Supplementary information to the paper, “Knowing Me, Knowing You: An Experiment on Mutual Payoff Information in the Stag Hunt and Prisoner’s Dilemma”

Hazem Alshaikhmubarak, David Hales, Maria Kogelnik, Molly Schwarz, and C. Kent Strauss

Appendices

A. Extended Review of Related Experiments

Games under incomplete information. Most similar to our design is Feltovich and Oda (2014), in which subjects play incomplete information versions of six games, including the SH and PD, with treatments for random re-matching and fixed matching. Their results suggest that the matching mechanism does matter in the incomplete information environment, with fixed-pair often leading to increased coordination on pure strategy equilibria, higher payoffs, and faster convergence. Since the authors run no full-information treatments for comparison, however, the effect of mutual payoff information on behavior cannot be studied with their design. Furthermore, in contrast to Feltovich and Oda (2014), who primarily interested in how subjects learn under random re-matching versus with fixed partners when players have incomplete information, we employ a learning model to disentangle if the effect of mutual payoff information operates through initial play or the way subjects learn. Our study is also closely related to the work of McKelvey and Palfrey (2001), who run an ambitious number of games and information treatments, including different versions of the PD and the SH. Besides employing a "Playing in the Dark" treatment where subjects do not even observe their own payoffs, they have a full- and partial-information treatment similar to ours. Since choice frequencies are not reported by the authors, however, a direct comparison with our results is unfortunately not possible.

A closely related collection of experiments studies how subjects behave in incomplete information environments similar to ours. We are not aware of an experiment, however, that employs the same information treatments we use to the PD and the SH game. Mookherjee and Sopher (1994) find that in a matching pennies game with fixed partners, choice frequencies tend towards the unique mixed strategy equilibrium whether or not subjects are presented with opponent payoff information. Oechssler and Schipper (2003) have subjects play incomplete information treatments in the SH and PD; however, our experiments are difficult to compare as their design creates incentives for experimentation in the initial rounds by paying subjects more later and by providing rewards for correct answers regarding the opponents' payoffs. Perhaps most notable for our experiment, subjects' play converges towards Nash equilibrium despite the inability of subjects to fully perceive the game structure.

Additional studies have altered the information structure in ways that make them less comparable to our environment. For example, Cox et al. (2001) and Danz et al. (2012) inform subjects about a set of payoffs from which the actual opponent payoffs may be drawn, Nicklisch (2011) introduce information asymmetries into the environment, and Nikolaychuk (2012) match subjects with a computer following a learning algorithm and let them observe their own earnings after each round, or the whole payoff matrix, in versions of the PD, SH, and Battle of the Sexes. Andreoni et al. (2007) vary the information that bidders have about their rivals' valuation in first- and second-price auctions, and document that subjects' behavior in response to this information is consistent with theoretical predictions.

Payoff information and cooperation. Our paper also relates to a literature examining the impact of payoff information on the formation of cooperation. In both Friedman et al. (2015) and Huck et al. (2017), subjects play Cournot games that exhibit tension similar to our PD between competition and cooperation. Subjects do not have access to their own or others' payoff functions but are told that payoff functions are symmetric and time invariant, and they receive feedback on their own and others' payoffs and actions. Friedman et al. (2015) find that subjects in duopolies and triopolies choose highly competitive quantities in the first dozens of rounds, which is consistent with teoretical predictions and is echoed by our finding that most subjects choose to defect in PD-Partial. In contrast to our setting, they employ 1200 rounds of play and observe that subjects reach almost fully collusive levels in duopolies and converge toward collusive quantities in triopolies after a few hundred rounds of play. We leave it to future research to explore if subjects would eventually learn to cooperate in our PD environment after hundreds of rounds of play.

In addition to the setting where subjects do not have access to payoff functions, Huck et al. (2017) introduce a comparison treatment where subjects are shown the possible payoffs they could have received based on their partner's last action. Having access to this information facilitates cooperation during the first 30 rounds of play, but starting at around round 150 until the final round 600, they observe significantly less cooperation when subjects are given this information. In sum, while our findings are consistent with their findings in the first dozens of rounds, Huck et al. (2017) present evidence that payoff information can hinder cooperation in the long run. They argue that this is because having access to payoff information makes subjects less likely to adopt heuristics that foster cooperation.

Nax et al. (2016) study cooperation in a voluntary contribution game under different information structures. When the financial return from contributing to a public good is low, they find similar contribution rates across information treatments, however they document that when returns to contributing are high, more cooperation emerges when players have full information about the game than when they only get to observe their own payoffs. The latter is consistent with our finding that payoff information facilitates reaching the socially optimal outcome in both the SH and the PD.

Fiala and Suetens (2017) conduct a meta study on how information on payoffs and choices affects cooperative behavior in public good games and collusion games. The authors focus on different games, and are interested in repeated interactions of the same groups. They conclude that transparency regarding group earnings reduces contributions and collusion, while transparency about choices generally increases contributions and collusion, although the effect size varies across settings. In sum, the literature suggests that payoff information may help or hinder the formation of cooperation, depending on the context. Our study contributes to the existing literature by providing evidence that in the games we employ, mutual payoff information appears to facilitate cooperative behavior.

Equilibrium selection in the SH. Another strand of literature we contribute to addresses the question of which equilibrium play converges to in the SH. With one payoff-dominant and one risk-dominant equilibrium, the SH embodies a tradeoff between maximizing social efficiency and minimizing personal risk. Many experiments (e.g., Battalio et al., 2001; Schmidt et al., 2003; Dubois et al., 2012; Kendall, 2022) have been designed to better understand the conditions under which play converges to either equilibrium. A common feature of these studies is that subjects play SH games where payoffs are commonly known, and by varying these payoffs across games, diverse theoretical predictions can be disentangled. Our paper is related to this literature but takes a different approach; We keep payoffs constant across treatments but vary whether players get to observe the other's payoffs. Doing so allows us to identify mutual payoff information as an important factor for the payoff-dominant equilibrium to arise in the SH.

B. Further Details: Learning Model and Simulations

B.1. Model choice considerations

The main advantage of using the EWA model in our setting is that it allows us to investigate whether the effect of having mutual payoff information operates through initial play, ongoing learning, or both. This is our purpose in estimating the model (rather than to test what models match the data most accurately). In describing the EWA learning model, Camerer and Ho (1999) comment that they "consider the scientific problem of figuring out how people choose their initial strategies as being fundamentally different than explaining how they learn." The model we apply is a belief-learning model, more precisely described as weighted fictitious play in which the initial attraction is assigned an "observation equivalency" of 1 period of play.[1,2] We note that belief-learning models generally perform favorably compared with alternatives; for example, see Nyarko and Schotter (2003).

