Online Appendices to: Testing Models of Strategic Uncertainty: Equilibrium Selection in Repeated Games

Marta Boczoń
Copenhagen Business School

Emanuel Vespa
UC San Diego

Taylor Weidman
University of Pittsburgh

Alistair J. Wilson
University of Pittsburgh

June 2024

Abstract
We provide additional materials that support our paper "Testing Strategic Uncertainty..." in the Jornal of the European Economic Association. The Online Appendices contain:
A. Additional Tables and Figures;
B. Further Analysis of the Within-Subject Treatments;
C. General Basin of Attraction Calculations;
D. Estimation for the partial correlation σ term;
E. Interface Screenshots;
F. Experimental Instructions;
G. Strategy Frequency Estimation.

Acknowledgments: We would like to thank: David Cooper, Guillaume Fréchette, Daniella Puzzello, Giancarlo Spagnolo, and Lise Vesterlund. This research was funded with support from the National Science Foundation (SES:1629193)

E-mail: mbo.eco@cbs.dk (Boczoń); evespa@ucsd.edu (Vespa); taylorjweidman@pitt.edu (Weidman); alistair@pitt.edu (Wilson)

Appendix A: Additional Tables and Figures

TABLE A.1. Cooperation and success rates across all supergames

Action and signal rates	$X = \$9$		$X = \$1$	
	$N = 2$	$N = 4$	$N = 4$	$N = 10$
	Cooperation			
Initial	0.466 (0.046)	0.100 (0.021)	0.719 (0.039)	0.457 (0.044)
Ongoing	0.296 (0.029)	0.044 (0.012)	0.433 (0.034)	0.245 (0.039)
	Success			
Initial	0.466	0.003	0.408	0.010
Ongoing	0.296	0.002	0.275	0.009

Results are calculated using data from all supergames, with subject-clustered standard errors in parentheses. Cooperation rates present raw proportions.

TABLE A.2. Cooperation in reaction to previous round's history

History	$X = \$9$		$X = \$1$		Chat ($X = \$9, N = 4$)	
	$N = 2$	$N = 4$	$N = 4$	$N = 10$	$\delta = 3/4$	$\delta = 1/2$
(C, S)	0.977 (0.011)	–	0.988 (0.013)	–	0.980 (0.006)	0.750 (0.217)
(C, F)	0.317 (0.063)	0.000	0.521 (0.085)	0.733 (0.077)	0.342 (0.0073)	0.255 (0.103)
(D, S)	0.150 (0.060)	–	0.263 (0.110)	–	0.143 (0.136)	0.000 ()
(D, F)	0.033 (0.013)	0.006 (0.003)	0.023 (0.009)	0.025 (0.009)	0.019 (0.019)	0.006 (0.004)

Data are taken from the last-five supergames in each treatment, with subject-clustered standard errors in parentheses. Cells marked "–" have no observations at the relevant history. History shows the own-action-signal pair from the previous round, (a_{t-1}, σ_{t-1}).

TABLE A.3. Cardinal comparisons to meta-study predictions

Treatment	Independent basin		Correlated basin		Cooperation	
	Initial	Ongoing	Initial	Ongoing	Initial	Ongoing
(N=2; X=\$9)	0.495	0.373	0.495	0.373	0.503	0.450
	($p = 0.886$)	($p = 0.164$)	($p = 0.886$)	($p = 0.164$)		
(N=4; X=\$9)	0.237	0.163	0.495	0.373	0.035	0.006
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)		
(N=4; X=\$1)	0.495	0.373	0.842	0.718	0.792	0.409
	($p < 0.001$)	($p = 0.472$)	($p = 0.229$)	($p < 0.001$)		
(N=10; X=\$1)	0.237	0.163	0.842	0.718	0.357	0.187
	($p = 0.031$)	($p = 0.656$)	($p < 0.001$)	($p < 0.001$)		
$F-$test, all	48.7	668.9	200.4	3,693.5		
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)		
$F-$test, all but (N=4; X=\$9)	18.4	0.89	26.2	53.8		
	($p < 0.001$)	($p = 0.447$)	($p < 0.001$)	($p < 0.001$)		

Results are calculated using data from the last five supergames using subject-clustered standard errors. In the first column, we present our four main treatments. The second and third columns display corresponding rates of initial and ongoing cooperation as predicted by the meta-study for the independent-basin measure. The fourth and fifth columns show cooperation rates predicted by the meta-study for the correlated basin measure. In the last column, we present observed cooperation rates in our main treatments (reproduced from Table 2). The p-values listed in the first four rows result from testing the null hypothesis of statistical equivalence between the observed and predicted cooperation rates. In the last two rows, we report F-statistics and p-values from testing the null hypothesis that observed and predicted cooperation rates across the treatments are statistically equal. The first F-test is conducted across all treatments, and the second is performed for all treatments except (N=4; X=\$9).

Table A.4. Ordinal pairwise treatment comparisons

Treatment pair	Basin		Cooperation	
	Independent	Correlated	Initial	Ongoing
$(N{=}2; X{=}\$9)$ vs. $(N{=}4; X{=}\$9)$	$\succ$	$\sim$	$\hat{\succ}$ $(p < 0.001)$	$\hat{\succ}$ $(p < 0.001)$
$(N{=}2; X{=}\$9)$ vs. $(N{=}4; X{=}\$1)$	$\sim$	$\prec$	$\hat{\prec}$ $(p < 0.001)$	$\hat{\sim}$ $(p = 0.583)$
$(N{=}2; X{=}\$9)$ vs. $(N{=}10; X{=}\$1)$	$\succ$	$\prec$	$\hat{\succ}$ $(p = 0.068)$	$\hat{\succ}$ $(p < 0.001)$
$(N{=}4; X{=}\$9)$ vs. $(N{=}4; X{=}\$1)$	$\prec$	$\prec$	$\hat{\prec}$ $(p < 0.001)$	$\hat{\prec}$ $(p < 0.001)$
$(N{=}4; X{=}\$9)$ vs. $(N{=}10; X{=}\$1)$	$\sim$	$\prec$	$\hat{\prec}$ $(p < 0.001)$	$\hat{\prec}$ $(p < 0.001)$
$(N{=}4; X{=}\$1)$ vs. $(N{=}10; X{=}\$1)$	$\succ$	$\sim$	$\hat{\succ}$ $(p < 0.001)$	$\hat{\succ}$ $(p = 0.001)$
# correct directional predictions (out of 4)				
Independent basin			4	4
Correlated basin			3	3
# correct null predictions (out of 2)				
Independent basin			0	1
Correlated basin			0	0

Results are calculated using data from the last five supergames. In the first column, we present six treatment comparisons. The next two columns indicate which of the two treatments in each pair has a higher predicted cooperation rate under the independent basin (column 2) and the correlated basin (column 3). The symbol $\succ$ signifies that the treatment listed first has a higher predicted cooperation rate, $\prec$ indicates that the treatment listed second has a higher predicted cooperation rate, and $\sim$ denotes that the two treatments have the same predicted cooperation rate. In the last two columns, we conduct tests to assess statistically significant differences in observed cooperation rates within each treatment pair, both for initial (column 4) and ongoing (column 5) cooperation, separately. All tests utilize subject-clustered standard errors and adhere to the 10 percent significance level. $\hat{\succ}$ indicates that the treatment listed first has empirically higher cooperation rate, $\hat{\prec}$ indicates that the treatment listed second has empirically higher cooperation rate, and $\hat{\sim}$ indicates that the cooperation rates observed in the two treatments are statistically indistinguishable. In the bottom half of the table, we present summary statistics for the total number of correct directional and null predictions under each basin extension.

Table A.5. Cooperation and success rates with implicit vs. explicit coordination

Action and signal rates	Implicit	Explicit	
	NoChat(3/4)	Chat(3/4)	Chat(1/2)
	Cooperation		
Initial	0.035 (0.017)	0.988 (0.007)	0.300 (0.037)
Ongoing	0.006 (0.003)	0.806 (0.030)	0.044 (0.018)
	Success		
Initial	0.000	0.971	0.094
Ongoing	0.000	0.757	0.002

Results are calculated using data from the last-five supergames, with subject-clustered standard errors in parentheses. Cooperation rates present raw proportions. All treatments have $X = \$9, N = 4$, where NoChat(3/4) refers to the core 2×2 between-subject design discussed in Section 4.

