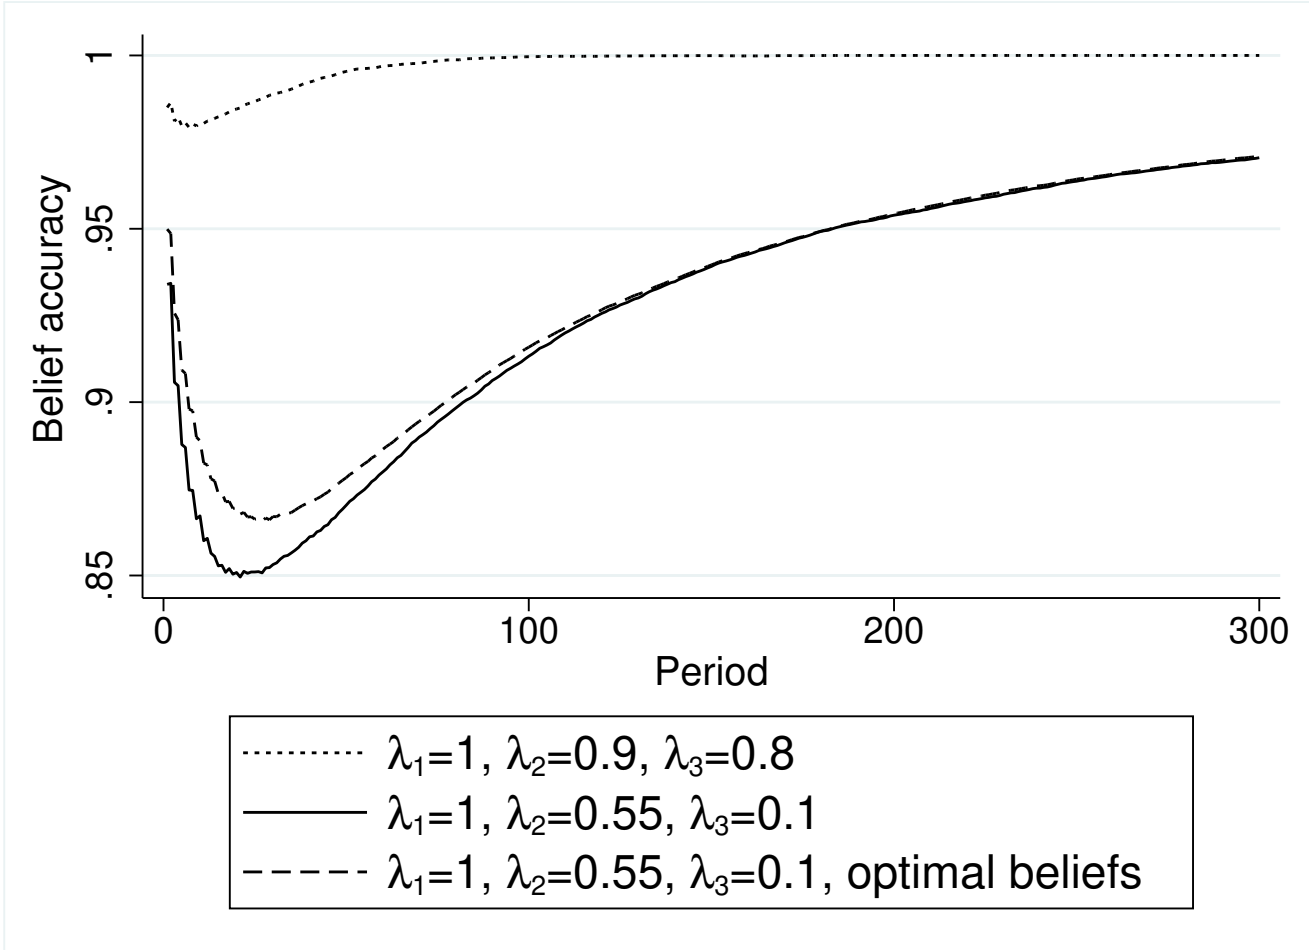

Results

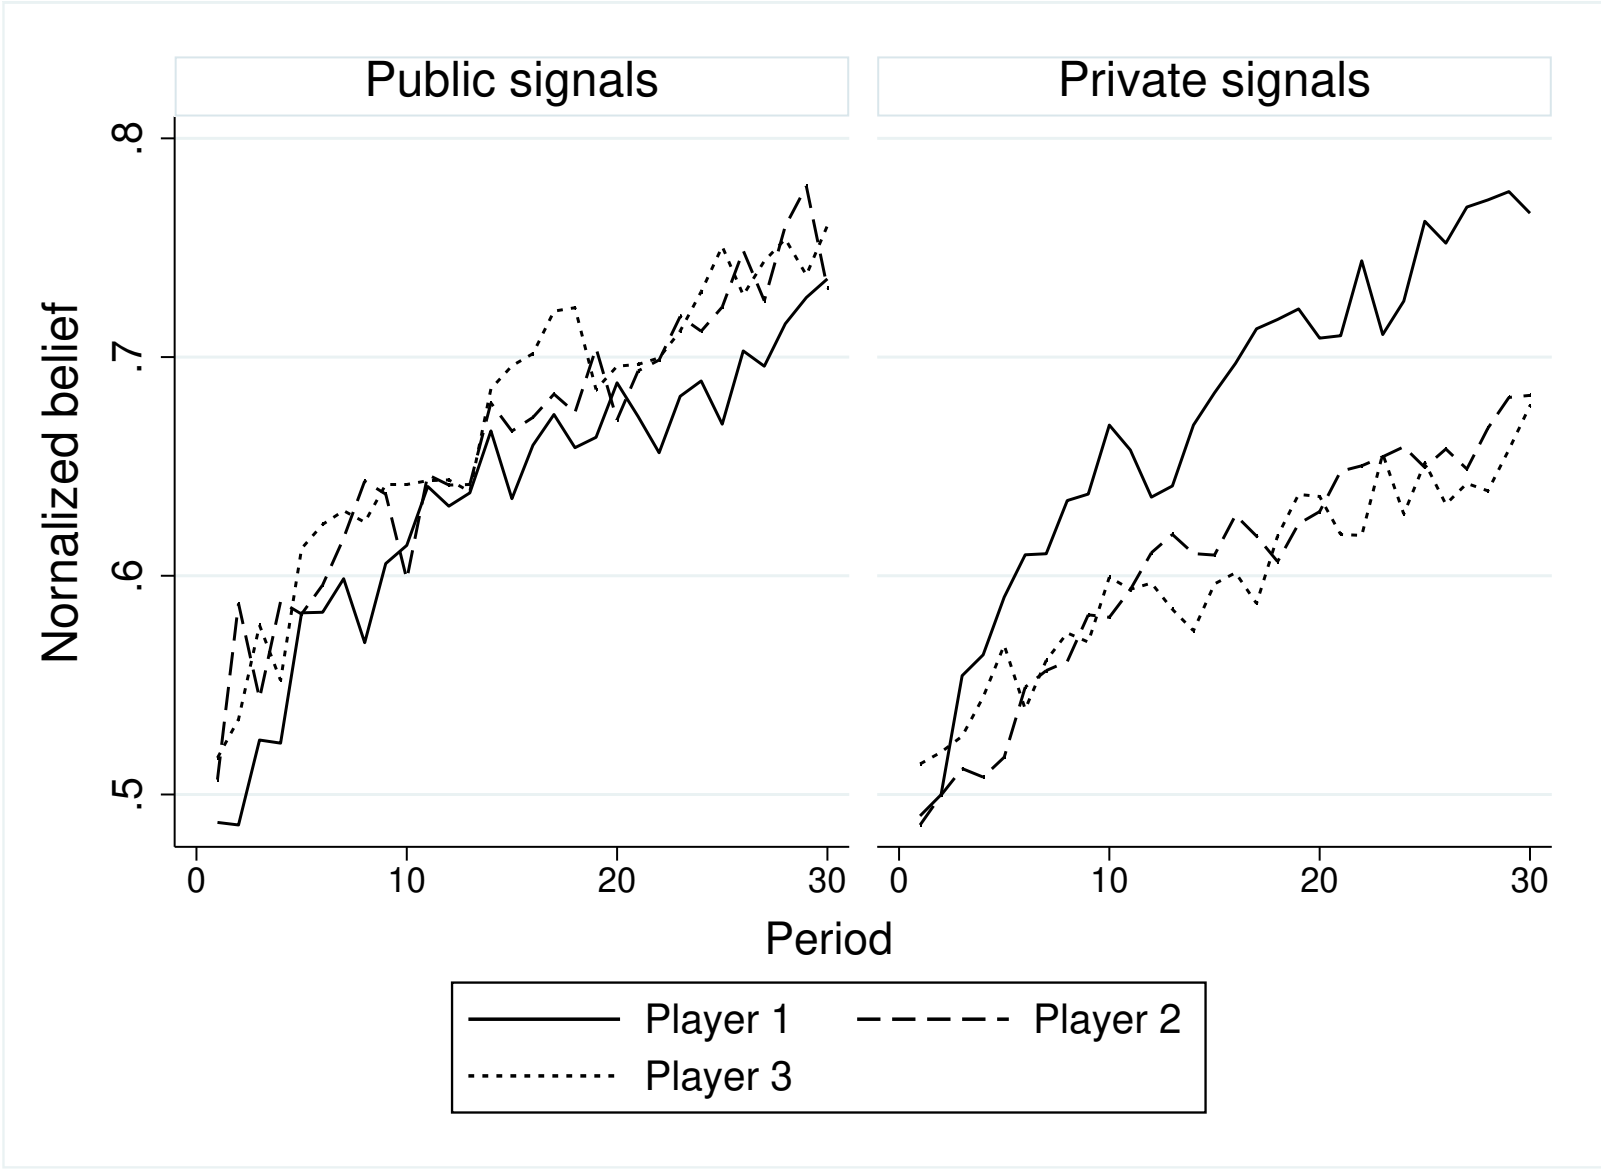
Public signals
Private signals
Normalized belief
.5
.6
.7
.8
0
10
20
30
Period
Player 1
Player 2
Player 3