B.2. Estimation and simulation details

Parameter estimation. The estimation is performed numerically, using maximum likelihood techniques. We employ the bootstrap procedure for estimating parameter sampling distributions, with $B = 2,000$ bootstrap samples per estimation, from which we then conduct inference using the Bias Corrected-accelerated (BCa) confidence interval method pioneered by Efron and Tibshirani (1993).[3]

Consider a treatment with N subjects and $T = 40$ rounds. Then, the likelihood of observing subject i's action history $\{a_i(1), a_i(2), ..., a_i(T-1), a_i(T)\}$, given $(A^x(0), A^y(0), \phi, \lambda)$ is

$$\Pi_{t=1}^{T} P_i^{a_i(t)}(t|A^x(0), A^y(0), \phi, \lambda), \tag{1}$$

where $P_i^{a_i(t)}$ corresponds to the logistic probability function defined in equation (**??**). The joint likelihood function $\mathcal{L}(A^x(0), A^y(0), \phi, \lambda)$ of observing all subjects' action histories is given by

$$\mathcal{L}(A^x(0), A^y(0), \phi, \lambda) = \Pi_i^N \{\Pi_{t=1}^{T} P_i^{a_i(t)}(t|A^x(0), A^y(0), \phi, \lambda)\}.$$

Simulations. To test whether our parameter estimates lead to predicted behavior that is consistent with actual observed behavior, we conduct simulations of $1,000$ sessions per treatment. Each one of these sessions has an even number of subjects between 14 and 20, chosen randomly. (The simulation results are robust to changes in the number of subjects in each session.)

[1] Specific parameteric restrictions we applied to the EWA model are: $\delta = 1$, $\rho = \phi$, and $N(0) = 1$. In introducing the EWA model, Camerer and Ho (1998) state that "substituting $\delta = 1$ and $\rho = \phi$ into the attraction updating equation...gives EWA attractions which are equal to updated expected payoffs using weighted fictitious play." Earlier in the same paper, the authors state that "$N(0)$ can be interpreted as the number of periods of actual experience which is equivalent in attraction impact to the pregame thinking." The implication of our parameter restriction of $N(0) = 1$ is thus that we implicitly assume our subjects assign an equivalent amount of "attraction impact" to their initial process of thinking as they do to each subsequent period's post-play process of observation and reflection. We note the authors also provide an example in which $N(0) = 1$, which we take to imply that this is a reasonable assignment.

[2] This version of EWA is similar to earlier belief learning and fictitious play models. See for example Mookherjee and Sopher (1997) and Cheung and Friedman (1997), and Fudenberg et al. (1998).

[3] More commonly used standard error techniques are inapplicable to our dataset, due to the extreme levels of skew and kurtosis in the sample distribution of bootstrapped parameter estimates. These high levels of 3rd and 4th moments render the normality assumptions underlying standard error techniques invalid, and necessitate adoption of a more robust approach to inference (e.g. the BCa confidence interval method).

C. Experimental Materials

C.1. Instructions

Welcome!

You are about to participate in an experiment on decision-making. In this experiment, you can earn a considerable amount of **money**, which will be paid to you in cash, privately, at the end of the experiment. How much you earn will depend on your decisions, the decisions of other participants, and chance.

Please **do not communicate** with the other participants at any point during the experiment. Make sure that your phone is turned off now.

To make sure that everybody understands the tasks in this experiment, we will begin with some basic instructions. **If you have any questions**, or need assistance of any kind, raise your hand and the experimenter will come and help you.

Basic Instructions

At the beginning of the experiment, the computer will randomly assign each participant to one of **two groups: *Group A* and *Group B***. Once assigned, participants will remain in the same group throughout the entire experiment. There will be the same number of participants in each group.

Participants assigned to the **same group** will each see the **same payoff table and other information** on their computer screen at the beginning of the experiment.

The experiment consists of **2 parts**. For now, we will explain to you what is happening in the first part of the experiment. Once the first part is completed, we will explain to you what is happening in the second part.

The first part of the experiment consists of **40 rounds**. In each round, you will be **randomly paired** with a participant from the other group. You will not know who of the other participants is assigned to which group, and you will also not know with whom you are randomly paired in any given round. In each round, it is equally likely that you will be paired with any of the participants from the other group. You will never be paired with somebody from your own group.

In each of the 40 rounds, you and the person you are currently paired with will be asked to make a decision on the computer. In what follows, we will explain to you how you can make these decisions.

The Decision Tasks

In each of the 40 rounds, you will be able to **choose one of two actions**. The participant you are paired with will also be able to choose one out of two actions. In each round, everybody will have to **choose an action before seeing the action that the other participant has chosen**.

Below, we show you an example of how a decision task could look like on the computer. In the experiment, you will see a similar table on your computer screen, but with different numbers.

Example of a Payoff Table

The **table below shows the payoffs** associated with each combination of your choice and the choice of the participant you are paired with. This is an example of how a decision task could look like on the computer; please note that the actual numbers you will see in the experiment will be different from those shown in this example. We will now explain to you how you can interpret the numbers in the table.

The **first entry** in each cell (i.e. the number before the comma) represents **your payoff**. The **second entry** in each cell (i.e. the number after the comma) represents the **payoff of the person you are paired with**.

Please choose a row			
	The Other's Choice		
Your Choice		B1	B2
	A1	5, 7	4, 3
	A2	2, 4	10, 6

All **cell entries** of the table show the **payoffs that are associated with each combination of your choice and the other participant's choice**:

- For example, if you select "A1" and the other participant selects "B1", you earn 5 Dollars and the other participant earns 7 Dollars.
- As another example, if you select "A2" and the other participant selects "B1", you earn 2 Dollars and the other participant earns 4 Dollars.
- Another example: if you select "A1" and the other participant selects "B2", you earn 4 Dollars and the other participant earns 3 Dollars.
- Another example: if you select "A2" and the other participant selects "B2", you earn 10 Dollars and the other participant earns 6 Dollars.

How to Make Decisions

Suppose that the computer assigned you to *Group A*. In this example, you will be asked to choose either "A1" or "A2". Remember that if you are in *Group A*, then in each round, you will be paired with somebody from *Group B*, and they will be asked to choose either "B1" or "B2". If you should get assigned to *Group B*, however, then **you** will be asked to choose between "B1" and "B2", and the **other participant** will be asked to choose between "A1" and "A2", as they would be assigned to *Group A*.

In the experiment, you will see a table similar to the example above on your computer screen. **To make a choice**, you will **click on one of the rows in the table**.

Once you select a row, it will change color and a red *SUBMIT* button will appear. Your choice will be finalized once you click on the *SUBMIT* button. After submitting your choice, you will need to wait until the other participant you are paired with has also made their choice. Once you and the participant you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear. **Remember that you will only see the choice of the other participant once you have submitted your own choice.** After each of the 40 rounds, you will see an overview of the choice you made, the choice the other participant made, and your payoffs of the round.