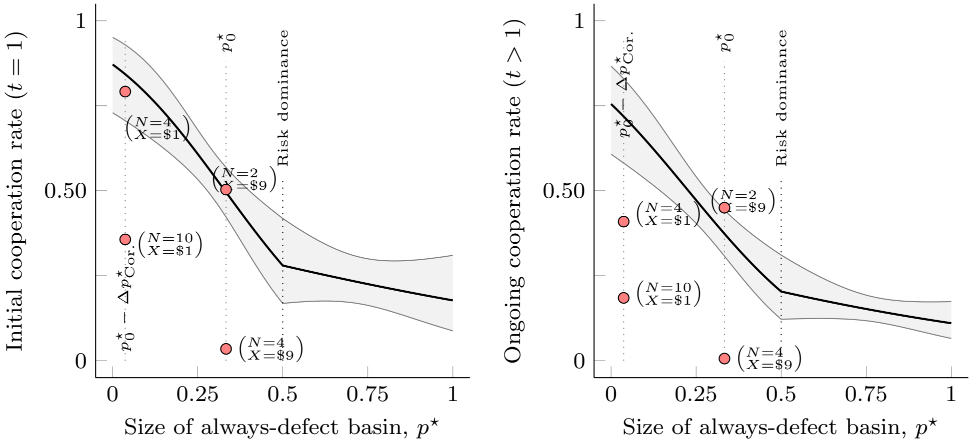

(A) Initial cooperation

(B) Ongoing cooperation

Figure A.1. Cooperation under the correlated basin-size model
Filled circles indicate separate treatments and filled diamonds treatments pooled over each value of the independent-basin measure. See Figure 2 in Section 4 for analogous results under the independent basin.

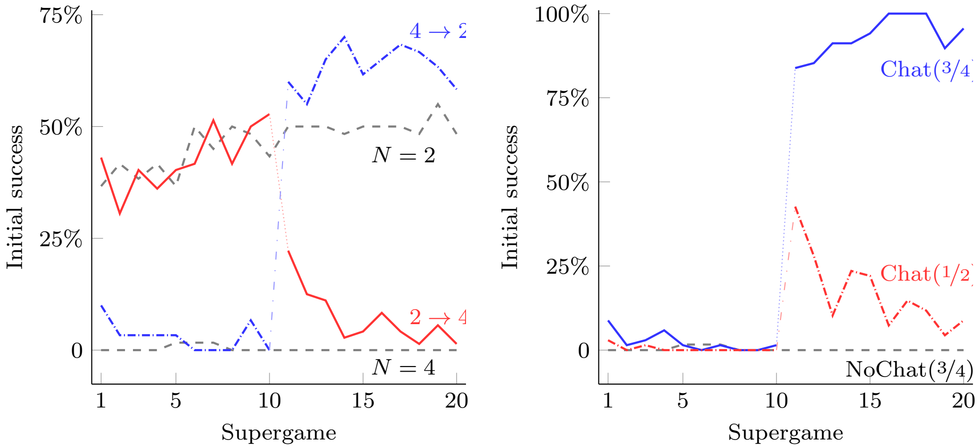

(A) Between vs. within ($X = \$9$)

(B) Explicit vs. implicit (N=4; X=\$9)

Figure A.2. Initial success rates in extensions (by supergame)

Appendix B: Further Analysis of the Within-Subject Treatments

In the within-subject treatments we find evidence of hysteresis. We observe a large and immediate jump in cooperation as N changes from $N = 4$ to $N = 2$, and no initial response as N moves in the opposite direction. In supergame 11 of our within-subject sessions (with prior experience at an alternate value of N) the initial cooperation rates at $N = 2$ and $N = 4$ are significantly greater than the initial cooperation rates in supergame one.[B.1] This suggests that in the short run, subjects respond to a change in the environment with a strong intent to cooperate after accumulating experience at another parameter value.

In the first four columns of Table B.1 we compare behavior of within- and between-treatment subjects after five rounds of experience holding the payoff cost of cooperating fixed at $X = \$9$. In the first two columns we present average behavior of subjects in supergames 6–10 for $N = 2$ and $N = 4$ (pooling as between/within treatments are identical to this point). In the next two columns we present average behavior of within-treatment subjects in supergames 16–20 in the $N = 4 \rightarrow 2$ and $N = 2 \rightarrow 4$ treatments. Examining the differences across the *within* and *between* subjects, we find: (i) No statistically significant differences for $N = 4$ and $N = 2 \rightarrow 4$ ($p = 0.117/p = 0.539$ for initial/ongoing cooperation), (ii) Statistically significant differences for $N = 2$ and $N = 4 \rightarrow 2$ ($p = 0.011$ for initial, $p < 0.001$ for ongoing). The significant differences reflect the substantially greater upward shift in the $4 \rightarrow 2$ treatment.

In the last three columns of Table B.1 we compare changes in behavior of within- and between-treatment subjects in response to a change in N. In column Δ_{Btwn} we present the change in average behavior of between-treatment subjects in supergames 16–20 as N increases from $N = 2$ to $N = 4$.[B.2] In the last two columns (jointly labeled as$\Delta_{\text{Wthn.}}$) we present the within-subject change in average behavior for the $2 \rightarrow 4$ and $4 \rightarrow 2$ treatments in supergames 6–10 and 16–20. While the three measures agree qualitatively—and exhibit economically large effects in N in the same direction—there are differences, particularly in the comparisons to the $2 \rightarrow 4$ case. However, we note that there are two effects at play here. In the $2 \rightarrow 4$ comparison, reduced magnitudes are driven primarily by the fact that behavior in this treatment has not converged. To see this, consider the assessed between-subject effect if we used data from supergames 6–10: a -33.5 percentage point effect on initial cooperation, which is not significantly different from the -26.0 percent effect identified in the within

B.1. Given the disjoint subject groups and identical treatment in supergames 1–10, we compare proportions using t-tests without clustering. We then compare the initial response under each value of N in the within-subject supergame eleven to all subjects at that N in supergame one. Using these tests, we reject equivalence with $p = 0.021$ for $N = 2$ and $p < 0.001$ for $N = 4$.

B.2. These results are analogous to the marginal effects attributable to an increase in the independent basin of $\Delta p^{\star}_{\text{Ind.}} = +0.36$ in Table 3 once we remove the $X = \$1$ treatments.

comparison ($p = 0.117$).[B.3] In contrast, the greater assessed effect in the $4 \to 2$ comparison is the composite of the same *reduction* in the effect from looking at the still-converging data for $N = 4$, with a substantial increase in cooperation at $N = 2$ in the second half over the between-subject levels.

TABLE B.1. Cooperation and success rates with between vs. within identification

Action and signal rates	Between (SG 6–10)		Within (SG 16–20)		$\Delta_{\text{Btwn.}}$	$\Delta_{\text{Wthn.}}$	
	$N = 2$	$N = 4$	$N = 2$	$N = 4$		$2 \to 4$	$4 \to 2$
			Cooperation				
Initial	0.474 (0.036)	0.139 (0.025)	0.643 (0.056)	0.214 (0.040)	−0.469 (0.061)	−0.261 (0.044)	−0.473 (0.059)
Ongoing	0.299 (0.026)	0.054 (0.012)	0.598 (0.051)	0.042 (0.016)	−0.444 (0.055)	−0.276 (0.036)	−0.536 (0.050)
			Success				
Initial	0.474	0.011	0.643	0.042	-0.503	-0.433	-0.630
Ongoing	0.299	0.004	0.598	0.008	-0.450	-0.310	-0.598

Comparisons at the same experience level are generated using supergames 6–10 across all sessions (fixing N, between and within sessions are identical until supergame 11). For the within change we measure the cooperation rates in supergames 16–20. All cooperation rates are raw proportions, with subject-clustered standard errors in parentheses. The last three columns measure the corresponding cooperation rate when $N = 4$ minus the cooperation rate when $N = 2$.

B.3. Similarly for ongoing cooperation the between-effect assessed in supergames 6–10 is -24.6 percent compared to -25.8 percent within ($p = 0.539$).

Appendix C: General Basin Calculation

Consider the N-player environment where we require cooperation from $M-1$ other players in order to receive a success signal. As such, if M players cooperate, then all N players will be guaranteed to receive a success signal. In the main body of the paper we consider a boundary case where $N = M$. In this appendix, we consider a generalized basin of attraction calculation where $M \leq N$.

