

Online Appendix
Slow to Anger and Fast to Forgive:
Cooperation in an Uncertain World

Post-publication note (December 4th, 2015)

Fudenberg, Rand and Dreber (2012) analyzed experimental play of a noisy repeated prisoner's dilemma where the 'cooperative' action corresponds to paying a cost c to give a benefit $b > c$ to the other player, the per-period discount factor is $\frac{7}{8}$, and each period each player's intended action is implemented with probability $\frac{7}{8}$. The paper stated without proof that when $b/c=1.5$, the only Nash equilibrium is for both players to defect in every period. However, Mira Frick and Johannes Hörner have shown us that this claim is not correct, and indeed, they have constructed a perfect public equilibrium (Fudenberg, Maskin, and Levine [1994]) in which cooperation occurs in the first period.

The online Appendix E shows that of the strategies we consider in the paper, only ALLD is a Nash equilibrium when $b/c = 1.5$. For example, it is not an equilibrium for both players to use the "Grim" strategy when $b/c=1.5$, while this is an equilibrium with the other b/c ratios the paper considered. More generally, ALLD is the unique strongly symmetric equilibrium, but not (as the counterexample mentioned above shows) the unique perfect public equilibrium.

Appendix 0-A: Sample instructions

Instructions:

Thank you for participating in this experiment.

Please read the following instructions carefully. If you have any questions, do not hesitate to ask us. Aside from this, no communication is allowed during the experiment.

This experiment is about decision making. You will be randomly matched with other people in the room. None of you will ever know the identity of the others. Everyone will receive a fixed show-up amount of \$10 for participating in the experiment. In addition, you will be able to earn more money based on the decisions you and others make in the experiment. Everything will be paid to you in cash immediately after the experiment.

You will interact numerous times with different people. Based on the choices made by you and the other participants over the course of these interactions, you will receive between \$0 and \$30, in addition to the \$10 show-up amount.

You begin the session with 50 units in your account. Units are then added and/or subtracted to that amount over the course of the session as described below. At the end of the session, the total number of units in your account will be converted into cash at an exchange rate of 30 units = \$1.

The Session:

The session is divided into a series of interactions between you and other participants in the room.

In each interaction, you play a random number of rounds with another person. In each round you and the person you are interacting with can choose one of two options. Once the interaction ends, you get randomly re-matched with another person in the room to play another interaction.

The setup will now be explained in more detail.

The round

In each round of the experiment, the same two possible options are available to both you and the other person you interact with: A or B.

The payoffs of the options (in units)

Option	You will get	The other person will get
A:	-2	+8
B:	0	0

If your move is A then you will get -2 units, and the other person will get +8 units.

If your move is B then you will get 0 units, and the other person will get 0 units.

Calculation of your income in each round:

Your income in each round is the sum of two components:

- the number of units you get from the move you played
- the number of units you get from the move played by the other person.

Your round-total income for each possible action by you and the other player is thus

		Other person	
		A	B
You	A	+6	-2
	B	+8	0

For example:

If you play A and the other person plays A, you would both get +6 units.

If you play A and the other person plays B, you would get -2 units, and they would get +8 units.

If you play B and the other person plays A, you would get +8 units, and they would get -2 units.

If you play B and the other person plays B, you would both get 0 units.

Your income for each round will be calculated and presented to you on your computer screen.

The total number of units you have at the end of the session will determine how much money you earn, at an exchange rate of 30 units = \$1.

Each round you must enter your choice within 30 seconds, or a random choice will be made.

A chance that the your choice is changed

There is a $7/8$ probability that the move you choose actually occurs. But with probability $1/8$, your move is changed to the opposite of what you picked. That is:

When you choose A, there is a $7/8$ chance that you will actually play A, and $1/8$ chance that instead you play B. The same is true for the other player.

When you choose B, there is a $7/8$ chance that you will actually play B, and $1/8$ chance that instead you play A. The same is true for the other player.

Both players are informed of the moves which actually occur. Neither player is informed of the move chosen by the other. Thus with $1/8$ probability, an error in execution occurs, and you never know whether the other person's action was what they chose, or an error.

For example, if you choose A and the other player chooses B then:

- With probability $(7/8)*(7/8)=0.766$, no changes occur. You will both be told that your move is A and the other person's move is B. You will get -2 units, and the other player will get +8 units.
- With probability $(7/8)*(1/8)=0.109$, the other person's move is changed. You will both be told that your move is A and the other person's move is A. You both will get +6 units.
- With probability $(1/8)*(7/8)=0.109$, your move is changed. You will both be told that your move is B and the other person's move is B. You will both get +0 units.
- With probability $(1/8)*(1/8)=0.016$, both your move and the other person's moves are changed. You will both be told that your move is B and the other person's move is A. You will get +8 units and the other person will get -2 units.

Random number of rounds in each interaction

After each round, there is a $7/8$ probability of another round, and $1/8$ probability that the interaction will end. Successive rounds will occur with probability $7/8$ each time, until the interaction ends (with probability $1/8$ after each round). Once the interaction ends, you will be randomly re-matched with a different person in the room for another interaction. Each interaction has the same setup. You will play a number of such interactions with different people.

You will not be paired twice with the same person during the session, or with a person that was previously paired with someone that was paired with you, or with someone that was paired with someone that was paired with someone that was paired with you, and so on. Thus, the pairing is done in such a way that the decisions you make in one interaction cannot affect the decisions of the people you will be paired with later in the session.

Summary

To summarize, every interaction you have with another person in the experiment includes a random number of rounds. After every round, there is a $7/8$ probability of another round. There will be a number of such interactions, and your behavior has no effect on the number of rounds or the number of interactions.

There is a $1/8$ probability that the option you choose will not happen and the opposite option occurs instead, and the same is true for the person you interact with. You will be told which moves actually occur, but you will not know what move the other person actually chose.

At the beginning of the session, you have 50 units in your account. At the end of the session, you will receive \$1 for every 30 units in your account.

You will now take a very short quiz to make sure you understand the setup.

The session will then begin with one practice round. This round will not count towards your final payoff.