Example of How a Payoff Table Will Look Like in the First Part of the Experiment

In the example above, we showed you an example of a decision task. In each cell of the table above, you could see both your own payoff and the other's payoff for each combination of your and the other participant's choices. In the actual experiment, however, **you will only see your own payoffs**. The payoffs of the other person will be covered.

Here is an example of what a table in the experiment could actually look like:

Please choose a row

		The Other's Choice	
		B1	B2
Your Choice	A1	5,	4,
	A2	2,	10,

This **table shows your payoffs** associated with each combination of your choice and the choice of the participant you are paired with. It does not show, however, the payoffs of the person you are paired with. The **first entry in each cell** (i.e. the number before the comma sign) **represents your payoff**. For the first part of the experiment, you will never know how much the person you are paired with earns in each combination of your and their choice.

- For example, if you select "A1" and the other participant selects "B1", you earn 5 Dollars, but you don't know how much the other participant earns.
- As another example, if you select "A2" and the other participant selects "B1", you earn 2 Dollars, but you don't know how much the other participant earns.
- Another example: if you select "A1" and the other participant selects "B2", you earn 4 Dollars, but you don't know how much the other participant earns.
- Another example: if you select "A2" and the other participant selects "B2", you earn 10 Dollars, but you don't know how much the other participant earns.

Summary: What It Means to Not See the Other's Payoffs

1. You only **know your own payoffs**, and you know that everybody in your group has the same payoffs.
2. You **do not know** the payoffs of the person you are paired with.
3. This means that you do not know what payoffs participants in the other group are getting. You know, however, that every participant in the other group is getting the same payoffs.

How much will you get paid in the end?

At the end of the experiment, for each participant the computer will randomly select a number between 1 and 80, corresponding to each of the rounds of the experiment. Every participant will get paid, in US Dollars, the amount of their payoff in that particular round, PLUS the show-up fee of 7 dollars. Before that, a short questionnaire will appear on your screen.

Summary

- There are a total of 80 rounds in the experiment, divided into two parts of 40 rounds each.
- You will make a decision in each of these 80 rounds.
- **At the beginning of the experiment**, half of all participants will be randomly assigned to *Group A*, and the other half will be assigned to *Group B*.
- Participants stay assigned to the same group throughout the experiment.
- All *Group A* participants have the same payoffs, and all *Group B* participants have the same payoffs. These payoffs remain the same throughout all 40 rounds of that part of the experiment.
- **In each round**, you will be randomly paired with someone from the other group.
- This means that before each decision round, a new random pair will be formed.

Comprehension Quiz First Game (40 Rounds)

Part II *(Handed out to subjects after they completed the first 40 rounds of the experiment.)*

The second part of the experiment has a similar setup to the first part. You will again be presented with a payoff table and will be asked to make choices by clicking on the rows of the table. This part of the experiment will consist of another **40 rounds**. As before, in each round, you will be **randomly paired** with a participant from the other group.

The table you see in this part of the experiment **shows both your payoff and the payoff of the person you are paired with**.

Example of a Payoff Table

The **first entry** in each cell (i.e. the number before the comma) represents **your payoff**. The **second entry** in each cell (i.e. the number after the comma) represents the **payoff of the person you are paired with**.

Please choose a row

		The Other's Choice	
		B1	B2
Your Choice	A1	5, 7	4, 3
	A2	2, 4	10, 6

Before we begin, let me briefly remind you of the following:

- Once you select a row, you need to click on the red *SUBMIT* button to confirm your choice.
- After that, please don't forget to press the *MOVE ON* button so that the next round of the experiment can begin.
- Remember that **all participants have to make their choice before they can observe the choice of the person they are paired with**.
- After completing all 40 rounds, a short questionnaire will appear. You will get paid after that.

C.2. Comprehension Quiz

Partial Information - Version 1

<u>Comprehension Quiz</u>

To make sure that you understand the instructions of this experiment, please answer the questions below.

Below, you see an example of a decision task, similar to the one that you might encounter in the experiment. In this example, you got assigned to Group A, and the person you are randomly paired with got assigned to Group B.

Please choose a row

		The Other's Choice	
		B1	B2
Your Choice	A1	2,	3,
	A2	7,	5,

Please answer the following questions. If you don't know the answer for sure, please insert a question mark (?) into the blank space.

1. If you choose "A2" and the other chooses "B2", what is the other's payoff?
2. If you choose "A2" and the other chooses "B1", what is your payoff?
3. If you choose "A1" and the other chooses "B2", what is your payoff?
4. If you choose "A2" and the other chooses "B1", what is the other's payoff?
5. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 5?
6. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 7?

Full Information - Version 1

Comprehension Quiz

To make sure that you understand the instructions of this experiment, please answer the questions below.

Below, you see an example of a decision task, similar to the one that you might encounter in the experiment. In this example, you got assigned to Group A, and the person you are randomly paired with got assigned to Group B.

Please choose a row

		The Other's Choice	
		B1	B2
Your Choice	A1	2, 0	3, 3
	A2	7, 11	5, 3

Please answer the following questions. If you don't know the answer for sure, please insert a question mark (?) into the blank space.

1. If you choose "A2" and the other chooses "B2", what is the other's payoff?
2. If you choose "A2" and the other chooses "B1", what is your payoff?
3. If you choose "A1" and the other chooses "B2", what is your payoff?
4. If you choose "A2" and the other chooses "B1", what is the other's payoff?
5. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 5?
6. In this example, which combination of your action and the other's action needs to happen so that the other gets a payoff of 3?