Define $F(k)$ as the CDF over k, which is the number of other players choosing grim trigger. An agent will prefer grim trigger over always defect as long as:

$$(1 - F(M-2)) - F(M-2)\delta x \geq (1 - F(M-1))\,(1+x) - (F(M-1) - F(M-2))\delta(1+x),$$

which is equivalent to:

$$\Pr\left(\text{Exactly } M-1 \text{ choose } \alpha_{\text{Grim}}\right) \geq x\frac{(1-\delta)}{\delta} + x\Pr\left(\text{More than } M-1 \text{ choose } \alpha_{\text{Grim}}\right).$$

Given independent beliefs and probability p of others playing grim trigger, F is the CDF of a Binomial$(N-1,\, p)$. As such, we can simplify the preference for conditional cooperation to:

$$\binom{N-1}{M-1} p^{M-1}(1-p)^{N-M} \geq x\frac{(1-\delta)}{\delta}\sum_{k=0}^{M-1}\binom{N-1}{k} p^{k}(1-p)^{N-k-1}.$$

For $M = N$ cooperation is preferred for $p \in [p^\star, 1]$.

For $1 < M < N$ cooperation is preferred for $p \in \left[\underline{p}^\star, \overline{p}^\star\right]$, where $0 < \underline{p}^\star \leq \overline{p}^\star < 1$. This is the case because cooperation is never a best response if no one else cooperates (impossible to get a success, as $M \geq 2$) or if everyone else cooperates (success regardless of own action, $M < N$).

In Figure C.1 we outline the values of p and x for which cooperation is preferred at $\delta = 3/4$. In Panel (A) we do so for $N = M \in \{2, 4, 10\}$, and in Panel (B) for $N = 4$ and $M \in \{2, 3, 4\}$. The figure illustrates that for $M < N$ the basin size is given by $\underline{p}^\star + 1 - \overline{p}^\star$ instead of $p^\star$. Moreover, the figure illustrates that the basin size is full for our extension treatment with $M = 2$ and $X = \$9$.

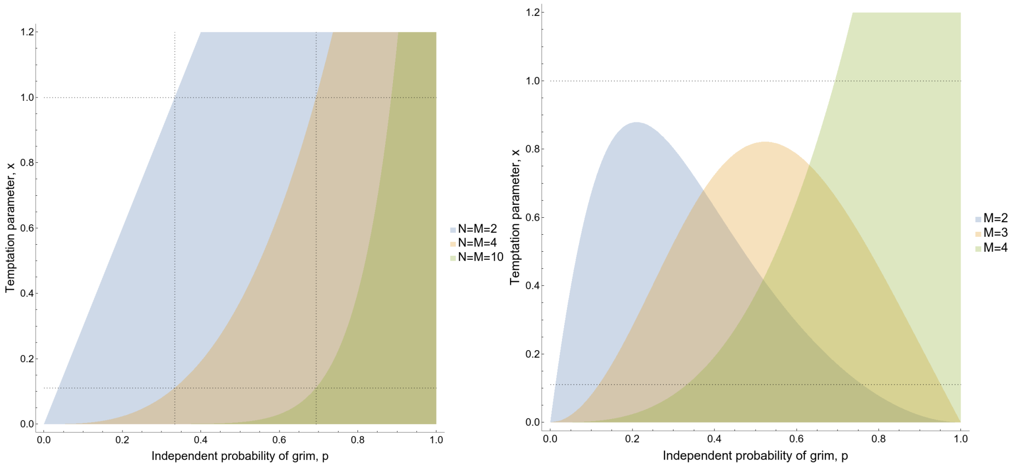

(A) $N = M$ (B) $M < N = 4$

FIGURE C.1. Grim best response to independent belief p and temptation x

Appendix D: Estimation of σ

Our measure of how much belief correlation is necessary to rationalize the data uses a convex combination of the independent and correlated models. Specifically, the probability that the other $N-1$ players coordinate is given by

$$\sigma \cdot p + (1-\sigma) \cdot p^{N-1},$$

with the critical belief denoted by $p^{\star}(\sigma, x, N)$. The additional parameter σ nests the two extremes: $\sigma = 0$ for full independence, $\sigma = 1$ for perfect correlation.

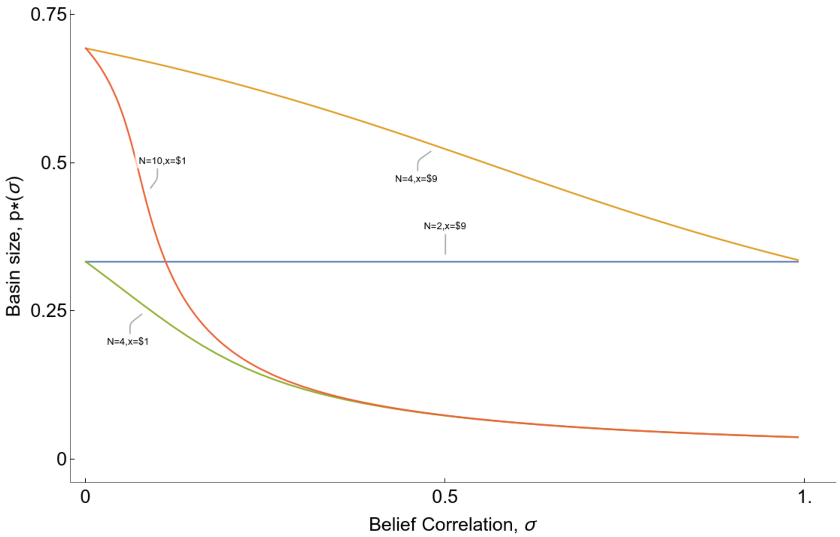

Figure D.1. Belief correlation and basin size

The figure displays analytic solutions for $p^{\star}(\sigma)$ for any value of $\sigma \in [0,1]$

To provide some insight into the estimator of σ, consider Figure D.1. The figure outlines the analytic solution for $p^{\star}(\sigma)$, the basin of attraction in each treatment, as a function of σ. For instance, when $\sigma = 0$ ($\sigma = 1$) the value of $p^{\star}$ for each treatment coincide with the predictions for the independent (correlated) extension discussed in Table 1 Panel B. Between the two extremes, the figure displays the predicated intermediate values of $p^{\star}$ for each σ within the range of $(0,1)$.

For each $p^{\star}$, we can use the probit estimates from the meta-study, and obtain predicted cooperation, $\hat{\mathcal{C}}_{\text{Meta}}(p^{\star})$; see Section 2 for details. Figure D.2(A) shows the predictions for the case of initial cooperation and Figure D.2(B) for ongoing cooperation. In addition, notice that both figures display in dashed lines the observed cooperation rates in our treatments.

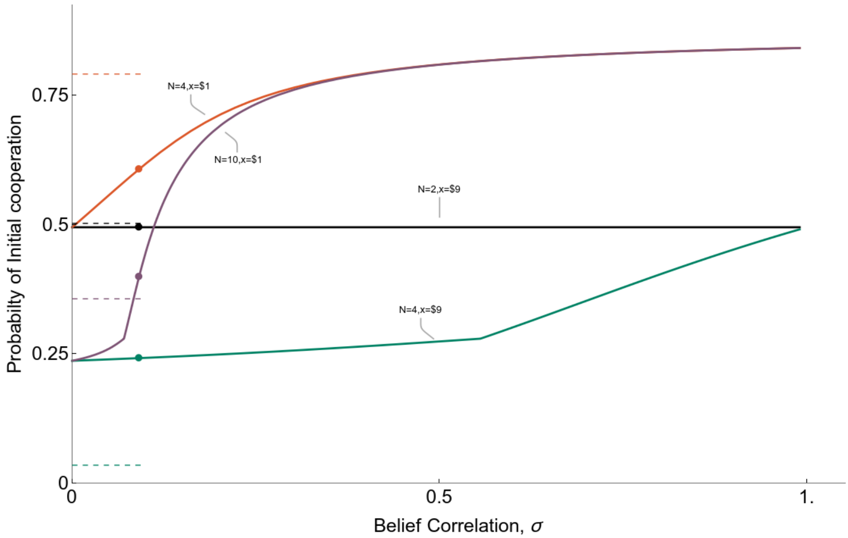

(A) Initial cooperation

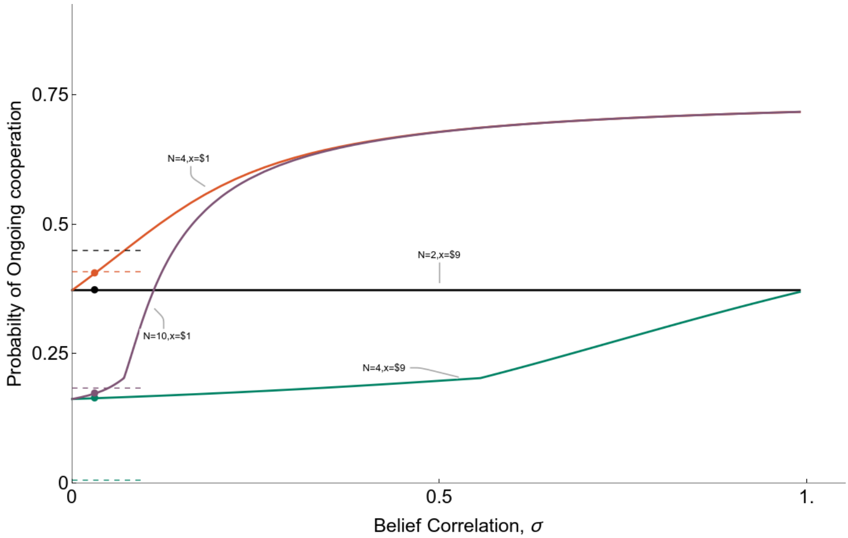

(B) Ongoing cooperation

Figure D.2. Belief correlation and cooperation

Using the estimated initial (ongoing) cooperation rates from the meta study as a function of basin-size $\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\right)$ we can therefore indicate the expected cooperation level for any value of σ, which we illustrate in panel A (B). The dashed lines (with color matching the corresponding treatment) mark the observed cooperation rates in our treatments. The dots indicate the corresponding value of σ that maximizes the likelihood function.