	Players 2 and 3		All players, private		All players, public	
Private	-0.0651***	-0.0400**				
	(0.0221)	(0.0174)				
Period		0.00706****		0.00854****		0.00710****
		(0.000720)		(0.00113)		(0.00117)
Period × Private		-0.00162				
		(0.00109)				
Player 2			-0.0697*	-0.0330	0.0264	0.0277
			(0.0358)	(0.0257)	(0.0373)	(0.0283)
Player 3			-0.0742**	-0.0146	0.0341	0.0342
			(0.0343)	(0.0255)	(0.0366)	(0.0295)
Period × Player 2				-0.00237		-0.0000800
				(0.00159)		(0.00166)
Period × Player 3				-0.00385**		-0.00000986
				(0.00165)		(0.00144)
Constant	0.666****	0.557****	0.673****	0.541****	0.636****	0.526****
	(0.0160)	(0.0136)	(0.0276)	(0.0205)	(0.0292)	(0.0215)
Observations	14940	14940	11610	11610	10800	10800

Subject-clustered standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Belief accuracy

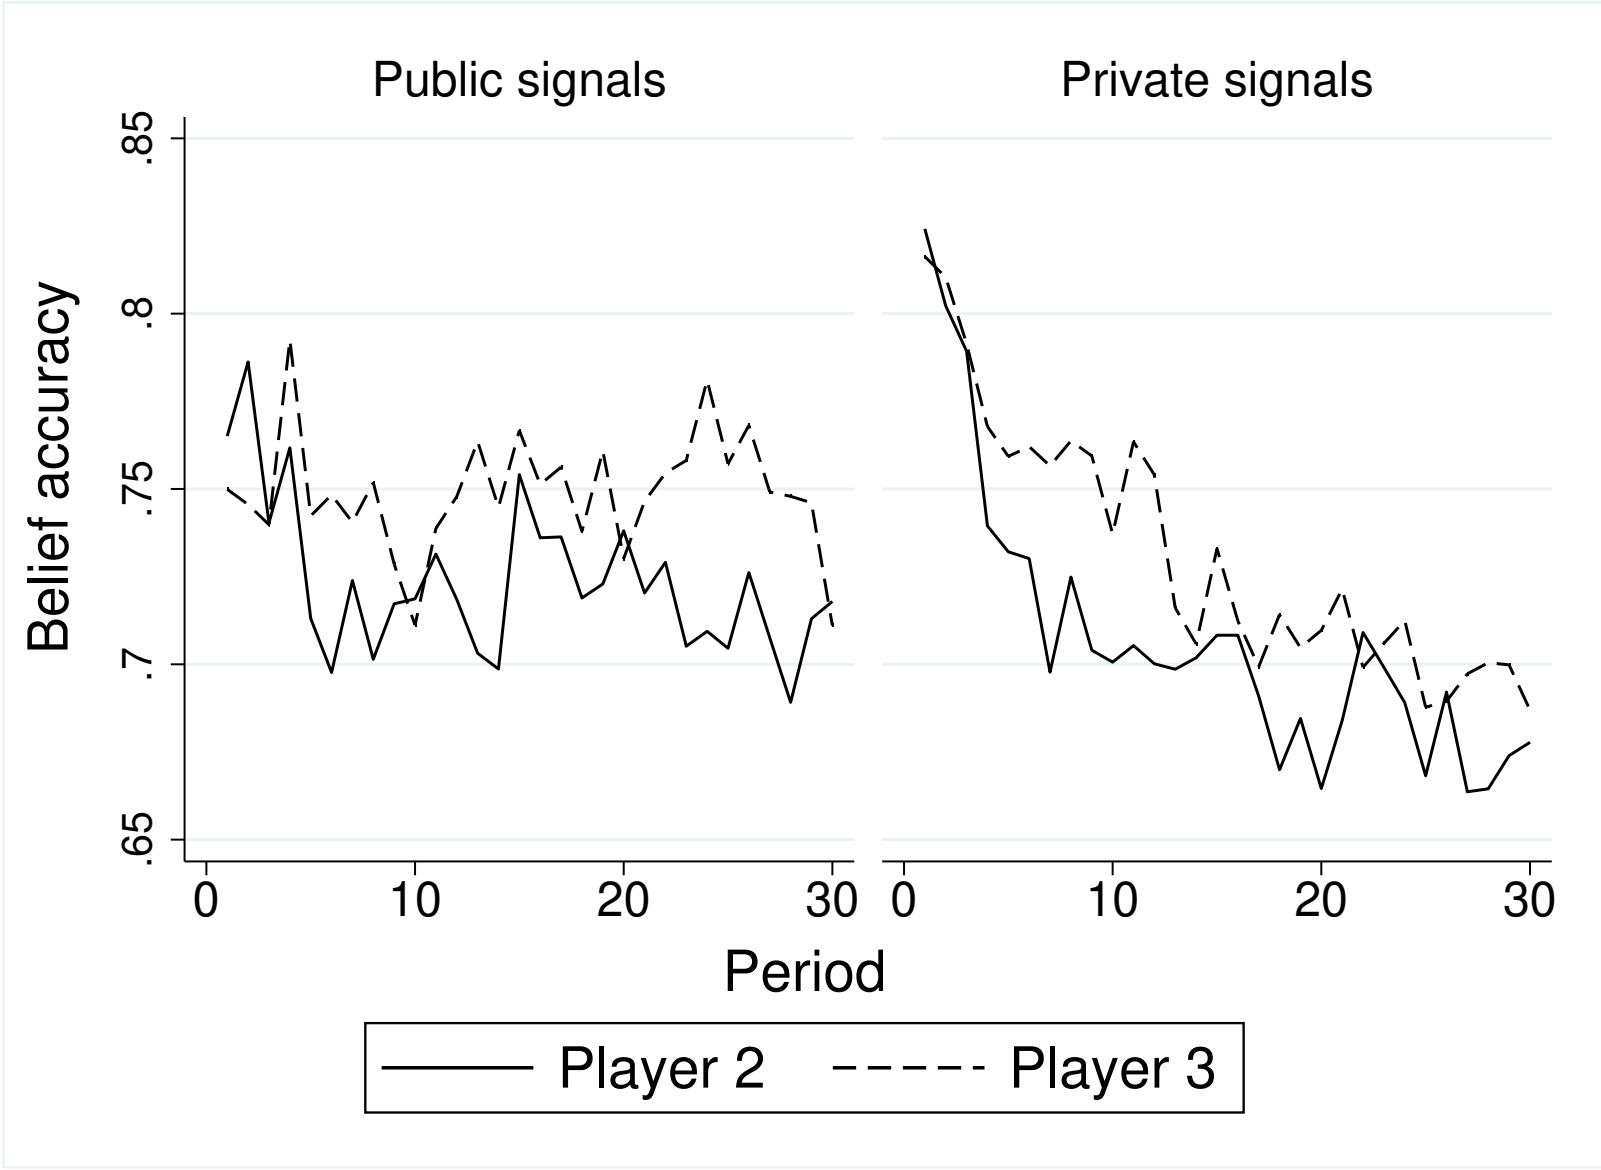

Belief accuracy

	Lab				MTurk	
	Players 2 and 3	Player 2	Player 3	Players 2 and 3	Player 2	Player 2
Private	-0.0173	-0.0169	-0.0177	0.0307	-0.0136	-0.0402**
	(0.0190)	(0.0282)	(0.0255)	(0.0194)	(0.0172)	(0.0189)
Period				-0.000556		-0.00196***
				(0.000610)		(0.000701)
Period × Private				-0.00310***		0.00172*
				(0.00101)		(0.00101)
Constant	0.736****	0.724****	0.749****	0.745****	0.672****	0.703****
	(0.0150)	(0.0225)	(0.0199)	(0.0136)	(0.0123)	(0.0131)
Observations	14940	7470	7470	14940	10620	10620

Subject-clustered standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Estimation of updating rules

- Estimation following the approach of Grether (1980):

$$\ln\left(\frac{\mu_n}{1-\mu_n}\right) = \beta_0 + \beta_{Prior}\ln\left(\frac{\mu_{n-1}}{1-\mu_{n-1}}\right) + \beta_{LR}\ln(LR) + \epsilon.$$

- For simplicity, only use beliefs in the public treatment ($N = 120$, 40 in each player role)