Appendix O-B – Demographic statistics by session

b/c	error	Female	Economics	Age
1.5	1/8	52%	21%	20.9
1.5	1/8	48%	18%	21.0
1.5	1/8	50%	10%	19.6
<i>Average</i>		<i>50%</i>	<i>17%</i>	<i>20.5</i>
2	1/8	45%	14%	20.5
2	1/8	43%	10%	20.4
<i>Average</i>		<i>44%</i>	<i>12%</i>	<i>20.5</i>
2.5	1/8	38%	13%	21.3
2.5	1/8	64%	9%	19.5
2.5	1/8	52%	17%	20.7
<i>Average</i>		<i>52%</i>	<i>13%</i>	<i>20.4</i>
4	1/8	62%	17%	22.8
4	1/8	69%	6%	22.1
4	1/8	61%	13%	21.8
4	1/8	42%	30%	20.3
<i>Average</i>		<i>59%</i>	<i>16%</i>	<i>21.7</i>
4	1/16	35%	17%	22.7
4	1/16	44%	19%	21.0
4	1/16	69%	21%	21.2
<i>Average</i>		<i>47%</i>	<i>19%</i>	<i>21.6</i>
4	0	41%	6%	24.9
4	0	44%	6%	22.8
4	0	36%	8%	21.0
<i>Average</i>		<i>41%</i>	<i>7%</i>	<i>23.2</i>

Appendix O-C – Sequence of game lengths

b/c	e	t0 [†]	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10	t11	t12	t13	t14	t15
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16		
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8					
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8					
2	1/8	3	8	7	10	7	8	9	5	11	9						
2	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16	4	
2.5	1/8	3	8	7	10	7	9	10	5	11							
2.5	1/8	3	8	7	10	7	9	10	5	11	9	8	7				
2.5	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	11		
4	1/8	3	8	7	10	1	8	9	5	11	9	8	7	8			
4	1/8	3	8	7	10	8	9	6	11	7							
4	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16	4	9
4	1/8	3	8	7	10	7	8	9	5	11	9	8					
4	1/16	3	8	7	10	7	8	9	5	11	9						
4	1/16	3	8	7	10	7	8	9	5	11	9	8	7	8			
4	1/16	3	8	7	10	7	8	9	5								
4	0	3	8	7	10	7	8	9	5	11	9						
4	0	3	8	7	10	7	8	9									
4	0	3	8	7	10	7	8	9	5	11	9						

[†] t0 was a practice round that did not count toward the players' earnings.

The sequence of games in each session is shown in the preceding table. The starting place in the sequence of random game lengths that was used in the experiment was picked by the programmer, and the sequence following the chosen starting place had an unusually low number of short games. Although the average probability over all rounds of the game continuing was not significantly different from 7/8 (t-test: $p > 0.10$ for all sessions), the overall distribution of game lengths differed significantly from what would be expected using a geometric distribution (Chi² goodness of fit test, $p < 0.05$ in all sessions).

This raises the concern that over the course of the session, subjects might have come to believe that the game was more likely to end in later rounds than in earlier ones,

and adjusted their play accordingly. Particularly, we might expect that the tendency for cooperation to decrease over the course of an interaction would be greater in later interactions; or put differently, that the relationship between round number and cooperation will become more negative as interaction number increases.

It is not clear why such an effect would alter our findings, but nonetheless we check for evidence of this occurring. To do so, we ran a logistic regression with robust standard errors clustered on subject and session, including controls for b/c ratio and error rate. In addition to round number and interaction number as independent variables, we also add a [round X interaction] term. A significant negative coefficient on the [round X interaction] term would indicate that in later interactions, subjects are less likely to cooperate in later rounds, suggesting that after subjects have had time to learn that games are more likely to end in later rounds, they become more likely to defect in those later rounds. However, we find no evidence of a such relationship between round number and interaction number, (the coefficient for the [round X interaction] term in the regression is not significantly different from 0), either when considering all histories (coeff=-0.003, p=0.330), only considering histories where in the previous round both players cooperated (coeff=-0.004, p=0.418), only considering histories with the possibility of leniency (coeff=-0.004, p=0.698) or only considering histories with the possibility of forgiveness (coeff=-0.001, p=0.832). This suggests that increasing experience with the game length distribution did not affect subjects' probability to cooperate in later rounds, and in particular did not affect their levels of leniency or forgiveness.

A second complication with the game lengths is that due to technical difficulties with the computer software, the actual sequence of game lengths deviated somewhat from the pre-generated sequence in 4 out of the 18 sessions. We did not inform subjects of this error (which we were unaware of at the time) and they were most likely not aware of it either; no subject commented on the issue, and our experience is that subjects in our experiments do communicate with us when they are aware of software errors.

Appendix O-D – Equilibrium calculations

O-D.1 Equilibrium calculation for TFT with error

If both players use TFT then all histories fall into one of 4 phases defined by play in the previous round, while what happened 2 rounds ago doesn't matter either for current play or for continuation:

- P1: CC yesterday. Here the strategy says to play C.
- P2: CD yesterday. Here the strategy says to play D.
- P3: DC yesterday. Here the strategy says to play C.
- P4: DD yesterday. Here the strategy says to play D

Regardless of the current phase, next round's phase is P1 if both play C; P2 if player 1 plays C while player 2 plays D; P3 if player 1 plays D while player 2 plays C; and P4 if both play D.

Let the payoffs from following TFT in these phases be v_1, v_2, v_3, v_4 respectively, and let the error rate and discount factor be e and δ respectively. Then

$$v_1 = (1 - e)^2(b - c + \delta v_1) + e(1 - e)(b - c + \delta(v_2 + v_3)) + e^2 \delta v_4$$

$$v_2 = e(1 - e)(b - c + \delta(v_1 + v_4)) + (1 - e)^2(b + \delta v_3) + e^2(-c + \delta v_2)$$

$$v_3 = e(1 - e)(b - c + \delta(v_1 + v_4)) + (1 - e)^2(-c + \delta v_2) + e^2(b + \delta v_3)$$

$$v_4 = e^2(b - c + \delta v_1) + e(1 - e)(b - c + \delta(v_2 + v_3)) + (1 - e)^2 \delta v_4.$$

For example, consider v_1 . Here both players intend to play C. Thus with probability $(1-e)^2$, no errors occur, both players play C and remain in phase 1 for the next round; and therefore player 1 earns $(b-c)$ now plus the value of v_1 discounted by δ . With probability $e(1-e)$, player 1 (only) makes an error and accidentally plays D, shifting to phase 3 for the next round; therefore player 1 earns b today plus the value of v_3 discounted by δ . Also with probability $e(1-e)$, player 2 (only) makes an error, exploiting player 1 and shifting to phase 2 for the next round; therefore player 1 earns $-c$ today plus the value of v_2 . Together this results in the 2nd term $e(1-e)(b-c+\delta(v_2+v_3))$. Finally, with probability e^2 both players error and play D, shifting to phase 4 in the next round; here both earn 0 today, plus the value of v_4 discounted by δ .

It is enough to consider deviations in P1 and P2 to show that TFT is never an equilibrium in the presence of noise; for TFT to be an equilibrium, it is necessary (but not sufficient) to have

$$v_1 \geq e(1-e)(b-c+\delta v_1) + (1-e)^2(b+\delta v_3) + e^2(-c+\delta v_2) + e(1-e)\delta v_4$$

and

$$v_2 \geq (1-e)^2(b-c+\delta v_1) + e(1-e)(b+\delta v_3) + e(1-e)(-c+\delta v_2) + e^2\delta v_4.$$

However, these two conditions are mutually exclusive. Thus TFT is never an equilibrium.¹

O-D.2 Equilibrium calculation for PTFT with error

If both players use PTFT then all histories fall into one of 2 phases:

- P1: CC or DD yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C or both play D, else P2.
- P2: CD or DC yesterday. Here the strategy says to play D, and what happened 2 days ago again doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C or both play D, else P2.