With this background we can construct a log-likelihood equation across our four treatments:

$$l(\sigma) = l\left(\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\left(\sigma,1,2\right)\right)\right) + l\left(\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\left(\sigma,1,4\right)\right)\right)$$
$$+\, l\left(\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\left(\sigma,1/9,4\right)\right)\right) + l\left(\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\left(\sigma,1/9,10\right)\right)\right).$$

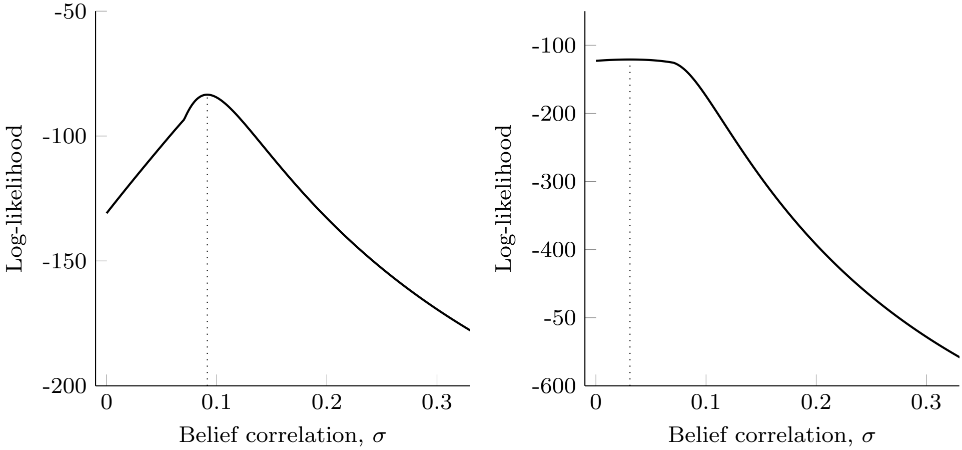

(A) Initial cooperation (B) Ongoing cooperation

FIGURE D.3. Belief correlation

Data are taken from the last five supergames in each between-subject treatment. Log-likelihoods are calculated using the imputed cooperation rate from the Dal Bó and Fréchette (2018) meta-study with belief correlation σ calculated as σ proportion of independent beliefs and $(1-\sigma)$ proportion of perfectly correlated beliefs.

This is a single equation in σ, where the likelihood of our data from each treatment is measured as a binomial under a cooperation probability of $\hat{\mathcal{C}}_{\text{Meta}}\left(p^{\star}\left(\sigma, x, N\right)\right)$.[D.1] Finally, we estimate σ via maximum likelihood.

Figure D.3 shows the log-likelihood as a function of σ for initial and ongoing cooperation.

D.1. Notice that the expression uses the normalized value x, so that $x = 1, N = 2$ corresponds to the (N=2; X=\$9) treatment.

Appendix E: Interface Screenshots

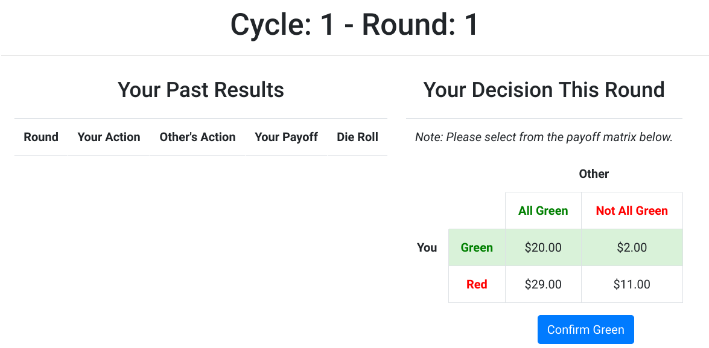

(A) Action selection

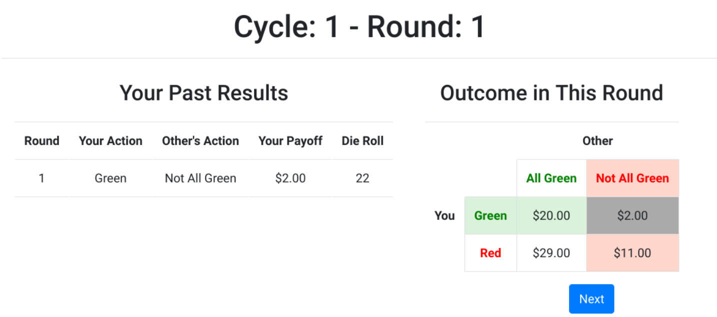

(B) Round feedback

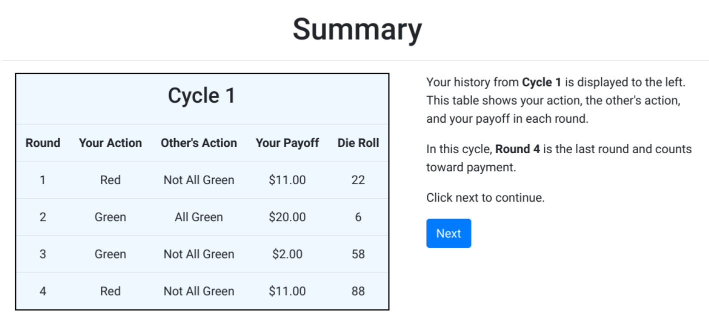

(C) Supergame feedback

Figure E.1. Interface screenshots

Appendix F: Provided Instructions

Below, we include the instructions provided to participants. All language deltas/treatment-specific language is enclosed in braces. Text in red pertains to the $N = 2$ treatment, while text in blue pertains to the $N > 2$ treatments (here we provide the $N = 4$ implementation, where $N = 10$ has only minor changes). Payoff text for $X = 9$ is presented in green, and for $X = 1$ in orange. Separate instructions for {Part two} are provided to treatments where N changes within a session. In the Chat(1/2) treatment, the only changes are for the critical die rolls in the *Study Organization & Payment* section, where the supergame cutoff changes from 75 to 50. In the extension treatment in which only two out of four players are needed for a success signal, adjustments are made to the description in the *Round Choices and Payoffs* to accommodate this change.

Instructions

Welcome

You are about to participate in a study on decision-making. What you earn depends on your decisions, and the decisions of others in this room. Please turn off your cell phones and any similar devices now. Please do not talk or in any way try to communicate with other participants. We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the study. If you have any questions during this period, raise your hand and your question will be answered in private at your computer carrel.

Study Organization & Payment

- The study has two Parts, where each Part has 10 decision-making **Cycles**. Each Cycle consists of a random number of **Rounds** where you make decisions.
- At the end of the study, one of the two Parts will be selected for payment with equal probability. For the selected Part, one of the 10 Cycles will be randomly selected for payment. Your payment for this randomly selected Cycle will be based on your decision's in that Cycle's last Round.
- The number of Rounds in each Cycle is random, where only the last Round in each Cycle counts for payment. Which Round is the last is determined as follows:
  - In every Round, after participants make their decisions, the computer will roll a fair 100-sided die. If the die roll is greater than 75 (so 76–100) the round just completed is the one that is used to determine the current Cycle's payment, and the Cycle ends. If instead the computer's roll is less than 75 (so 1– 75) then the Cycle continues into another Round.
  - Because of this rule, after every Round decision there is a 25 percent chance that the current Round is the ones that count for the Cycle's payment, and a 75 percent chance that the Cycle continues and the decisions in a subsequent round will count for that Cycle payment.
- Your final payment for the study will be made up of a $6 show-up fee, and your payment from the last Round in the randomly selected Cycle.