- Model estimated separately for each individual

Heterogeneity

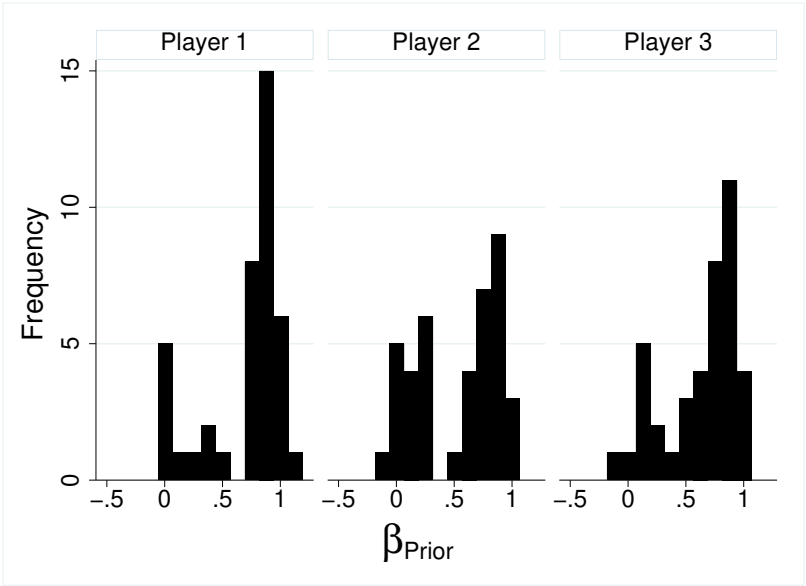

(a) PDF of β_{prior}

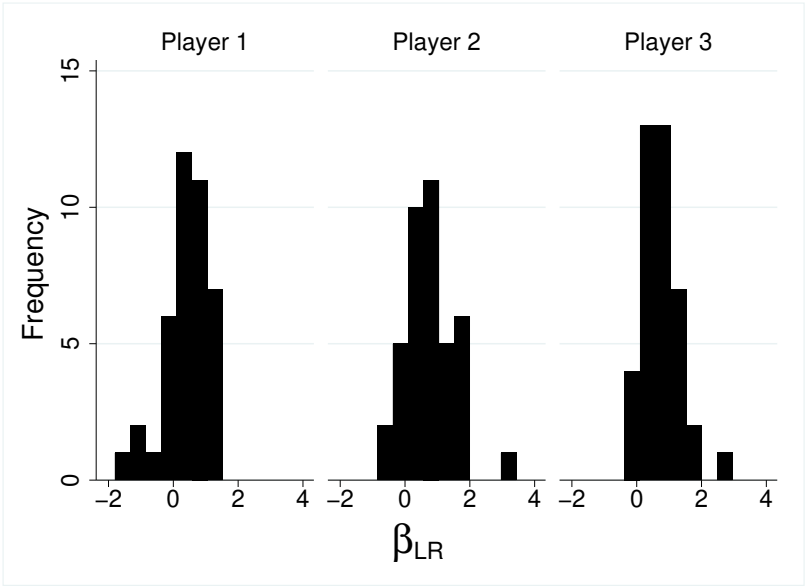

(b) PDF of β_{LR}

Counterfactual exercise

$$E[\mu_n|\mu_{n-1}, LR] = \frac{\mu_{n-1}^{\beta_{Prior}}}