D. Additional Figures and Tables

Figure D1: Screenshots of Experimental Interface, PD-Partial Treatment

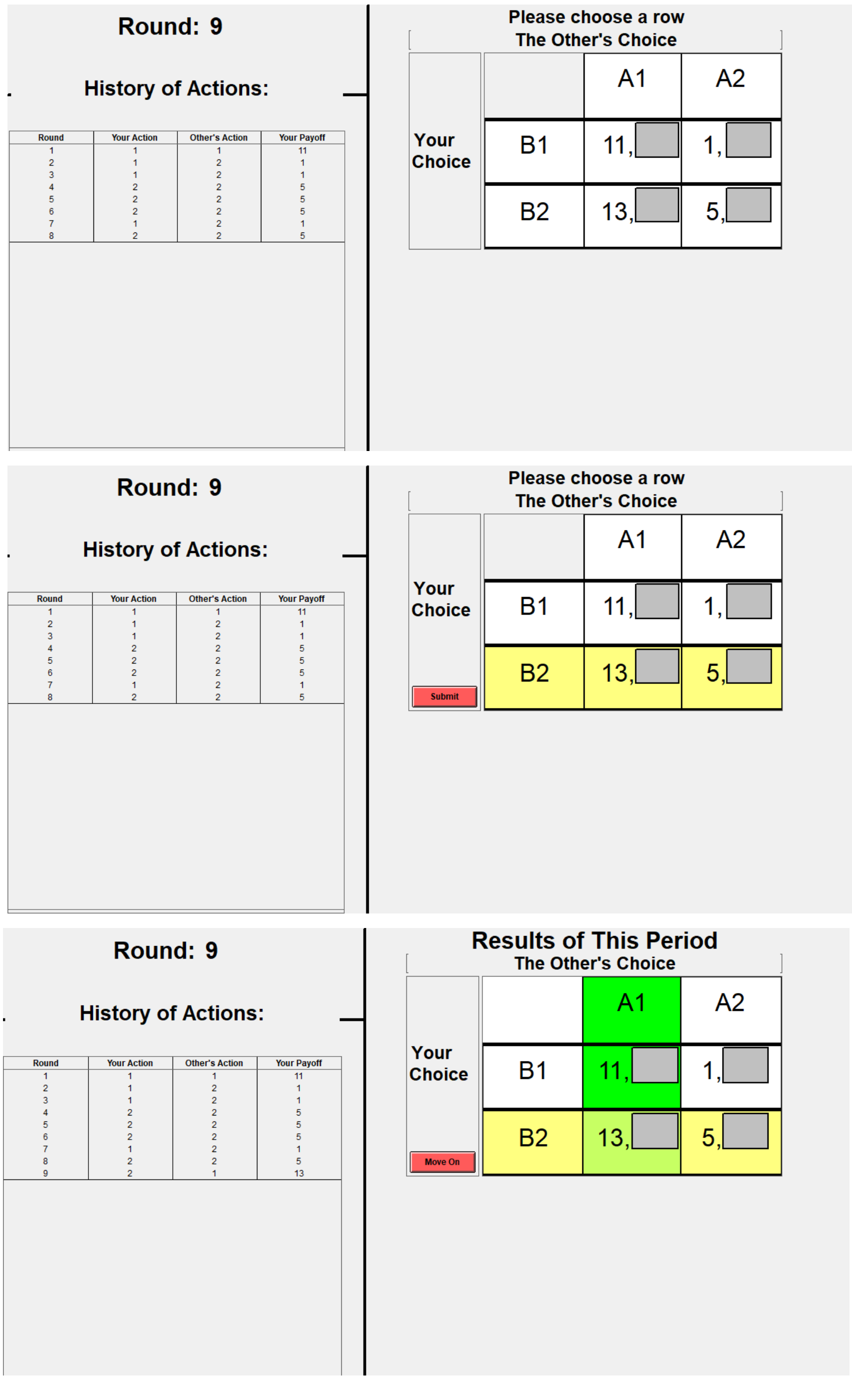

Note: This figure shows screenshots of the experimental interface of the PD-Partial treatment for a subject that was assigned to Group B, before (upper panel), during (middle panel), and after (bottom panel) choosing their action in a round. See Figure D2 for corresponding screenshots from the SH-Full treatment.

Figure D2: Screenshots of Experimental Interface, SH-Full Treatment

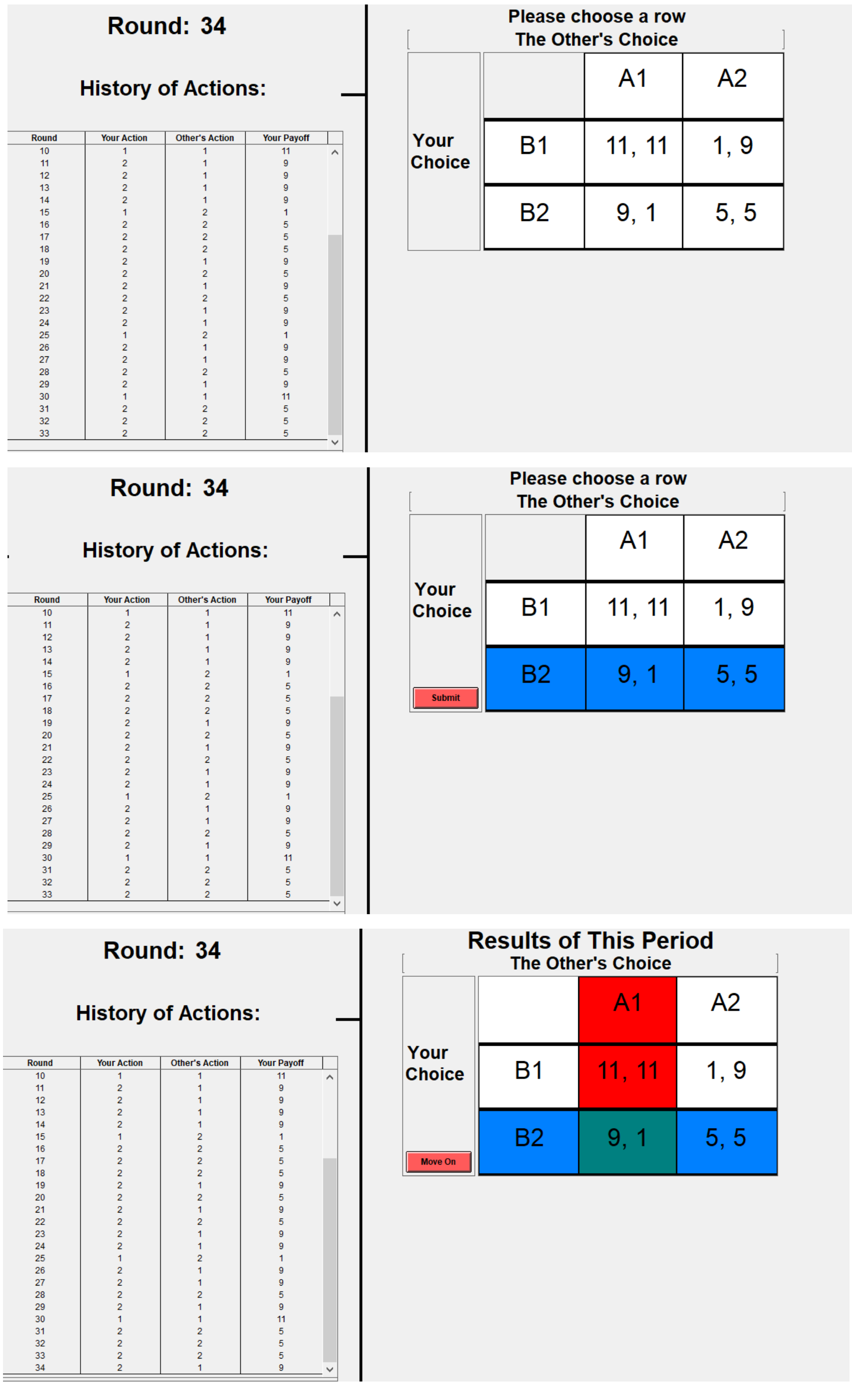

Note: We used a different color scheme for each of the two games in each experiment. We did this to help subjects understand and recall, when they are playing the second game of the experiment, that the current game they are playing is distinct from the first game they have already completed. See Figure D1 for the corresponding screenshots of the PD-Partial treatment.