Part 1

- In the first part of the study you will make decisions in 10 Cycles. In each Cycle you will be matched with {another participant}{a group of three other participants} in the room for a sequence of Rounds. You will interact with the same {other participant}{group of three other participants} in all rounds of the cycle.

- Once a Cycle is completed, you will be randomly matched to a new {participant}{group of three participants} for the next Cycle.
- While the specific {participant}{participants} you are matched to is fixed across all Rounds in the Cycle, the computer interface in which you make your decisions is anonymous, so you will never find out which participants in the room you interacted with in a particular Cycle, nor will others be able to find out that they interacted with you.

Round Choices and Payoffs

For each Round in each Cycle, you and the matched {participant}{participants in your group} will make simultaneous choices. {Both}{All four} of you must choose between either the **Green** action or the **Red** action. After you and the other {participant}{three participants} have made your choices, you will be given feedback on the {other participant's}{other participants'} choices that Round, alongside the Computer's die roll to determine if that Round counts for the Cycle payment.

If a particular Round is the Cycle's last, and that Cycle is the one selected for final payment, there are four possible payoff outcomes.

1. If both you and {the other participant}{all three of the other participants} choose the Green action, you get a round payoff of $20.
2. If you choose the Green action and {the other participant chooses}{any of the other participants choose} Red, you get a round payoff of {$2}{$10}.
3. If you choose the Red action and {the other participant chooses}{all of the three other participants choose} Green, you get a round payoff of {$29}{$21}.
4. If both you and {the other participant}{any of the other three participants} choose the Red action, you get a round payoff of $11.

These four payoffs are summarized in the following table:

		Other {Participant's Action:}{Participants' Actions:}	
		{Green}{All 3 Green}	{Red}{Any of 3 Red}
Your Action:	Green	$20	{$2}{$10}
	Red	{$29}{$21}	$11

Some examples of these payoffs:

Case 1. Suppose you choose Green and {the other participant}{all three of the other participants} in the Cycle also choose Green. If that Round is the final one in the Cycle {both}{all four} of you would get a payoff of $20.

Case 2. Suppose {you}{you and two of the other participants} choose Green while the other participant chooses Red. If that Round is the final one in the Cycle {you}{you and the other two participants who chose Green} would get a payoff of {$2}{$10}, while the other participant would get a payoff of {$29}{$21}.

Case 3. Suppose you chooseRed while {the other participant chooses}{all three of the other participants choose} Green. If that Round is the final one in the Cycle you would get a payoff of {$29}{$21}, while the other {participant}{three participants} would get a payoff of {$2}{$10}.

Case 4. Suppose you and {the other participant choose Red.}{another participant choose Red while the other two participants choose Green.} If that Round is the final one in the Cycle {you}{you and the other participant that chose Red} would get a payoff of {$11}{$11, while the other two participants would get a payoff of{$2}{$10}}
.

Part 2

After Part 1 is concluded, you will be given instructions on Part 2, which will have a very similar structure to the task in Part 1.

{END OF PART 1 HANDOUT}

Part 2 Instructions {Between Only, handed out Supergame 11}

Part 2 is identical to Part 1. In each of the 10 Cycles in Part 2 you will again be matched to {another participant}{three other participants} in the room.

Similar to Part 1, the Cycle payoff is determined by the last round in the Cycle, where the payoff depends on the action you chose and the {action chosen by the matched participant}{actions chosen by the three matched participants} for that Cycle. Similar to Part 1, the below Table summarizes the payoff based upon the choices made in the Cycle's last round.

		Other {Participant's Action:}{Participants' Actions:}	
		{Green}{All 3 Green}	{Red}{Any of 3 Red}
Your Action:	Green	\$20	{\$2}{\$10}
	Red	{\$29}{\$21}	\$11

{END OF PART 2 HANDOUT}

Part 2 Instructions {Within Only, handed out Supergame 11}

Part 2 is very similar to Part 1. However, in each of the 10 Cycles in Part 2 you will instead be matched to three other participants in the room for each Cycle.

Similar to Part 1, the Cycle payoff is determined by the last round in the Cycle, where the payoff depends on the action you chose and the actions chosen by the three matched participants for that Cycle. If a particular Round is the Cycle's last, and that Cycle is the one selected for final payment, there are four possible payoff outcomes.

1. If both you and all three of the other participants choose the Green action, you get a round payoff of $20.
2. If you choose the Green action and any of the other participants chooses Red, you get a round payoff of $2.
3. If you choose the Red action and all three other participants choose Green, you get a round payoff of $29.
4. If both you and any of the other three participants choose the Red action, you get a round payoff of $11.

These four payoffs are summarized in the following table:

		Other Participant's Action:	
		All 3 Green	Any of 3 Red
Your Action:	Green	$20	$2
	Red	$29	$11

Some examples of these payoffs:

Case 1. Suppose you choose Green and all three of the other participants in the Cycle also choose Green. If that Round is the final one in the Cycle all four of you would get a payoff of $20.

Case 2. Suppose you and two of the other participants choose Green while the other participant chooses Red. If that Round is the final one in the Cycle you and the other two participants who chose Green would get a payoff of $2, while the other participant would get a payoff of $29.

Case 3. Suppose you choose Red while all three of the other participants choose Green. If that Round is the final one in the Cycle you would get a payoff of $29, while the other three participants would get a Round payoff of $2.

Case 4. Suppose you and another participant choose Red while the other two participants choose Green. If that Round is the final one in the Cycle you and the other participant that chose Red would get a payoff of $11, while the other two participants would get a payoff of $2.

{END OF PART 2 HANDOUT}

Part 2 Instructions {Chat Only, handed out Supergame 11}

Part 2 is identical to Part 1 except for the beginning of each cycle where we will now allow the matched participants to chat to one another before the cycle begins. In each of the 10 Cycles in Part 2 you will again be matched to three other participants in the room.

Similar to Part 1, the Cycle payoff is determined by the last round in the Cycle, where the payoff depends on the action you chose and the actions chosen by the three matched participants for that Cycle. Similar to Part 1, the below Table summarizes the payoff based upon the choices made in the Cycle's last round.

		Other Participants' Actions:	
		All 3 Green	Any of 3 Red
Your Action:	Green	$20	$2
	Red	$29	$11

In contrast to Part 1 though, at the beginning of each new cycle, a chat window will be given to you, which will stay open for two minutes, or until all group members close it.

You may not use the chat to discuss details about your previous earnings, nor are you to provide any details that may help other participants in this room identify you. This is important to the validity of this study and will be not tolerated. However, you are encouraged to use the chat window to discuss the upcoming Cycle.

If at any point within the two-minute limit you wish to leave the chat, you can click the "Finish Chat" button. The other participants will be informed that you left.

{END OF PART 2 HANDOUT}

Instructions{Two-from-four, difference shown to Four-from-Four}

Welcome

You are about to participate in a study on decision-making. What you earn depends on your decisions, and the decisions of others in this room. Please turn off your cell phones and any similar devices now. Please do not talk or in any way try to communicate with other participants. We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the study. If you have any questions during this period, raise your hand and your question will be answered in private at your computer carrel.

Study Organization & Payment

- The study has two Parts, where each Part has 10 decision-making **Cycles**. Each Cycle consists of a random number of **Rounds** where you make decisions.
- At the end of the study, one of the two Parts will be selected for payment with equal probability. For the selected Part, one of the 10 Cycles will be randomly selected for payment. Your payment for this randomly selected Cycle will be based on your decision's in that Cycle's last Round.
- The number of Rounds in each Cycle is random, where only the last Round in each Cycle counts for payment. Which Round is the last is determined as follows:
  - In every Round, after participants make their decisions, the computer will roll a fair 100-sided die. If the die roll is greater than 75 (so 76–100) the round just completed is the one that is used to determine the current Cycle's payment, and the Cycle ends. If instead the computer's roll is less than 75 (so 1– 75) then the Cycle continues into another Round.
  - Because of this rule, after every Round decision there is a 25 percent chance that the current Round is the ones that count for the Cycle's payment, and a 75 percent chance that the Cycle continues and the decisions in a subsequent round will count for that Cycle payment.
- Your final payment for the study will be made up of a $6 show-up fee, and your payment from the last Round in the randomly selected Cycle.

Part 1

- In the first part of the study you will make decisions in 10 Cycles. In each Cycle you will be matched with a group of three other participants in the room for a sequence of Rounds. You will interact with the same group of three other participants in all rounds of the cycle.