{\mu_{n-1}^{\beta_{Prior}} + e^{-\beta_0}(1-\mu_{n-1})^{\beta_{Prior}} LR^{-\beta_{LR}}}.$$

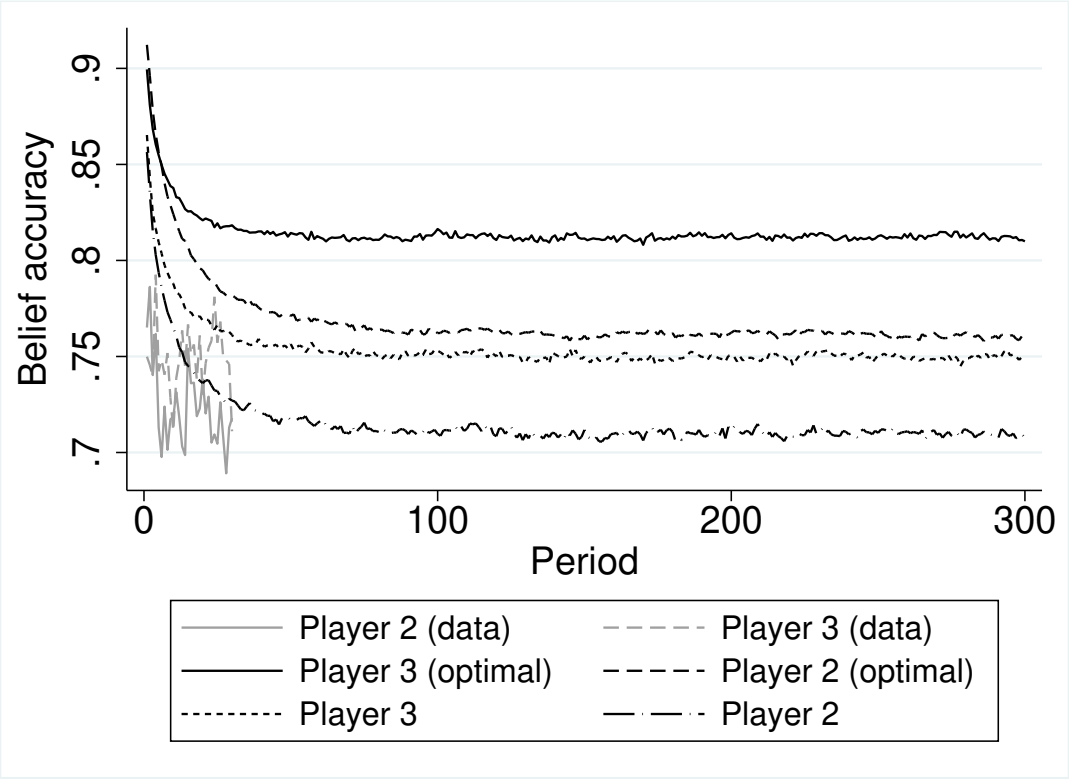

Base-rate neglect impedes convergence (Benjamin et al, 2019)

Simulated beliefs of an agent with $\beta_{LR} = 1$, $\beta_{prior} = 0.9$:

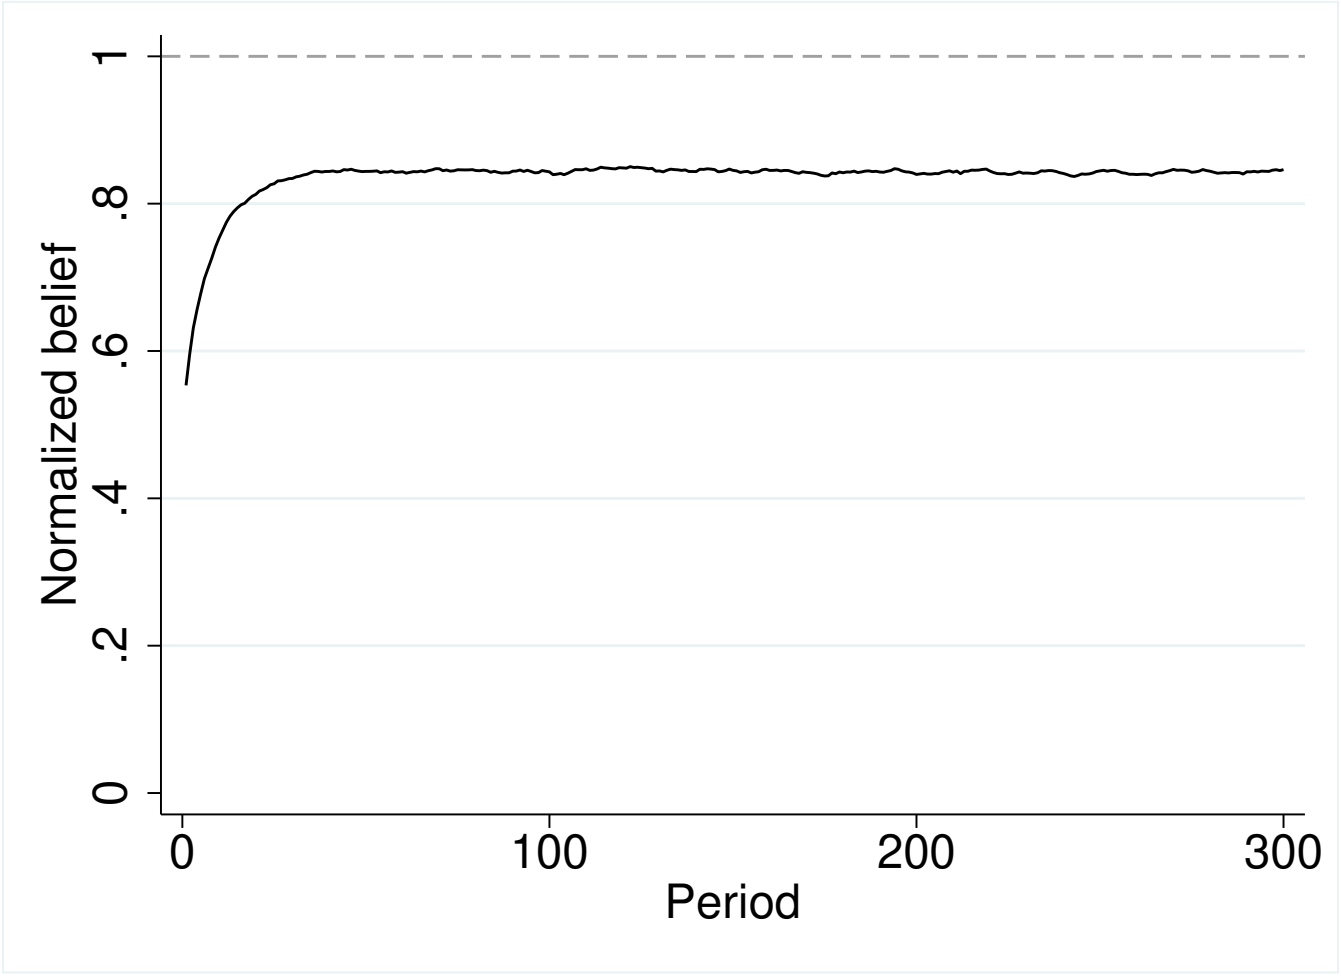

Prediction: The accuracy of higher-order beliefs does not improve any more after 300 than 30 signals.