Figure D3: Rate of Playing X by Session and Order

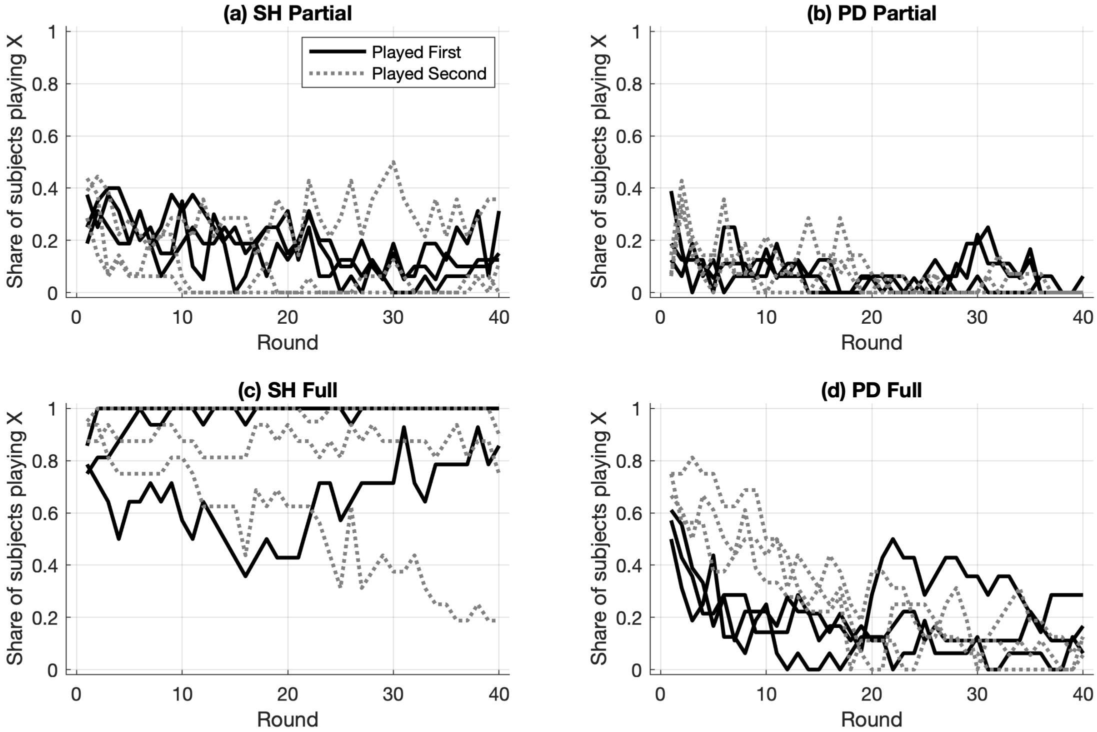

Note: This figure shows the share of subjects playing action X by session and order over the 40 rounds, separately for each game and treatment. The solid lines show data from sessions where the represented treatment was played in the first 40 rounds of the experimental session, and the dotted lines show data from sessions where the represented treatment was played in the last 40 rounds of the experimental session. Each line represents the mean rate of the action X. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD).

Figure D4: Mean Action X Rates in Pooled Data vs. Simulations

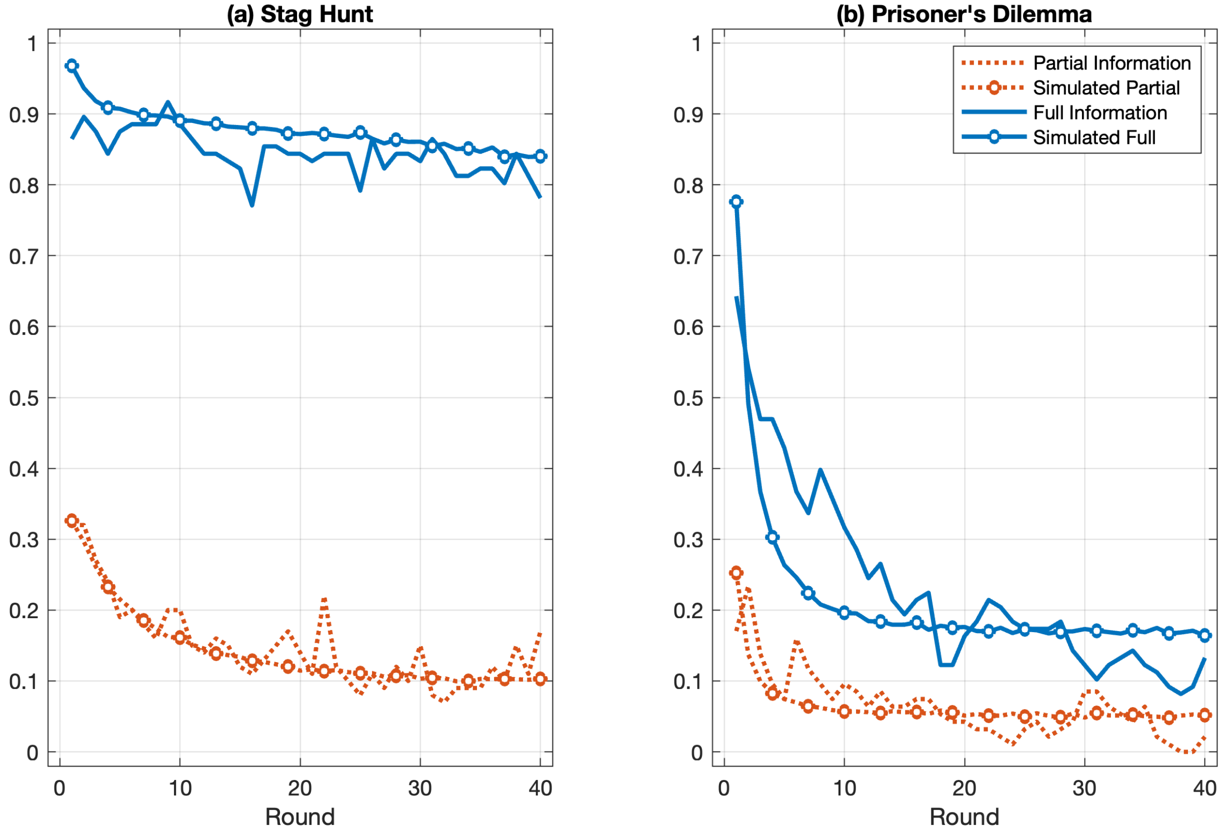

Note: Solid lines show the shares of subjects playing action X in our observed data and dashed lines show the same shares in simulated data of 1,000 sessions per treatment. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD).