- Once a Cycle is completed, you will be randomly matched to a new group of three participants for the next Cycle.
- While the specific participants you are matched to is fixed across all Rounds in the Cycle, the computer interface in which you make your decisions is anonymous, so you will never find out which participants in the room you interacted with in a particular Cycle, nor will others be able to find out that they interacted with you.

Round Choices and Payoffs

For each Round in each Cycle, you and the matched participants in your group will make simultaneous choices. All four of you must choose between either the **Green** action or the **Red** action. After you and the other three participants have made your choices, you will be given feedback on the other participants' choices that Round, alongside the Computer's die roll to determine if that Round counts for the Cycle payment.

If a particular Round is the Cycle's last, and that Cycle is the one selected for final payment, there are four possible payoff outcomes.

1. If both you and {all three}{any} of the other participants choose the Green action, you get a round payoff of $20.
2. If you choose the Green action and {any}{all} of the other participants choose Red, you get a round payoff of $2.
3. If you choose the Red action and {all}{any} of the three other participants choose Green, you get a round payoff of $29.
4. If both you and {any}{all} of the other three participants choose the Red action, you get a round payoff of $11.

These four payoffs are summarized in the following table:

		Other Participants' Actions:	
		{All}{Any of} 3 Green	{Any of}{All} 3 Red
Your Action:	Green	$20	$2
	Red	$29	$11

Some examples of these payoffs:

Case 1. Suppose you choose Green and all three of the other participants in the Cycle also choose Green. If that Round is the final one in the Cycle all four of you would get a payoff of $20.

Case 2. Suppose you and two of the other participants choose Green while the other participant chooses Red. If that Round is the final one in the Cycle you and the other two participants who chose Green would get a payoff of {$2}{$2}, while the other participant would get a payoff of $29.

Case 3. Suppose you choose{Red}{Green} while all three of the other participants choose {Green}{Red}. If that Round is the final one in the Cycle you would get a payoff of {\$29}{\$2}, while the other three participants would get a payoff of {\$2}{\$29}.

Case 4. Suppose you and{another participant choose Red while the other two participants choose Green. }{the other participants choose Red. }If that Round is the final one in the Cycle{you and the other participant that chose Red would get a payoff of \$11, while the other two participants would get a payoff of\$2 .}{you and the other participants would get a payoff of \$11.}

Appendix G: Strategies and the selection index

An RPD is characterized by a large number of possible strategies. However, the meta-study of RPD lab experiments of Dal Bó and Fréchette (2018) shows that a small set of strategies rationalizes choices well for a large number of parameterizations. The five strategies that capture most choices are: (i) always cooperate, (ii) always defect; and three strategies in which cooperation is conditional, (iii) grim trigger, (iv) tit for tat, and (v) suspicious tit for tat. The difference between tit for tat and suspicious tit for tat is limited to the first interaction, where tit for tat starts with cooperation and suspicious tit for that starts with defection. In all subsequent rounds, both players cooperate as long as their opponents cooperate and defect otherwise.

In this paper, and more broadly in the RPD literature, the selection index focuses on two strategies: always defect and grim trigger. A first reason to focus on these two strategies is that they capture very distinct types of behavior (non-cooperative and conditionally cooperative) that both may be supported in equilibrium. Always defect is the only non-cooperative strategy that is subgame perfect, and grim trigger is the only conditionally cooperative *and* empirically-relevant strategy that depending on δ can be subgame perfect. Note that for tit for tat, which is a Nash equilibrium of the supergame, there can be incentives to deviate from the punishment path. In addition, if one player chooses always defect and the other plays tit for tat, the outcome is cooperation in the first round and defection from the second round on. This is the same outcome that will realize if the other player chooses grim trigger instead. Similarly, if both players choose tit for tat, the outcome is the same (cooperation in every round) as when the two players use grim trigger instead. Therefore, one can begin to argue that there is little loss in focusing solely on grim trigger as the conditionally cooperative strategy.

In this appendix we show why focusing on always defect and grim trigger also extends to our setting. Here, we start with a brief description of the Strategy Frequency Estimation Method (SFEM), which was introduced in Dal Bó and Fréchette (2011).[G.1] From a big-picture perspective, the method takes choices made by subjects and contrasts them against hypothetical choices that would have been made were the subjects using a different strategy from a pre-determined strategy set. Using a mixture model that allows for errors in choices, the procedure reports the proportion of choices that are better rationalized by each strategy. Dal Bó and Fréchette (2019) also use an alternative procedure to study strategies: an experimental design that familiarizes subjects with a set of strategies and asks them to select one

G.1. Further details on the procedure are available in the online appendix of Embrey, Fréchette, and Stacchetti (2013). A Monte-Carlo-style analysis was also performed in Fudenberg, Rand, and Dreber (2012). The procedure has also been used to study strategies in other repeated-game experiments, for example, Aoyagi, Bhaskar, and Fréchette (2019), Vespa (2020), and Vespa and Wilson (2020).

to be played in their name. The authors contrast this elicitation procedure with the SFEM and find consistency across the two methods.

Strategy Frequency Estimation Method

The goal of the procedure is to recover ϕ_k, which represents the frequency attributed to strategy k in the data. To illustrate how the procedure works, consider a set of strategies $\mathcal{K}$. Let $d^i_{gr}(\mathbf{h})$ be the choice of subject i and $k^i_{gr}(\mathbf{h})$ the decision prescribed for subject i by strategy $k \in \mathcal{K}$ in round r of supergame g for a given history $\mathbf{h}$. Strategy k is a perfect fit for round r if $d^i_{gr}(\mathbf{h}) = k^i_{gr}(\mathbf{h})$. The procedure models the probability that the choice (d) corresponds to the prescribed decision (k) as:

$$\Pr\left(d^i_{gr}(\mathbf{h}) = k^i_{gr}(\mathbf{h})\right) = {}^1\!/_{1+\exp\left(\frac{-1}{\gamma}\right)} = \beta, \tag{G.1}$$

where β captures the probability that the subject does not make make mental errors in the implementation of a strategy and $\gamma > 0$ is a parameter of interest. In the limit, as $\gamma \to 0$ and $\beta \to 1$ the model fully rationalizes the data. On the other hand, as $\gamma \to \infty$, $\beta \to \frac{1}{2}$, the model has no explanatory power.

Now, let y^i_{gr} be an indicator that takes value one if the subject's choice matches the decision prescribed by the strategy. It follows from Equation (G.1) that the likelihood of observing strategy k for subject i is given by:

$$p_i(k) = \prod_g \prod_r \beta^{y^i_{gr}} (1-\beta)^{1-y^i_{gr}}. \tag{G.2}$$

Aggregating over subjects we arrive at the log-likelihood of the following form: $\sum_i \ln\left(\sum_k \phi_k p_i(k)\right)$.[G.2,G.3]

For illustration, consider a case in which the set of strategies includes always defect and always cooperate. The fit of the model will be good (high β) if the population is composed of subjects who either almost always defect or almost always cooperate. If a large proportion of subjects shifts between cooperation and defection within a supergame, neither strategy will accommodate their choices and the estimation will return a low estimate of β.

The SFEM depends on the pre-specified set of strategies $\mathcal{K}$. The information that subjects receive at the end of each round in our environments with $N > 2$ is similar

G.2. To construct $p_i(k)$, consider a subject who is implementing the prescriptions of strategy k with mistake rate given by $1-\beta$. If $y^i_{gr} = 1$, the subject's choice matches the prescription; if $y^i_{gp} = 0$, the subject's choice does not coincide with the one prescribed by the strategy.

G.3. Since $\sum_k \phi_k = 1$, the procedure provides $|\mathcal{K}| - 1$ estimates and the $|\mathcal{K}|$-th strategy is computed by difference. The procedure also estimates γ. Following Equation (G.1) there is a one-to-one mapping between γ and β, so we will refer to the estimate of γ directly as an estimate of β.

to the information that subjects receive in a two-player RPD. The reason is that in our multi-player game subjects do not learn the specific choices of others, but instead, receive an aggregate signal of either a success or a failure. Therefore, to study behavior of multiple players we focus on the same set of five strategies identified in Dal Bó and Fréchette (2011) for a two-player game. In fact, as we will show later, these five strategies suffice to obtain relatively high goodness of fit estimates (as captured by β).