Let the payoffs from following PTFT in these phases be v_1, v_2

respectively, and let the error rate and discount factor be e and δ respectively. Then

$$v_1 = (1-e)^2(b-c) + e(1-e)(b-c) + \delta((1-e)^2 + e^2)v_1 + 2e(1-e)v_2$$

$$v_2 = (1-e)^2\delta v_1 + e^2(b-c+\delta v_1) + e(1-e)((b-c) + 2\delta v_2).$$

¹ Moreover, a forgetful version of TFT that forgets the current state and picks a new state randomly with some non-zero probability is also never an equilibrium. For TFT, forgetting is equivalent to increasing the error rate e , as follows. From the point of view of an individual TFT player, there are two states: s1 (the other played C last round) and s2 (the other played D last round). In state s1, TFT plays C with probability $1-e$ and play D with probability e ; in state s2, TFT plays C with probability e and play D with probability $1-e$. Now imagine that players have a probability $2p$ of forgetting the state and randomly picking a new state, such that with probability p a player switches state. When the state is forgotten in s1 (with probability p), the player gets 'confused' and switches to state s2, and therefore intends to defect. Thus a TFT player in state s1 plays C with probability $(1-p)(1-e)+ep$ and plays D with probability $p(1-e)+(1-p)e$. This is equivalent to non-forgetful TFT with error rate $e'=e+p-2ep$. The same is true when considering a player that forgets in state s2, who plays C with probability $p(1-e)+(1-p)e$ and plays D with probability $(1-p)(1-e)+ep$. As shown above, TFT is never an equilibrium regardless of the value of e . Therefore forgetful TFT is never an equilibrium.

Following PTFT is clearly optimal in P2, so it is enough to check for a one-stage deviation in P1. Thus PTFT is an equilibrium if


$$v_1 \geq (1 - e)^2(b + \delta v_2) + e^2(-c + \delta v_2) + e(1 - e)(b - c + 2\delta v_1)$$

For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .

e	b/c	v_1	v_2	Phase 1 deviation payoff	Is PTFT an equilibrium?
1/8	1.5	5.85	5.1	6.98	No
1/8	2	11.7	10.2	12.46	No
1/8	2.5	17.55	15.3	17.94	No
1/8	4	35.11	30.61	34.39	Yes
1/16	4	40.69	35.44	38.93	Yes
0	4	48	42	44.75	Yes


To more fully explore that range of e values over which PTFT is an equilibrium, we plot the payoff advantage of deviating in P1 as a function of e , for $\delta = 7/8$, $c = 2$ and various b . We see that the for payoff specifications where PTFT is not an equilibrium at $e = 1/8$, increasing e does not lead to PTFT becoming an equilibrium.²

Advantage of Phase 1 Deviation, $b/c \in [1.5, 2.5]$




² A forgetful version of PTFT that forgets the current state and picks a new state randomly with some non-zero probability is also not an equilibrium at $b/c=1.5$, 2 or 2.5, as forgetting for PTFT is equivalent to increasing the error rate e , for similar reasons as for TFT. As shown above, increasing the error rate e cannot make PTFT an equilibrium for $b/c=1.5$, $b/c=2$ or $b/c=2.5$ with $\delta=7/8$, and therefore forgetful PTFT is not an equilibrium for those value.


Advantage of Phase 1 Deviation, $b|c|2$



Advantage of Phase 1 Deviation, $b|c|2.5$



Advantage of Phase 1 Deviation, $b|c|4$



O-D.3 Payoffs of 2x2 game between TFT and ALLD and between Grim and ALLD with error probability $e=1/8$ and $\delta=7/8$

TFT vs. ALLD

Grim vs. ALLD

b/c = 1.5

b/c = 1.5

	TFT	ALLD
TFT	5.09	-1.81
ALLD	5.23	1

	Grim	ALLD
Grim	3.27	-0.66
ALLD	3.49	1

b/c = 2

b/c = 2

	TFT	ALLD
TFT	10.18	-0.81
ALLD	7.625	2

	Grim	ALLD
Grim	6.54	0.34
ALLD	5.32	2

b/c = 2.5

b/c = 2.5

	TFT	ALLD
TFT	15.27	0.19
ALLD	10.03	3

	Grim	ALLD
Grim	9.82	1.34
ALLD	7.15	3

b/c = 4

b/c = 4

	TFT	ALLD
TFT	30.55	3.19
ALLD	17.25	6

	Grim	ALLD
Grim	19.63	4.34
ALLD	12.64	6

O-D.4 Payoffs of 2x2 game between TFT and ALLD and between Grim and ALLD with error probability $e=1/16$ and $\delta=7/8$

TFT vs. ALLD

b/c = 1.5

	TFT	ALLD
TFT	5.87	-2.02
ALLD	4.27	0.50

b/c = 2

	TFT	ALLD
TFT	11.73	-1.52
ALLD	6.03	1.00

b/c = 2.5

	TFT	ALLD
TFT	17.60	-1.02
ALLD	7.79	1.50

b/c = 4

	TFT	ALLD
TFT	35.20	0.48
ALLD	13.06	3.00

Grim vs. ALLD

b/c = 1.5

	Grim	ALLD
Grim	4.29	-1.34
ALLD	3.27	0.50

b/c = 2

	Grim	ALLD
Grim	8.58	-0.84
ALLD	4.69	1.00

b/c = 2.5

	Grim	ALLD
Grim	12.87	-0.34
ALLD	6.11	1.50

b/c = 4

	Grim	ALLD
Grim	25.73	1.16
ALLD	10.38	3.00

O-D.5 Equilibrium calculation for Grim with error

If both players use Grim then all histories fall into one of 2 phases:

- P1: No D yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C else P2
- P2: at least one D in the last round: play D. This phase is absorbing.