Figure D5: Simulated Mean Action X Rates with Swapped Model Parameters

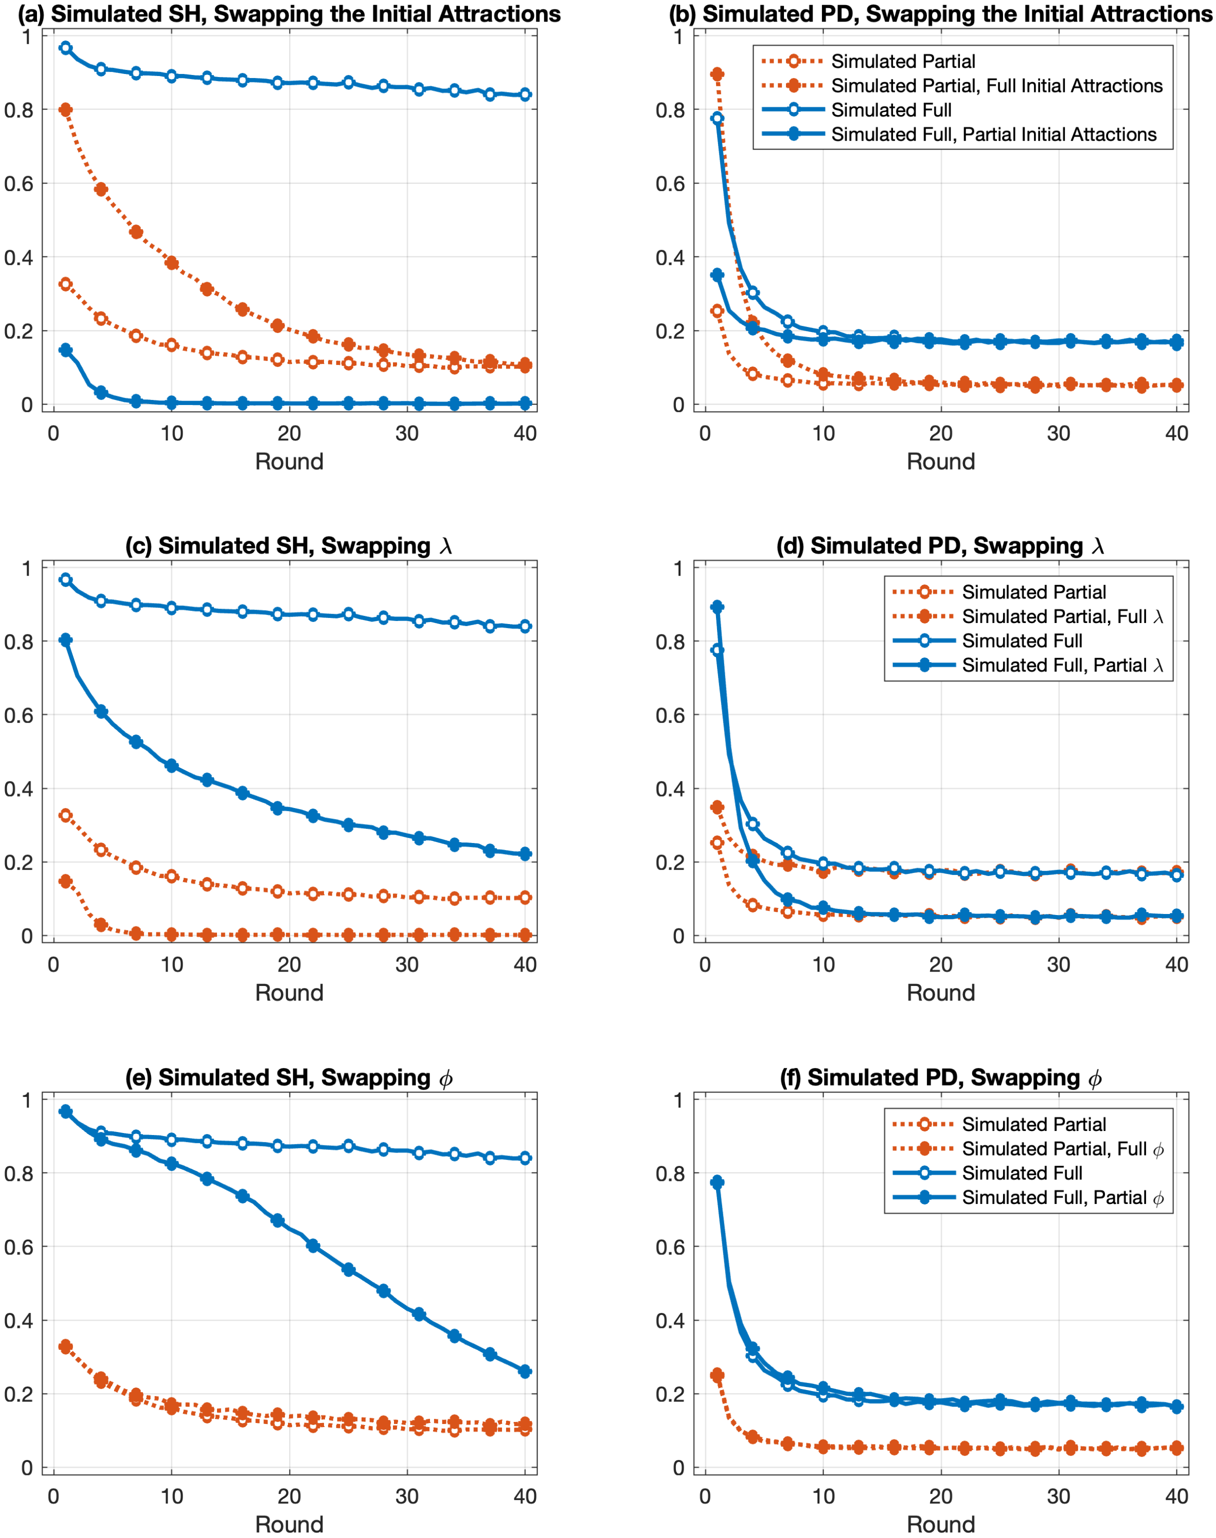

Note: Mean action X rates in simulated data of 1000 sessions per treatment, versus same rates with swapped learning model parameters from the other treatment. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD).