Results

In this section, we will present results obtained using the SFEM. First, we will show estimates for the last seven supergames of the session and then, we will provide results for the first seven supergames of the session.[G.4]

Final Supergames. In Panel (A) of Table G.1 we present estimation results for each of our nine treatments, including the estimates of β. For each treatment and each of the five strategies identified as focal in Dal Bó and Fréchette (2018) the table reports the estimates and (whenever possible) the bootstrap-estimated standard errors.[G.5,G.6]

In Panel (B) of Table G.1 we report goodness-of-fit estimates obtained after restricting the set of available strategies. In particular, $\beta^{\dagger}$ corresponds to the β estimate after eliminating tit for tat and suspicious tit for tat from the strategy set. In this case, the only conditional-cooperation strategy remaining in the set is grim trigger. Since the model uses maximum likelihood and the restricted and unrestricted models are nested, we use a likelihood-ratio test with a $p-$value referenced with † to evaluate the null hypothesis that the restriction does not bind. In the last two rows in Table G.1 Panel (B) we report results of an analogous exercise but where the only two strategies included in the strategy set are always cooperate and always defect. The corresponding β estimate and $p-$value are marked with ‡. In what follows, we discuss the estimation results across different strategy sets treatment by treatment.

In the (N=2; X=\$9) treatment, two strategies with the most mass are always defect (45.2 percent) and grim trigger (28.9 percent). However, there is a non-negligible yet not-significant mass captured by tit for tat (18.0 percent). In the estimation that excludes tit for tat and suspicious tit for tat, there is a small reduction in terms of goodness of fit: a drop from 0.929 to 0.911. On the one hand, the reduction in

G.4. The results that we report qualitatively do not depend on having seven supergames among the early and late samples. The focus on seven is intended for two reasons. First, it allows for six supergames in between, so that it is possible to see if behavior early on changes relative to behavior much later in the session. Second, there is enough data in each seven-supergame sample.

G.5. Recall that the procedure recovers standard errors for all the strategies but one. (See footnote G.3 for details).

G.6. Observations for the chat treatment with $\delta = 1/2$ are lower than in other treatments because in this case with a higher termination probability after each round, supergames are shorter.

goodness of fit seems marginal. On the other hand, the likelihood ratio test suggests that excluding tit for tat and suspicious tit for tat from the strategy set leads to a statistically significant loss. To put this result into perspective, consider a SFEM estimation that only accounts for always cooperate and always defect. Here, we observe a relatively large decrease in the goodness-of-fit measure: a drop from 0.929 to 0.804. This suggests that the noise component needed to rationalize the data without grim trigger is substantially larger than with grim trigger in the strategy set.

The (N=4; X=\$9) treatment is one with the least amount of cooperation, as confirmed by our estimation results. About a third of the mass corresponds to always defect and about two thirds to suspicious tit for tat. Here, the goodness-of-fit measure is close to one (at 0.981). Moreover, there is no evidence of a loss in carrying out the estimation without tit for tat and suspicious tit for tat, as shown by the $\beta^{\dagger}$ estimates and the likelihood ratio test statistics.

The treatment with the highest degree of initial cooperation in our data is (N=4; X=\$1). Here, about a little less than a quarter of the mass corresponds to always cooperate. While this suggests a relatively large amount of unconditional cooperation, whether a player cooperates unconditionally cannot be determined unless their opponent defects.[G.7] However, the broader evidence suggests that cooperation is conditional. In treatments with more frequent defection there is essentially no evidence of large amounts of unconditional cooperation. In fact, in the (N=4; X=\$1) treatment the most popular strategy is tit for tat, which captures 53.5 percent of the mass. There is also a close to 20.0 percent of the mass that corresponds to always defect. We note that while the likelihood ratio test rejects the null, the loss in terms of goodness of fit is rather small: a drop from 0.939 to 0.931. A reduction in goodness of fit is much larger in the absence of grim trigger: a drop from 0.939 to 0.840.

In our last core treatment, (N=10; X=\$1), all strategies for which standard errors can be computed have statistically significant estimates. About a quarter of the mass corresponds to always defect and a tenth to always cooperate. Grim trigger, tit for tat, and suspicious tit for tat, jointly, account for close to 60.0 percent of the mass. However, when the estimation excludes tit for tat and suspicious tit for tat, the goodness-of-fit estimate does not change up to the third decimal. Consistently, the likelihood ratio test rejects the null at any typical significance level. However, the

G.7. While a strategy in an infinitely repeated game specifies what to do at each possible decision node (an infinite-dimensional object), the observed set of choices for a subject corresponds to a specific path of play. To increase possible identification, Vespa (2020) uses a one-period-ahead strategy method (OASM) in which subjects make choices in round r without knowing others' choices in round $r-1$. That is, the subject makes a choice in round r for each possible choice of the other player in round $r-1$. After making these choices, the subjects learn the actual history of play for round $r-1$, and their choices for round r are implemented. In this way, it is possible to retrieve in an incentivized manner choices that subjects would have made off the path of play. Implementing OASM is costly because it reduces the number of supergames that subjects can reasonably play within a session given that they must make more decisions per round. Since the goal of the current paper does not lie in identifying strategies, we decided not to include it in our design.

goodness-of-fit measure decreases from 0.934 to 0.897 without grim trigger in the strategy set.

Overall, we find that:

Result G.1. *Focusing on two strategies, always defect and grim trigger, when testing the extensions of the basin does not lead to a substantial loss. This is the case either because a likelihood ratio test directly points towards the restriction not binding or because when it binds the relative loss is small (as measured by the goodness-of-fit estimates).*

Result G.2. *Further restricting the strategy set to exclude grim trigger, in most cases, leads to a relatively large loss in goodness of fit.*

We now discuss the estimates for our extension treatments. In the $2 \to 4$ treatment the last seven supergames are played with $N = 4$, so that the results of this extension are comparable to those obtained in our core (N=4; X=\$9) treatment. However, the former estimates are noisier (relative to the latter). This is likely due to the fact that in the last seven supergames of the $2 \to 4$ treatment subjects are less experienced relative to their counterparts in (N=4; X=\$9).[G.8] However, the big picture is similar: (i) There is a very small reduction in the goodness-of-fit estimate when the set of strategies excludes tit for tat and suspicious tit for tat (a drop from 0.950 to 0.948);[G.9] (ii) There is a non-negligible loss in the goodness-of-fit measure after excluding grim trigger from the strategy set (a drop from 0.950 to 0.895); (iii) Most subjects appear to use strategies that are captured either by directly not cooperating (always defect) or by starting in a non-cooperative manner (suspicious tit for tat). Meanwhile, in the more competitive $4 \to 2$ treatment excluding both tit for tat and suspicious tit for that leads to a larger loss in goodness of fit relative to other treatments (a drop from 0.921 to 0.883). However, were we to additionally exclude grim trigger, it would lead to a much larger loss, with $\beta^{\ddagger}$ at 0.819.

In a treatment with chat and $\delta = 3/4$, grim trigger and tit for tat essentially capture almost all the mass. Excluding tit for that and suspicious tit for tat leads to a statistically significant yet negligible reduction in goodness of fit. A reduction in goodness of fit is much larger when grim trigger is also excluded: a drop from 0.975 to 0.883. On the contrary, in a chat treatment with $\delta = 1/2$, always defect captures more than 60.0 percent of the mass and suspicious tit for tat accounts for 11.0 percent of the mass. Here, excluding tit for that and suspicious tit for tat has virtually no effect

G.8. In fact, the estimates for $2 \to 4$ in Table G.1 are closer to the estimates for (N=4; X=\$9) using the *first* seven supergames, which are reported in Table G.2. In both cases, the largest mass corresponds to always defect (around 60 percent) and suspicious tit for tat.

G.9. The likelihood ratio test also leads to the same result using a 95 percent confidence level.

on the goodness of fit (a drop from 0.873 to 0.871). However, were we to additionally exclude grim trigger, it would lead to an economically and statistically significant loss, with $\beta^{\ddagger}$ at 0.809.[G.10]

In our final extension treatment always defect captures close to three-quarters of the mass. This evidence is consistent with what we have observed earlier that cooperation in our treatment with weakened cooperation requirement is quite unlikely despite the fact that the number of players needed for a cooperative outcome is smaller than N. Furthermore, we observe a small reduction in the goodness-of-fit estimate when the set of strategies excludes tit for tat and suspicious tit for tat (a drop from 0.899 to 0.896). The reduction in goodness of fit becomes much larger once we also exclude grim trigger (a drop from 0.899 to 0.863).

Early Supergames. We conclude this section of the appendix by describing SFEM estimates for the first seven supergames. We first note that the goodness-of-fit estimates (β) are still quite far from random, with the smallest estimate at 0.811. This suggests that even if constrained to few strategies, most of the data can be rationalized. However, in all but one treatment there are large reductions in goodness-of-fit estimates when compared to the last seven supergames (see Table G.1). This suggests that as subjects gather experience, their behavior becomes more consistently captured by the five focal strategies identified in Dal Bó and Fréchette (2018).