Let the payoffs from following Grim in these phases be v_1, v_2 respectively, and let the error rate and discount factor be e and δ respectively.

$$\text{Then } v_1 = (1 - e)^2(b - c) + e(1 - e)(b - c) + \delta[(1 - e)^2 v_1 + (2e - e^2)v_2],$$

$$v_2 = e^2(b - c) + e(1 - e)(b - c) + \delta v_2 \rightarrow v_2 = (e^2(b - c) + e(1 - e)(b - c)) / (1 - \delta)$$

Following Grim is clearly optimal in P2, so it is enough to check for a one-stage deviation in P1. Thus Grim is an equilibrium if

$$v_1 \geq (1 - e)^2(b + \delta v_2) + e(1 - e)(b - c + \delta v_1) + e(1 - e)(\delta v_2) + e^2(-c + \delta v_2)$$

For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .

e	b/c	v_1	v_2	Phase 1 deviation payoff	Is Grim an equilibrium?
1/8	1.5	3.27	1	3.47	No
1/8	2	6.54	2	5.43	Yes
1/8	2.5	9.82	3	7.40	Yes
1/8	4	19.63	6	13.30	Yes
1/16	4	25.73	3	11.17	Yes
0	4	48	0	8	Yes

O-D.6 Equilibrium calculation for Grim2 with error

If both players use Grim2 then all histories fall into one of 3 phases:

- P1: No D yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C else P2
- P2: D yesterday but none the day before: play C. Next round's phase is P1 if both play C else P3.
- P3: at least one D in each of the last two rounds: play D. This phase is absorbing.

Let the payoffs from following Grim2 in these phases be v_1, v_2, v_3 respectively, and let the error rate and discount factor be e and δ respectively.

$$\text{Then } v_1 = (1 - e)^2(b - c) + e(1 - e)(b - c) + \delta[(1 - e)^2 v_1 + (2e - e^2)v_2],$$

$$v_2 = (1 - e)^2(b - c + \delta v_1) + e(1 - e)(b - c) + (2e - e^2)\delta v_3]$$

$$v_3 = e^2(b - c) + e(1 - e)(b - c) + \delta v_3 \rightarrow v_3 = (e^2(b - c) + e(1 - e)(b - c))/(1 - \delta)$$

Following Grim2 is clearly optimal in P3, so it is enough to check for one-stage deviations in P1 and P2. Thus Grim2 is an equilibrium if

$$v_1 \geq (1 - e)^2(b + \delta v_2) + e(1 - e)(b - c + \delta v_1) + e(1 - e)(\delta v_2) + e^2(-c + \delta v_2) \text{ and}$$


$$v_2 \geq (1 - e)^2(b + \delta v_3) + e(1 - e)((b - c) + \delta v_1) + e(1 - e)(\delta v_3) + e^2(-c + \delta v_3)$$

For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .


e	b/c	v_1	v_2	v_3	Phase 1 deviation payoff	Phase 2 deviation payoff	Is Grim2 an equilibrium?
1/8	1.5	5.69	4.89	1	6.73	3.70	No
1/8	2	11.38	9.78	2	11.96	5.90	No
1/8	2.5	17.07	14.68	3	17.20	8.10	No
1/8	4	34.14	29.35	6	32.89	14.69	Yes
1/16	4	41.85	38.12	3	40.92	11.99	Yes
0	4	48	48	0	50	8	No

To more fully explore that range of e values over which Grim2 is an equilibrium when $b=8$, $c=2$ and $\delta=7/8$, we plot the payoff advantage of deviating in each state as a function of e . Numerical calculation shows that Grim2 is an equilibrium when $0.033 < e < 0.278$.

Advantage of Phase 1 Deviation



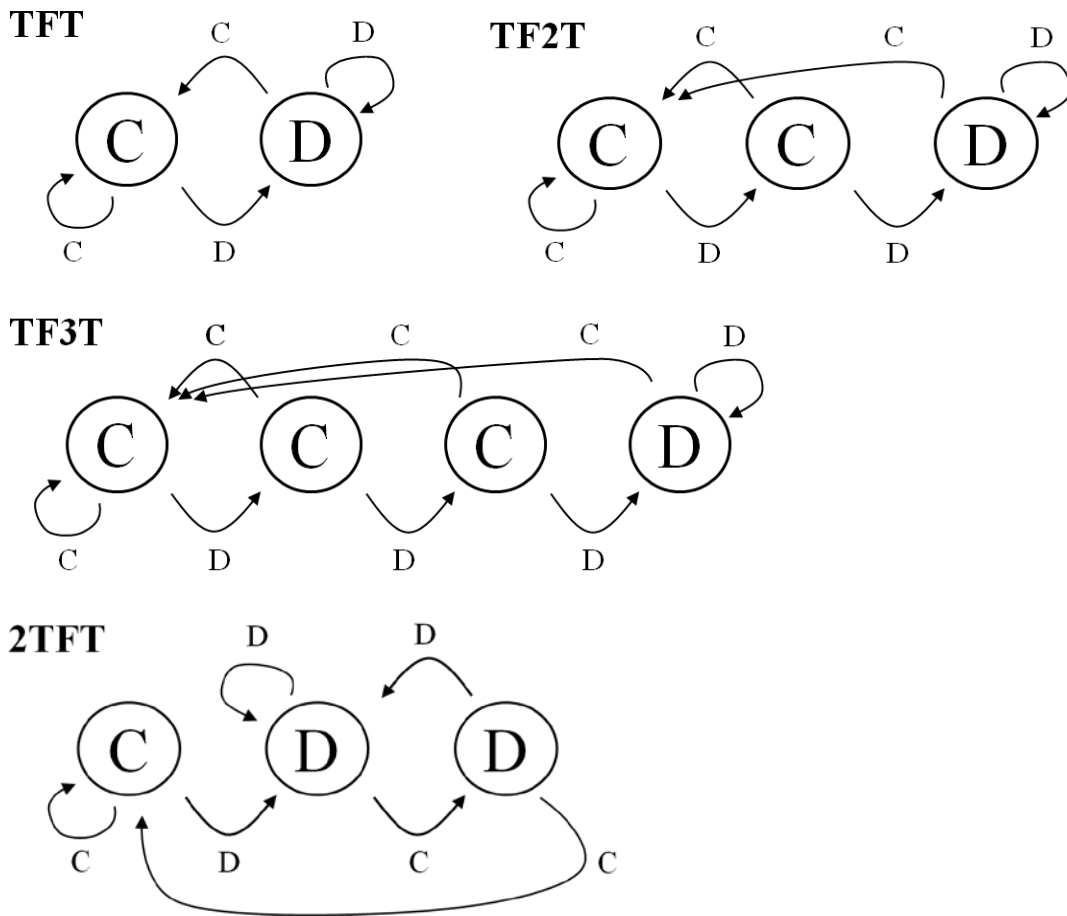
Advantage of Phase 2 Deviation




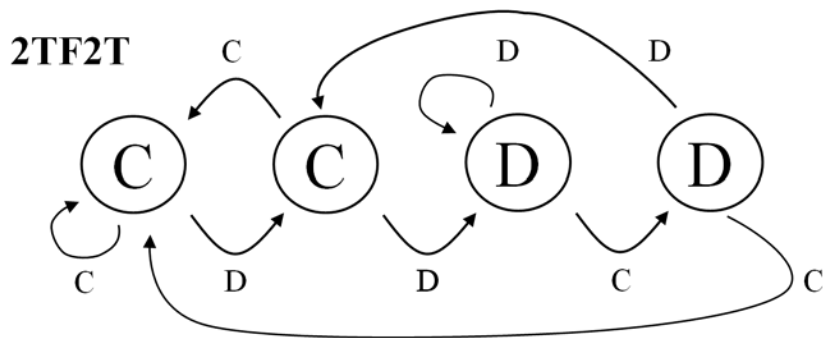


Appendix O-E – Strategy definitions

Here we define each strategy included in our analysis. Each phase is represented by a circle, with the strategy's move in that phase shown in the center of the circle, and transitions out of the phase indicated with arrows.