Table D1: Effect of Partial-information Treatment for Selection of Action X, Within-Subjects

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) SH					
Partial	-0.643	-0.646	-0.648	-0.663	-0.613
	(0.043)	(0.038)	(0.051)	(0.050)	(0.062)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Full mean	0.817	0.906	0.844	0.794	0.725
Number of clusters	52	52	52	52	52
N	4,160	1,040	1,040	1,040	1,040
b) PD					
Partial	-0.203	-0.460	-0.220	-0.100	-0.034
	(0.023)	(0.051)	(0.037)	(0.030)	(0.020)
Cluster p-value	0.000	0.000	0.000	0.002	0.088
Full mean	0.271	0.576	0.272	0.154	0.082
Number of clusters	50	50	50	50	50
N	4,000	1,000	1,000	1,000	1,000

Note: The sample uses the within-subjects data. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD). The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D2: Effect of Partial-information Treatment for Selection of Action X, Between-Subjects

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) SH					
Partial	-0.802	-0.659	-0.787	-0.881	-0.881
	(0.028)	(0.052)	(0.044)	(0.025)	(0.032)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Control mean	0.876	0.852	0.830	0.886	0.936
Number of clusters	96	96	96	96	96
N	3,840	960	960	960	960
b) PD					
Partial	-0.076	-0.157	-0.056	-0.057	-0.036
	(0.037)	(0.046)	(0.038)	(0.041)	(0.040)
Cluster p-value	0.042	0.001	0.149	0.173	0.371
Control mean	0.191	0.283	0.135	0.198	0.146
Number of clusters	98	98	98	98	98
N	3,920	980	980	980	980

Note: The sample uses the between-subjects data. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD). The regression includes controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D3: Order Effects for Stag Hunt - Choosing Action X

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) SH - Partial					
Game played first	0.053	0.066	0.113	0.022	0.009
	(0.034)	(0.044)	(0.043)	(0.043)	(0.045)
Cluster p-value	0.125	0.139	0.011	0.600	0.833
Outcome mean, 2nd game	0.122	0.194	0.083	0.108	0.102
Number of clusters	100	100	100	100	100
N	4,000	1,000	1,000	1,000	1,000
b) SH - Full					
Game played first	0.059	-0.053	-0.015	0.092	0.211
	(0.051)	(0.045)	(0.064)	(0.057)	(0.061)
Cluster p-value	0.247	0.236	0.818	0.108	0.001
Outcome mean, 2nd game	0.817	0.906	0.844	0.794	0.725
Number of clusters	96	96	96	96	96
N	3,840	960	960	960	960

Note: Sample includes all SH sessions. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD). Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D4: Order Effects for Prisoner's Dilemma - Choosing Action X

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) PD - Partial					
Game played first	0.004	-0.020	-0.028	0.040	0.025
	(0.015)	(0.031)	(0.022)	(0.016)	(0.020)
Cluster p-value	0.766	0.512	0.208	0.015	0.205
Outcome mean, 2nd game	0.063	0.136	0.080	0.014	0.023
Number of clusters	94	94	94	94	94
N	3,760	940	940	940	940
b) PD - Full					
Game played first	-0.080	-0.293	-0.137	0.044	0.064
	(0.043)	(0.061)	(0.051)	(0.056)	(0.047)
Cluster p-value	0.064	0.000	0.009	0.431	0.179
Outcome mean, 2nd game	0.271	0.576	0.272	0.154	0.082
Number of clusters	98	98	98	98	98
N	3,920	980	980	980	980

Note: Sample includes all PD sessions. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD). Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D5: Effect of Partial Information Treatment for Selection of Action X Using Logit Model

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) SH					
Partial	-3.387	-3.427	-3.361	-3.515	-3.646
	(0.252)	(0.264)	(0.307)	(0.299)	(0.355)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Marginal effect	-0.427	-0.474	-0.419	-0.406	-0.411
Control mean	0.844	0.881	0.838	0.836	0.822
Number of clusters	144	144	144	144	144
N	7,840	1,960	1,960	1,960	1,960
b) PD					
Partial	-1.523	-1.652	-1.376	-1.970	-1.395
	(0.166)	(0.192)	(0.234)	(0.362)	(0.395)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Marginal effect	-0.184	-0.295	-0.155	-0.175	-0.094
Control mean	0.232	0.433	0.205	0.176	0.113
Number of clusters	142	142	142	142	142
N	7,680	1,920	1,920	1,920	1,920

Note: The sample uses the pooled data. Action X is associated with the socially optimal outcome in both games (payoff-dominant equilibrium in SH and cooperation in PD). The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D6: Share of SH Subjects Choosing Close to Mixed Strategy Nash Equilibrium

	Rounds				
	Over-all	1-10	11-20	21-30	31-40
a) Partial					
Fraction playing within .10 of MSNE	0.17	0.27	0.16	0.14	0.07
b) Full					
Fraction playing within .10 of MSNE	0.22	0.38	0.20	0.18	0.08

Note: Given that the mixed strategy Nash equilibrium is to play action X 66.67% of the time, the fraction of subjects playing within 10 percentage points of the mixed strategy Nash equilibrium (MSNE) equals the total number of subjects playing action X between 56.67% and 76.67% of the time divided by the total number of subjects exposed to that information treatment.

Table D7: Effect of Partial-information Treatment for Efficiency Ratio

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) Stag Hunt					
Partial	-0.396	-0.412	-0.386	-0.390	-0.394
	(0.018)	(0.015)	(0.021)	(0.026)	(0.021)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Full-information mean	0.864	0.883	0.856	0.857	0.859
Number of clusters	144	144	144	144	144
N	7,840	1,960	1,960	1,960	1,960
b) Prisoner's Dilemma					
Partial	-0.078	-0.147	-0.063	-0.066	-0.034
	(0.011)	(0.020)	(0.014)	(0.015)	(0.015)
Cluster p-value	0.000	0.000	0.000	0.000	0.025
Full-information mean	0.554	0.650	0.539	0.527	0.498
Number of clusters	142	142	142	142	142
N	7,680	1,920	1,920	1,920	1,920

Note: The sample uses the pooled data. The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table D8: Comparison of Estimates of λ by Game and Information Treatment

Estimate	λ in the SH	λ in the PD
Full	1.5831	0.4297
	(1.2159 – 1.8776)	(0.3347 – 0.5242)
Partial	0.6516	0.7426
	(0.5184 – 0.7673)	(0.6615 – 0.9146)
$\underline{H_0 : \lambda_{partial} \geq \lambda_{full}}$		
BCa interval test p-val	0.0014	
$\underline{H_0 : \lambda_{full} \geq \lambda_{partial}}$		
BCa interval test p-val		0.0368
Mann-Whitney U test	0.0018	0.0018
Kolmogorov-Smirnov test	0.0000	0.0000

Note: One-sided BCa interval tests conducted using $B = 10,000$ bootstrap iterations, with bootstrapping being performed separately for full- and partial-information observations

Table D9: Comparison of Estimates of ϕ by Game and Information Treatment

Estimate	ϕ in the SH	ϕ in the PD
Full	0.9649	0.8011
	(0.9011 – 1.0223)	(0.3009 – 0.9398)
Partial	0.8290	0.8769
	(0.7112 – 0.9911)	(0.6019 – 85.013)
$\underline{H_0 : \phi_{partial} \geq \phi_{full}}$		
BCa interval test p-val	0.0624	
$\underline{H_0 : \phi_{full} \geq \phi_{partial}}$		
BCa interval test p-val		0.3005
Mann-Whitney U test	0.0017	0.0014
Kolmogorov-Smirnov test	0.0000	0.0000

Note: One-sided BCa interval tests conducted using $B = 10,000$ bootstrap iterations, with bootstrapping being performed separately for Full- and Partial-information observations.