Second, comparing between the first- and the last seven supergames we observe differences in strategies that best capture the observed behavior. For instance, in (N=4; X=\$9) the largest mass in the first seven supergames corresponds to always defect (60.0 percent) and the second largest to suspicious tit for tat (33.8 percent). In the last seven supergames, the order is reversed, with 67.0 percent for suspicious tit for tat and 31.3 percent for always defect. However, in both cases, the two strategies jointly capture more than 90.0 percent of the mass. That is, the odds of a cooperative outcome in the second round among inexperienced and experienced players are equally slim.

Another example is the (N=10; X=\$1) treatment, in which in the first seven supergames close to 50.0 percent of the mass corresponds to strategies that start by cooperating. This changes in the last seven supergames in which almost 70.0 percent of the mass corresponds to strategies that start by defecting. These results suggest that subjects in this treatment start by cooperating but in time switch to defect.

Moving to our extension treatments, we note that behavior in the first seven supergames of our within treatments does not coincide with that observed in the last seven supergames among within-treatment subjects. In the first seven supergames of $2 \rightarrow 4$ close to 40.0 percent of the mass corresponds to strategies that start by

G.10. While there is a non-negligible mass for grim trigger, any time a subject who plays a grim-trigger strategy is matched with a player who uses always defect or suspicious tit for tat they will start to defect in round two. Therefore, the likelihood of long-term cooperation is very small.

TABLE G.1. SFEM output in the last seven supergames

Strategies	(N=2; X=\$9)	(N=4; X=\$9)	(N=4; X=\$1)	(N=10; X=\$1)	$2 \to 4$	$4 \to 2$	Chat(3/4)	Chat(1/2)	*Two from four*
				Panel A.					
Always cooperate	0.017	0.000	0.231*	0.133***	0.014	0.073	0.083	0.016	0.017
	(0.025)	(0.009)	(0.123)	(0.044)	(0.017)	(0.067)	(0.066)	(0.023)	(0.033)
Always defect	0.452***	0.324***	0.182**	0.265***	0.549**	0.217***	0.000	0.613***	0.735***
	(0.153)	(0.017)	(0.092)	(0.016)	(0.245)	(0.077)	(0.003)	(0.116)	(0.067)
Grim trigger	0.289**	0.008	0.046	0.094***	0.185	0.140	0.669**	0.262***	0.101*
	(0.121)	(0.011)	(0.130)	(0.024)	(0.129)	(0.086)	(0.283)	(0.094)	(0.058)
Tit for tat	0.180	0.008	0.535***	0.094***	0.000	0.458***	0.248	0.000	0.148**
	(0.111)	(0.011)	(0.151)	(0.024)	(0.065)	(0.090)	(0.303)	(0.013)	(0.060)
Suspicious tit for tat	0.063	0.659	0.006	0.414	0.252	0.111	0.000	0.110	0.000
β	0.929	0.981	0.939	0.934	0.950	0.921	0.975	0.873	0.899
# Observations	1360	1320	1632	1500	1296	1304	1560	884	1152
				Panel B.					
$\beta^\dagger$	0.911	0.981	0.931	0.934	0.948	0.883	0.974	0.871	0.896
p-value†	¡ 0.000	1.000	¡ 0.000	1.000	0.196	¡ 0.000	0.094	0.794	0.008
$\beta^\ddagger$	0.804	0.978	0.840	0.897	0.895	0.819	0.883	0.809	0.863
p-value‡	¡ 0.000	¡0.000	¡ 0.000	¡ 0.000	¡ 0.000	¡ 0.000	¡ 0.000	¡ 0.000	¡ 0.000

(i) Bootstrapped standard errors in parentheses. Level of significance: *** 1 percent; ** 5 percent; * 10 percent. (ii) $\beta^\dagger$ corresponds to the β estimate in case tit for tat and suspicious tit for tat are excluded. (iii) p-value† reports the p-value of a likelihood ratio test in which the restricted model excludes tit for tat and suspicious tit for tat. (iv) $\beta^\ddagger$ corresponds to the β estimate in case grim trigger, tit for tat, and suspicious tit for tat are excluded. (v) p-value‡ reports the p-value of a likelihood ratio test in which the restricted model excludes grim trigger, tit for tat, and suspicious tit for tat.

cooperating. Meanwhile, in the last seven supergames, 80.0 percent of the mass is captured by strategies that start by defecting. This change in behavior is even more striking in our second within-subject treatment. In the first seven supergames of $4 \rightarrow 2$, slightly more than 80.0 percent of the mass corresponds to strategies that start by defecting. However, in the last seven supergames, only about a third of strategies start by defecting.

Similarly, we observe large differences between the the first and the last seven supergames in chat treatments in which the possibility to exchange messages is only introduced in the second half of the session. For example for $\delta = 3/4$, in the first seven supergames more than 80.0 percent of the mass corresponds to strategies that start with defection. In the last seven supergames more than 90.0 percent of the mass is captured by strategies that start with cooperation.

Finally, in the extension treatment where only two of four players are needed for a cooperative outcome we see less of an effect of learning as the session progresses. Between the first and the last seven supergames the mass associated with strategies that start with defecting decreases by less than 8 percentage points.

TABLE G.2. SFEM output in the first seven supergames

Strategies	$(N{=}2; X{=}\$9)$	$(N{=}4; X{=}\$9)$	$(N{=}4; X{=}\$1)$	$\binom{N=10}{X=\$1}$	$2 \to 4$	$4 \to 2$	Chat(3/4)	Chat(1/2)	*Two from four*
Always cooperate	0.048	0.017	0.290***	0.228***	0.028	0.036	0.015	0.016	0.013
	(0.034)	(0.024)	(0.083)	(0.055)	(0.036)	(0.027)	(0.019)	(0.021)	(0.017)
Always defect	0.517***	0.600**	0.287***	0.440**	0.511***	0.320	0.315	0.921***	0.716***
	(0.086)	(0.275)	(0.064)	(0.213)	(0.079)	(0.207)	(0.261)	(0.251)	(0.069)
Grim trigger	0.160*	0.000	0.000	0.308***	0.224**	0.000	0.179**	0.032	0.098**
	(0.091)	(0.021)	(0.066)	(0.077)	(0.102)	(0.024)	(0.085)	(0.024)	(0.047)
Tit for tat	0.146**	0.045	0.392***	0.024	0.120**	0.152	0.000	0.032	0.081*
	(0.058)	(0.044)	(0.114)	(0.090)	(0.056)	(0.114)	(0.048)	(0.023)	(0.043)
Suspicious tit for tat	0.129	0.338	0.032	0.000	0.117	0.492	0.491	0.000	0.092
β	0.874	0.921	0.851	0.811	0.874	0.893	0.912	0.900	0.844
# Observations	1840	1840	1992	1380	2088	1820	2080	1124	1868

Bootstrapped standard errors in parentheses. Level of significance: *** 1 percent; ** 5 percent; * 10 percent.

References

Aoyagi, Masaki, Venkataraman Bhaskar, and Guillaume R Fréchette (2019), “The impact of monitoring in infinitely repeated games: Perfect, public, and private.” *American Economic Journal: Microeconomics*, 11, 1–43.

Dal Bó, Pedro and Guillaume R Fréchette (2011), “The evolution of cooperation in infinitely repeated games: Experimental evidence.” *American Economic Review*, 101, 411–429.

Dal Bó, Pedro and Guillaume R Fréchette (2018), “On the determinants of cooperation in infinitely repeated games: A survey.” *Journal of Economic Literature*, 56, 60–114.

Dal Bó, Pedro and Guillaume R Fréchette (2019), “Strategy choice in the infinitely repeated prisoner’s dilemma.” *American Economic Review*, 109, 3929–3952.

Embrey, Matthew, Guillaume R Fréchette, and Ennio Stacchetti (2013), “An experimental study of imperfect public monitoring: Efficiency versus renegotiation-proofness.” NYU working paper.

Fudenberg, Drew, David G Rand, and Anna Dreber (2012), “Slow to anger and fast to forgive: Cooperation in an uncertain world.” *American Economic Review*, 102, 720–749.

Vespa, Emanuel (2020), “An experimental investigation of cooperation in the dynamic common pool game.” *International Economic Review*, 61, 417–440.

Vespa, Emanuel and Alistair J Wilson (2020), “Experimenting with equilibrium selection in dynamic games.” *Working Paper*.