We begin with the strategies where transitions between phases depend only on the partner's move in the previous round: TFT, TF2T, TF3T, 2TFT, 2TF2T, D-TFT, D-TF2T and D-TF3T. For clarity we indicate only the partner's move with each transition arrow. All strategies begin in the leftmost phase.







Next we define the strategies where transitions depend on both players' actions in the previous round: Grim, Grim2, Grim3, D-Grim2, D-Grim3, PTFT, 2PTFT and T2. The last round histories associated with each transition are indicated by the pair $X_i X_j$ where X_i is the strategy's move last round and X_j is the partner's move last round (i.e. CD represents histories where the strategy played C last round while the partner played D).

When only one transition out of a phase exists, irrespective of either player's action, the transition is labeled "All".


Grim




Grim2




Grim3




D-Grim2





D-Grim3







PTFT





Lastly, we define the strategies whose transitions do not depend on previous histories of play: ALLC, ALLD, C-to-ALLD and DC-Alt.



Appendix O-F Robustness of strategy analysis

In the main text, we analyze the last 4 interactions of each session to minimize the effects of learning. Here we replicate our main analyses considering instead the last 2 or last 6 interactions, and find little difference, as shown in the table below. Regardless of the cutoff, we find that cooperation and leniency increase substantially going from $b/c=1.5$ to $b/c=2$, while forgiveness is changes little between $b/c=1.5$ and $b/c=2$, and then steadily increases from $b/c=2$ to $b/c=4$.

Last 2 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	53%	79%	78%	77%
%C All Rounds	33%	45%	61%	64%
Leniency	30%	64%	67%	69%
Forgiveness	18%	16%	27%	45%
MLE aggregation				
Cooperative strategies	59%	84%	83%	78%
Lenient strategies	21%	53%	62%	61%
Forgiving strategies	27%	24%	41%	58%

Last 4 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	54%	75%	79%	76%
%C All Rounds	32%	49%	61%	59%
Leniency	29%	63%	67%	66%
Forgiveness	15%	18%	33%	32%
MLE aggregation				
Cooperative strategies	57%	83%	81%	77%
Lenient strategies	18%	62%	60%	63%
Forgiving strategies	31%	31%	44%	57%

Last 6 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	54%	75%	78%	75%
%C All Rounds	32%	49%	60%	58%
Leniency	30%	59%	68%	64%
Forgiveness	14%	19%	31%	32%
MLE aggregation				
Cooperative strategies	59%	82%	83%	77%
Lenient strategies	23%	65%	63%	64%
Forgiving strategies	36%	37%	46%	57%

Our MLE estimation procedure assumes that the probability of mental error in strategy implementation, parameterized by γ , is equal across strategies. It is possible,

however, that some strategies are more difficult to implement than others and therefore γ may vary across strategies. Here we replicate our MLE estimates from Table 3, now allowing each strategy to have a different γ . We find little difference.

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
	E=1/8	E=1/8	E=1/8	E=1/8
ALLC	0	0.0195	0.0486	0.062
TFT	0.1894	0.0668	0.0864	0.0601
TF2T	0.0473	0.019	0.2109	0.2137
TF3T	0.0122	0.0186	0.0343	0.0616
2TFT	0.069	0.0558	0.0234	0.0381
2TF2T	0.0048	0.0903	0.0887	0.1195
Grim	0.1551	0.1324	0.1054	0.0297
Grim2	0.0567	0.1122	0.0362	0.0902
Grim3	0.0429	0.3147	0.1685	0.1188
ALLD	0.2362	0.1408	0.1396	0.1669
D-TFT	0.1865	0.0299	0.0579	0.0396
γ -ALLC	0.6517	0.3617	0.0438	0.2101
γ -TFT	0.4615	0.505	0.4034	0.3392
γ -TF2T	0.4329	3.8609	0.6046	0.4275
γ -TF3T	0.2886	0	0.4152	0.296
γ -2TFT	0.7757	0.4042	0.6942	0.5697
γ -2TF2T	0.4689	0.4188	0.419	0.3959
γ -Grim	0.4602	0.6609	0.6798	0.4077
γ -Grim2	0.4536	0.4062	0.6554	1.2917
γ -Grim3	0.4474	0.5662	0.408	0.4535
γ -ALLD	0.3165	0.2935	0.356	0.2086
γ -D-TFT	0.6247	1.4884	0.8278	10
Cooperative strategies	58%	83%	80%	79%
Lenient strategies	16%	57%	59%	67%
Forgiving strategies	32%	27%	49%	56%

We now show that our MLE estimates are robust to including two alternative classes of simple strategies. Firstly, we consider forgetful memory-1 strategies (we noted above that these strategies are not equilibria at $b/c=2$ or $b/c=2.5$, but subjects might be playing them anyway). To include F-TFT and F-PTFT in the MLE, we take advantage of the fact that forgetting is equivalent to experiencing a higher error rate. Thus we add an additional parameter γ_F to the MLE which represents the additional probability of mental error for forgetful players (relative to non-forgetful players). The MLE terms for F-TFT and F-PTFT are therefore

$$p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + \exp(-s_{ikr}(s) / (\gamma + \gamma_F))} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s) / (\gamma + \gamma_F))} \right)^{1 - y_{ikr}}$$

where $p_i(s)$ is the likelihood of strategy s given the history of subject i , y_{ikr} is the actual decision of subject i in round r of interaction k ($0=D$, $1=C$), s_{ikr} is the predicted move of strategy s ($-1=D$, $1=C$) and γ parameterizes the mental error rate of non-forgetful strategies. A referee also suggested that subjects might be playing simple strategies which simply ignore their partner's first move before triggering (i.e. always cooperate in the first 2 interactions). To test for this, we also include the strategy C-TFT with always plays C for the first 2 periods then switches to TFT, and C-Grim with always plays C for the first 2 periods then switches to Grim. As in the main text, we estimate the frequency of each strategy using MLE and determine whether a strategy is present at frequency significantly greater than 0 using a t-test with bootstrapped standard errors.

As can be seen in the following table, we find little of any of these strategies. Bootstrapping standard errors shows that none are present at levels significantly greater than 0 ($p > 0.05$ for all 4 strategies in all 4 payoff specifications).