Table D10: Comparison of Observed and Simulated Data - Rate of Choosing Action X

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) SH-Partial					
Simulated	-0.007	-0.007	-0.010	-0.008	-0.002
	(0.028)	(0.021)	(0.040)	(0.039)	(0.037)
Cluster p-value	0.805	0.755	0.798	0.832	0.946
Bootstrap p-value	0.820	0.778	0.810	0.854	0.961
Outcome mean	0.149	0.228	0.142	0.120	0.107
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	686,400	171,600	171,600	171,600	171,600
b) SH-Full					
Simulated	0.032	0.031	0.045	0.028	0.024
	(0.073)	(0.045)	(0.079)	(0.083)	(0.108)
Cluster p-value	0.663	0.496	0.574	0.735	0.825
Bootstrap p-value	0.714	0.573	0.597	0.789	0.916
Outcome mean	0.844	0.881	0.838	0.836	0.822
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	683,120	170,780	170,780	170,780	170,780
c) PD-Partial					
Simulated	-0.002	-0.029	-0.010	0.017	0.016
	(0.006)	(0.011)	(0.012)	(0.010)	(0.010)
Cluster p-value	0.814	0.006	0.432	0.083	0.108
Bootstrap p-value	0.820	0.069	0.588	0.146	0.194
Outcome mean	0.065	0.126	0.065	0.035	0.036
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	687,360	171,840	171,840	171,840	171,840
d) PD-Full					
Simulated	-0.018	-0.102	-0.025	-0.003	0.058
	(0.028)	(0.064)	(0.036)	(0.040)	(0.034)
Cluster p-value	0.522	0.114	0.496	0.931	0.084
Bootstrap p-value	0.577	0.225	0.539	0.960	0.148
Outcome mean	0.232	0.433	0.205	0.176	0.113
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	686,080	171,520	171,520	171,520	171,520

Note: Action X is associated with the socially optimal outcome in both games. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. The level of clustering is at the session. Cluster p-value indicates the p-value from a two-sided t-test of the Null hypothesis that the difference in the rate of playing X between the simulated data and observed data is zero using the cluster-robust standard error. However, the asymptotic properties of the cluster-robust standard error and resulting test statistic rest on the assumption that the effective number of clusters is approaching infinity. When the number of effective clusters is low, the cluster-robust method may bias the standard errors downward, resulting in inflated type-I error that exceeds the nominal size of the test. Though we have a large number of clusters, the effective number is below 50, indicating this issue might be occurring. Thus, the wild bootstrap method with 999 replications and Webb weights is instead recommended for inference. Bootstrap p-value indicates the p-value from the empirical sampling distribution found with the bootstrapping method.

References

Andreoni, J., Y.-K. Che, and J. Kim (2007). Asymmetric information about rivals' types in standard auctions: An experiment. *Games and Economic Behavior 59*(2), 240–259.

Battalio, R., L. Samuelson, and J. Van Huyck (2001). Optimization incentives and coordination failure in laboratory Stag Hunt games. *Econometrica 69*(3), 749–764.

Camerer, C. and T.-H. Ho (1998). Experience-weighted attraction learning in coordination games: Probability rules, heterogeneity, and time-variation. *Journal of Mathematical Psychology 42*(2-3), 305–326.

Camerer, C. and T.-H. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica 67*(4), 827–874.

Cheung, Y.-W. and D. Friedman (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior 19*(1), 46–76.

Cox, J. C., J. Shachat, and M. Walker (2001). An experiment to evaluate Bayesian learning of Nash Equilibrium play. *Games and Economic Behavior 33*, 11–33.

Danz, D. N., D. Fehr, and D. Kübler (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics 15*(4), 622–640.

Dubois, D., M. Willinger, and P. Van Nguyen (2012). Optimization incentive and relative riskiness in experimental Stag-Hunt games. *International Journal of Game Theory 41*(2), 369–380.

Efron, B. and R. Tibshirani (1993). *An Introduction to the Bootstrap*. Chapman and Hall.

Feltovich, N. and S. H. Oda (2014). Effect of matching mechanism on learning in games played under limited information. *Pacific Economic Review 3*, 260–277.

Fiala, L. and S. Suetens (2017). Transparency and cooperation in repeated dilemma games: a meta study. *Experimental Economics 20*, 755–771.

Friedman, D., S. Huck, R. Oprea, and S. Weidenholzer (2015). From imitation to collusion: Long-run learning in a low-information environment. *Journal of Economic Theory 155*, 185–205.

Fudenberg, D., F. Drew, D. K. Levine, and D. K. Levine (1998). *The Theory of Learning in Games*, Volume 2. MIT Press.

Huck, S., J. Leutgeb, and R. Oprea (2017). Payoff information hampers the evolution of cooperation. *Nature Communications 8*, 1–5.

Kendall, R. (2022). Decomposing coordination failure in stag hunt games. *Experimental Economics 25*(4), 1109–1145.

McKelvey, R. D. and T. R. Palfrey (2001). Playing in the dark: Information, learning, and coordination in repeated games. Working paper.

Mookherjee, D. and B. Sopher (1994). Learning behavior in an experimental Matching Pennies game. *Games and Economic Behavior 7*(1), 62–91.

Mookherjee, D. and B. Sopher (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior 19*(1), 97–132.

Nax, H. H., M. N. Burton-Chellew, S. A. West, and H. P. Young (2016). Learning in a black box. *Journal of Economic Behavior & Organization 127*, 1–15.

Nicklisch, A. (2011). Learning strategic environments: an experimental study of strategy formation and transfer. *Theory and Decision 71*(4), 539–558.

Nikolaychuk, O. (2012). Does it pay to know more in games of incomplete information? Working paper.

Nyarko, Y. and A. Schotter (2003, December). An experimental study of belief learning using elicited beliefs. *Econometrica 70*(3), 971–1005.

Oechssler, J. and B. Schipper (2003). Can you guess the game you are playing? *Games and Economic Behavior 43*, 137–152.

Schmidt, D., R. Shupp, J. M. Walker, and E. Ostrom (2003). Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior 42*(2), 281–299.