The long treatment

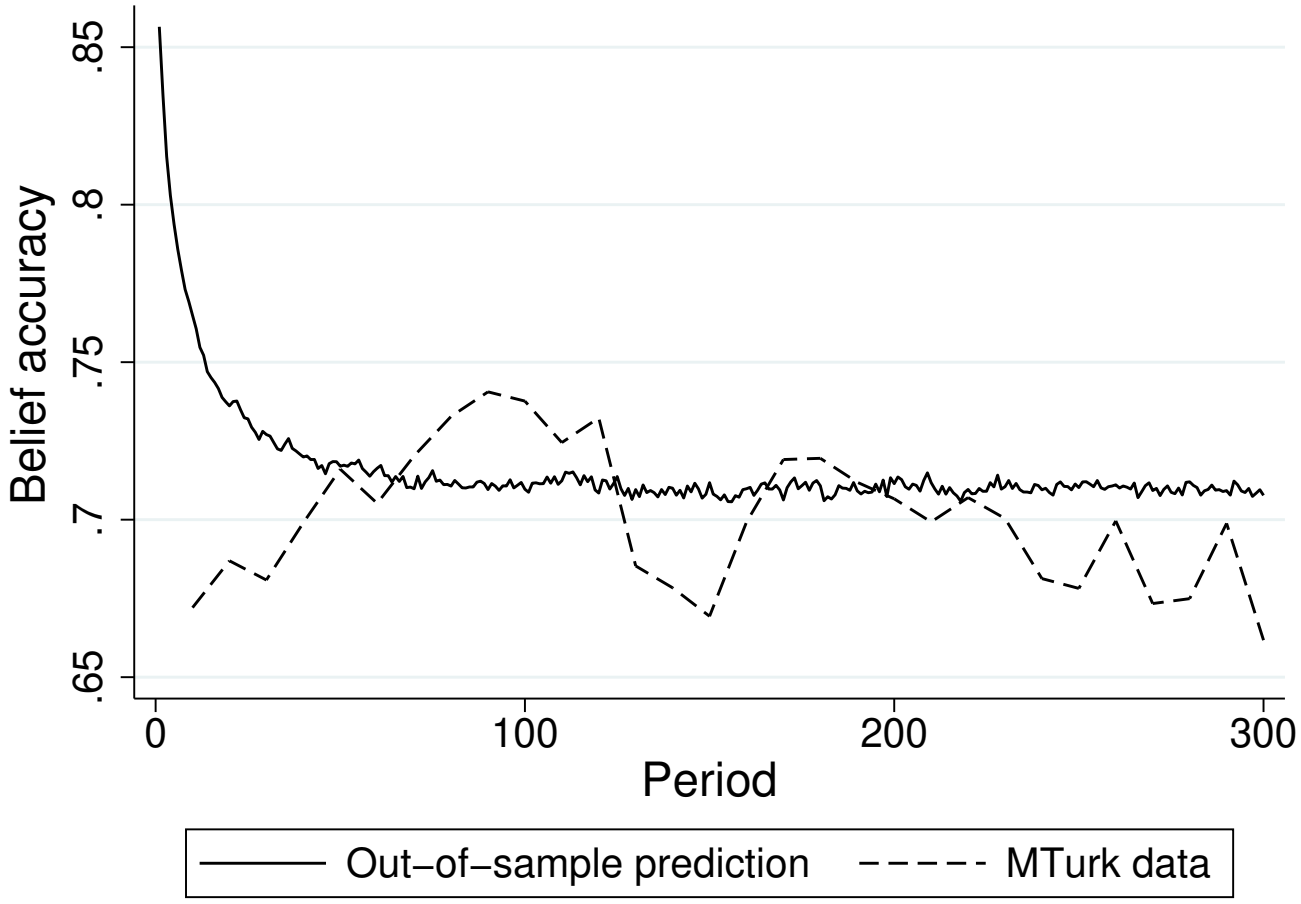

Discussion

- How can the failure of higher-order learning be mitigated?
  - Giving subjects information about others?
  - Feedback?
- Expand study to strategic setting