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
ALLC	0	0.02	0.01	0.06
TFT	0.15	0.04	0.08	0.06
TF2T	0.05	0	0.15	0.17
TF3T	0.01	0.01	0.05	0.08†
2TF2T	0	0.11	0.11	0.12
Grim	0.13	0.02	0.07	0.01
Grim2	0.05	0.15	0.02	0.01
Grim3	0.05	0.28	0.23	0.12
PTFT	0	0	0	0
2PTFT	0	0.03	0	0
2TFT	0.03	0.07	0.02	0.03
T2	0	0	0	0
ALLD	0.27	0.17	0.14	0.18
C-to-ALLD	0	0.02	0	0.01
D-TFT	0.09	0	0.02	0
D-TF2T	0	0	0.02	0
D-TF3T	0.01	0	0	0
D-Grim2	0.05	0	0	0
D-Grim3	0	0	0.01	0
DC-Alt	0	0	0	0
C-TFT	0.03	0.03	0	0.02
C-Grim	0.01	0.04	0.03	0.03
F-TFT	0.07	0.03	0.06	0.04
F-PTFT	0	0	0	0.05
<i>Gamma</i>	<i>0.43</i>	<i>0.48</i>	<i>0.46</i>	<i>0.37</i>
<i>Gamma_F</i>	<i>0.6</i>	<i>2.17</i>	<i>1.09</i>	<i>1.92</i>

A referee also suggested the strategy that cooperates until the fraction of D moves by the partner passes some threshold, at which point it switches permanently to defection. To test for this possibility, we re-analyze the data including this family of strategies. We include 9 strategies which stop cooperating once the fraction of Ds by the partner is greater than 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80% or 90%. We find that none of these strategies are present at a frequency significantly greater than 0 in any payoff specification ($p > 0.05$ for all). The MLE results are shown in the following table:

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
ALLC	0	0.012	0	0.04
TFT	0.1648	0.065	0.088	0.064
TF2T	0.0459	0	0.1719	0.194
TF3T	0.0099	0.038	0.0383	0.083
2TFT	0.0033	0.097	0.1038	0.093
2TF2T	0.126	0.066	0.1147	0.028
Grim	0.0497	0.187	0.0093	0.023
Grim2	0.0321	0.245	0.1997	0.065
Grim3	0.054	0.07	0	0.013
ALLD	0.2848	0.173	0.1407	0.227
D-TFT	0.1321	0	0.0478	0
FracD-10%	0	0	0	0.029
FracD-20%	0	0	0	0
FracD-30%	0.0273	0	0	0.016
FracD-40%	0.0224	0	0	0
FracD-50%	0	0.029	0.0859	0
FracD-60%	0.0446	0	0	0.077
FracD-70%	0.0031	0	0	0.048
FracD-80%	0	0	0	0
FracD-90%	0	0.017	0	0
<i>Gamma</i>	<i>0.453</i>	<i>0.497</i>	<i>0.49</i>	<i>0.423</i>

Appendix O-G MLE estimates including stochastically forgiving strategies

Our main analysis considers only deterministic strategies. Here we extend the MLE formulation to include mixed strategies. The original formulation is

$$p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + \exp(-s_{ikr}(s) / \gamma)} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s) / \gamma)} \right)^{1-y_{ikr}}$$

where $p_i(s)$ is the likelihood of strategy s given the history of subject i , y_{ikr} is the actual decision of subject i in round r of interaction k ($0=D$, $1=C$), s_{ikr} is the predicted move of strategy s ($-1=D$, $1=C$) and γ parameterizes the mental error rate. We replace this with a new formulation

$$p_i(s) = \prod_k \prod_r \left[s_{ikr} \left(\frac{1}{1 + \exp(-1 / \gamma)} \right) + (1 - s_{ikr}) \left(\frac{1}{1 + \exp(1 / \gamma)} \right) \right]^{y_{ikr}} \cdot \left[(1 - s_{ikr}) \left(\frac{1}{1 + \exp(-1 / \gamma)} \right) + s_{ikr} \left(\frac{1}{1 + \exp(1 / \gamma)} \right) \right]^{1-y_{ikr}}$$

where s_{ikr} now represents the probability that strategy s cooperates given the history preceding round r of interaction k (0 to 1).

We use this new formulation to consider stochastic conditional strategies. In particular we explore a family of ‘generous’ TFT (GTFT) strategies which have received significant attention in the evolutionary game theory literature (e.g. Nowak and Sigmund 1990). These strategies are stochastically forgiving. Like TFT, GTFT always responds to C with C. In response to an opponent’s D, however, GTFT stochastically cooperates with probability q .

First we analyze simulated data to see whether the MLE can differentiate between a GTFT that forgives defection 20% of the time and TF3T. We simulate a session with 10 ALLD players, 10 TF3T players, and 10 GTFT-2 players. We simulate 4 interactions, each lasting 8 rounds, and find that the MLE correctly assigns 1/3 to ALLD, TF3T and GTFT-2. Thus the MLE is able to distinguish between memory-1 stochastic forgiveness and longer deterministic forgiveness.

We now reanalyze our data using the 11 strategies from the main text Table 3 plus 9 GTFT variants- those which forgive defection with 10% probability (GTFT-1), 20% probability (GTFT-2), ... 90% probability (GTFT-9). We find a somewhat sizable fraction of people playing stochastically forgiving memory 1 strategies (between 10 and 19% in the treatments with cooperative equilibria). However, the inclusion of these strategies doesn't undermine the importance of lenient strategies with more than 1 period of memory – the longer memory lenient strategies are much more common:

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
GTFT-X	11%	12%	10%	19%
TF2T+TF3T+2TF2T+Grim2+Grim3	16%	57%	51%	43%

Thus we conclude that exploring stochastic strategies is an interesting avenue for future research, but that our main findings related to leniency are robust to including stochastic forgiveness.

Appendix O-H Logistic regression analyzing dependence on play 2 rounds ago

In the main text, we use a logistic regression to provide evidence that subjects' decisions are influenced by the partner's move two rounds ago as well as decisions in the previous round. A potential concern with this methodology lies in the possibility for such correlations to occur in heterogeneous populations of subjects with different strategies each of which conditions on at most play in the previous period. This is because the partner's decision two periods ago can give information about a subject's type even if it does not directly influence her decision. For example, consider a population of $\frac{1}{2}$ ALLD and $\frac{1}{2}$ TFT, and imagine that in the previous round both subject played C. If Player A played C two rounds ago then Player A is much more likely to be a TFT player than an ALLD player. This is because last round Player A played C in response to her partner's C two rounds ago, and this is consistent with TFT but not ALLD. Therefore Player A is also likely to play C now. If Player A's partner played D two rounds ago, however, then Player A is equally likely to be TFT or ALLD, because both strategies intend to play D in response to D. Therefore Player A is equally likely to play C or D in current round. And so if Player A's partner played C two rounds ago, she is more likely to play C now than if her partner played D last round, even though her strategy does not look back two periods.

Including controls for player type can help address this issue. For example, in a population of ALLD and TFT players, first round cooperation can do a good job of cleanly differentiating types. First round cooperation does not differentiate between TFT and GTFT, however, but GTFT players will have higher overall cooperation. Thus we include controls for both first round cooperation and overall cooperation in our regression.

To explore how pervasive of a problem bias stemming from heterogeneity might be for our analysis, and how effective our controls are for mitigating it, we conduct various simulations. We consider 5 different populations:

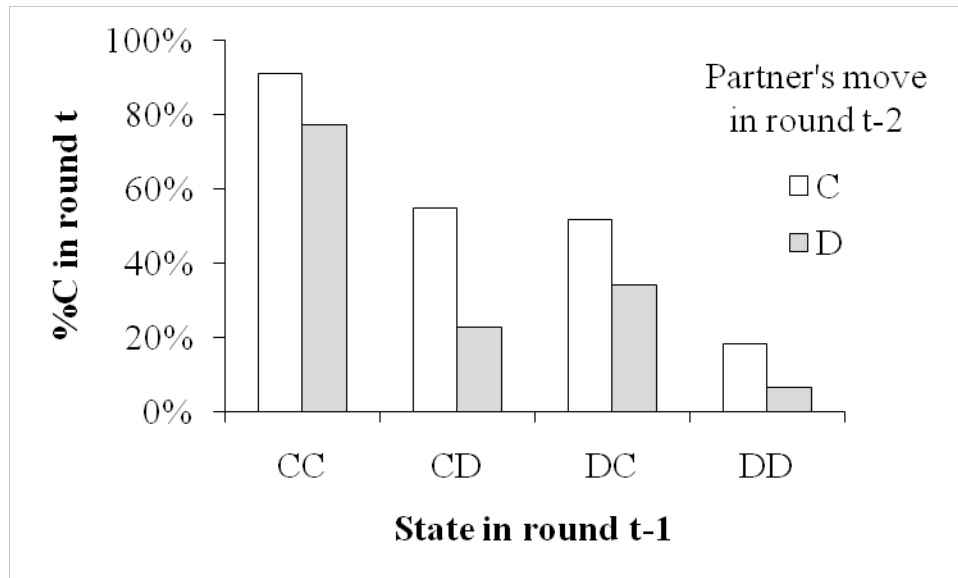
- i. $\frac{1}{3}$ ALLD, $\frac{1}{3}$ TFT, $\frac{1}{3}$ GTFT-5
- ii. $\frac{1}{3}$ ALLD, $\frac{1}{3}$ TFT, $\frac{1}{3}$ F-TFT (forgets the state with 10% probability)
- iii. $\frac{1}{2}$ ALLD, $\frac{1}{2}$ TFT

- iv. 1/2 75%-TFT 25%-CoinFlip, 1/2 25%-TFT 75%-CoinFlip (on each decision, these strategies randomly choose to either play according to TFT or select a move at random (i.e. CoinFlip) – one half the population plays TFT 75% of the time, the other half picks randomly 75% the time)
- v. 1/2 ALLD, 1/2 TF2T

Agents in populations (i), (ii), (iii) and (iv) do not condition on play two periods ago, while 1/2 of the agents in population (v) do. For each population, we simulate 3 sessions of 30 players each, playing 4 interactions each last 8 periods (for maximum comparability to our data). For each population we conduct 25 simulation replicates, and for each replicate we perform the logistic regression with correlated random effects described in the main text (own decision as a function of partner's move 2 rounds ago, own move 2 rounds ago, other's move last round, own move last round, own frequency of first round cooperation, and own frequency of overall cooperation). Across the 100 simulations using only memory 1 strategies, we find 4 instances in which other's play two rounds ago is significant at the $p < 0.05$ level, consistent with what is expected due to chance. In contrast, we consistently find a highly significant effect of play two periods ago for every replicate of (v), where agents are in fact conditioning on play two rounds ago.

Examining the size of the coefficient for partner's move 2 rounds ago, not just the p-value, gives further insight into the relative importance of earlier play. We find that in all 4 scenarios where agents are not actually conditioning on play 2 rounds ago, the coefficient for partner's move 2 rounds ago is at least an order of magnitude smaller than the coefficients for play 1 round ago (usually several orders of magnitude smaller). In the scenario where many agents are actually conditioning on play 2 rounds ago, conversely, all coefficients are of the same order of magnitude. In our analysis of the data from our experiments, we find the latter result rather than the former: a large coefficient on partner's move 2 rounds ago, on the same order of magnitude as the other coefficients. These results suggest that we are in fact picking up subjects explicitly conditioning on the outcome 2 periods ago, rather than only finding spurious correlations due to heterogeneity and stochastic strategies. Graphically displaying these relationships shows

a substantial and consistent effect of play two rounds ago, across all states in the previous round:



Appendix O-I Self-reported strategy descriptions

The post-experimental questionnaire included a free-response question in which subjects had the option of describing the strategy they used in the game. The responses of those subject who answered are reproduced here.

$b/c=1.5$, $E=1/8$

- I almost always started by choosing B and continued to play B for the subsequent rounds
- I chose B almost all the time because you are guaranteed not to lose any points with B
- I chose A to start with. If the other person also chose A, I switched to B the next round. If they too chose B, I stuck with it. If they chose A, I would make a 50/50 decision between A and B
- Once B was played by either player (including by me if my choice was switched) I played B for the rest of the interaction
- I chose B. I continued to chose B unless I felt the daring to press A. I chose B 98
- I'd start with B and continue if he played B. if he did A, I'd switch
- I started out choosing B, but if someone chose A twice I'd play A (I think). If someone was playing straight Bs, I'd occasionally play an A in an effort to guilt them into playing an A (at which point I'd have switched back to straight Bs)
- sometimes I started with B but rarely
- i started by choosing B. if the other person played A in the next round, i started playing A until the other person picked B in a later round. If, after that round, they chose A again, i continued to play A
- i chose B every time unless for two rounds in a row the other person played A
- i chose B most (if not all my rounds) because i wanted to maximize my points and minimize my losses
- random
- i chose A once then chose B the rest of the time
- i played A until they played B twice. Then i would switch to B. if they dont switch back to B, i would keep playing B or possibly do a one-off and switch to A (only if i have positive (or more than 2) points)
- i played B every chance i had because i know that the conservative strategy can win. However, i think i played A twice over the entire experiment because i was bored of pressing B
- always play A. switch to B if other player plays B more than 2 times in a row
- i basically always chose B
- always choose B. you cannot lose
- if it became clear we were going tit for tat with each other I'd try and break out of the cycle by playing A even if his or her last move were B

- i start by choosing B. if the other person plays A i will then choose A the next round. If the other person plays A twice in a row or every other time i then play A

b/c=2, E=1/8

- I chose B every time because it had the least risk
- start with A. always chose A unless the person goes B twice in a row
- you want to convey trustworthiness so when/if your response gets changed/changes you will remain with both As subsequently. I chose A until the other person played B 2x, unless it came in the middle of a string of As. I would return to A if my partner did
- I started with A. if I got B back, I would usually keep playing B unless they switched to A. if I got A back, I would usually play A, but slip in a few B's. hoping they would appear accidental.
- I played A until I thought the person was a B idiot, which was based on frequency not total amount of B's or A's chosen
- basically, I start with A and assume A from my opponent until otherwise noted. If opponent played B. I would switch to B until /unless my opponent went back to A indicating the B was a random switch by the comp from A.
- random plays: until the other player plays B. Then, play B afterwards.
- usually, I started by choosing B, then switched to A if the person played A twice in a row
- if the other person gave me 4 points and lost 2 points in the previous round I would try to do the same. if the other person consistently chose B I didn't want to chose A because I would only lose points and gain none. I would chose B after a round where I played A and the other played B, because in the previous round I gave up points to the other person so I expect them to do the same for me in this round
- AAABAABABB. always 3 As in the beginning unless the person played all Bs, in which case I did all Bs
- never played A
- along with A, build trust/rapport that would lead to best outcome for both
- wanted to avoid a B vs. B stalemate
- if I felt the round was ending and there was no incentive to further build the relationship, I might pick B. the other person might have chosen A, but his choice switched by chance. So to reciprocate his intentions, I might choose A
- I tried to use the tit for tat strategy to earn the most points in the long run. the first time the other player played B, I gave them the benefit of the doubt in that the computer had chosen for them.
- A makes the most sense. Both have incentive to gain whereas B only one person has incentive or both get nothing.
- lull them into thinking you were playing fairly (how barbaric)
- when I saw a string of "A" I felt I could trust them. String of "B"s, I didn't
- I would most likely play A for a few rounds as long as the other person was doing the same. Every now and then I would throw in a B to maximize my points. I would also play B if the other person had previously played B. I would never

choose B twice in a row because that would most likely lead to both persons choosing B for the remainder of the experiment. I would start each trial by choosing A. I would choose A until I decided to switch to gain extra points, or until the other person chose B. if that happened I would choose B the next round. If I chose B and the other person chose A, then I would always choose A the next round. rare exceptions: i would play B again only sometimes if the trial had been going for a while and i thought it would be the last round. if i ever chose B twice in a row, then i would choose B for the rest of the experiment unless it was really early and the other person played A twice

b/c=2.5, E=1/8

- my strategy was as follows: -start with A -after one B by other players, assume it was an error and continue with As -after two consecutive Bs by other player, assume he is a selfish jerk and stay with Bs for rest of interaction. Only when errors mixed up my responses and confused things did I have to become creative.
- I want to build trust with the other person. I wanted to give the other person the benefit of the doubt that his or her answer was changed by the computer.
- i didn't want to get taken advantage of.
- i wouldn't choose B after AA. After AB, I figured it could have been the glitch that made them possibly choose B 7/8 of the time - i wouldn't choose B, I'd give the other person one more chance. I always chose A first, and if my action was B it was because the 7/8 glitch made me or the other person continually played B.
- i never chose A, i always chose B.
- benefit of the doubt: maybe he meant to do A so he doesn't yet deserve retaliation
- i would play B after AA if i had a feeling the computer would change it to A.
- there was a possibility that the other person was picked A but B was what the result. everybody wanted to make as much money as they could and if they all picked B every round, nobody would get anything so it would be important to start the first few rounds with A that way I could let the other person know that if we cooperate, we can go a long way
- B could have been randomly thrown out and not their choice. to see if they randomly selected B and would throw more A choices and work cooperatively, I might choose A.
- choose A initially to earn trust so that I can eventually use B to earn more later.
- didn't really think about other person, but felt taking risk to gain money was better than getting \$0
- basically I took into consideration karma. I would click A again just to see if it wasn't the 1/8 probability that made the other person choose B, and knowing this determined the rest of my responses.
- collusion in the long run in game theory can lead to a greater maximum of welfare instead of Nash equilibrium
- I figured that if the other person had been choosing A, they would choose it again and I'd get 3 units or my 1/8th probability of my choice coming out the opposite had not yet occurred so I took my chances that it would w/ A.
- if there was no reason to believe the other person chose B on purpose/would continue to choose B, I would choose A.

- I would cooperate as long as I thought the other person was. If I had done several (2-3) rounds of A I would then do B because it would look more like it was unintentional.
- It depended a lot on the person I played. one person chose B every time so I had to do the same in order to not get taken advantage of. Other times I would mostly press A and throw in a few Bs if the person was cooperative. X .
- I pretty much always played A and assumed if the person played B it was an error. If a person played B twice in a row I would sometimes play B in the next round.
- Gain trust by starting with A for 4 rounds then change to B.
- If they play B more than once, attribute it to intention not error and respond accordingly.
- I would usually choose B during the first round, then A for a few rounds, and then B once or twice later on.
- I would always start with A, and play A the second round. If the person played B both times, I switched to B. If the person played A at least 1 of those items, I would continue with A until the 6th or 7th round. then I'd play B.
- I punished for multiple B choices, not when it was just one and could have been accidental.
- I just wanted to make the most money per round possible.
- sometimes I chose A to establish it as a pattern so deviations were more likely to be interpreted as error.
- If they had more than 2 rounds with B, I went with B too, as little chance this was due to error. Or i went with B if their choice of B was greater than 1/8.
- played A consistently until other person played B three times, then switched to B until end of round.
- Played for a number of periods and realized that the average number of periods was between 7-10, so played B to maximize profit closer to the end.
- i acted fully in self-interest. However this is a PD so my best interest goes alongside what of my "partner". the interpretation of a B depends on how many rounds of A-A have occurred. If only 1, i will think it is more probable that it was intentional. If we have gone through A-A 5 times or more, i feel it is very unlikely it was intentional.
- i gave very little thought to the feelings of the other player at all. I played B every time to maximize my points and assumed they were doing the same so that any playing of A was unintentional.
- I rarely chose A b/c it earned you least points and seemed most high-risk. risk averse: playing B nets me more points and rarely costs me.
- I gave everyone the benefit of a doubt 1 round before during and after they played B. defensively if after 3 of their B rounds in a row I would [illeg] B as well.
- by choosing A, both players were in a mutually beneficial state, but only as long as they both kept choosing.
- I tried tit for 2 tat.
- some people played B all the time so I didn't want to fall into that perpetual cycle of losing. if the person had played B before, I probably didn't think it was chance.

- mostly tit for tat strategy with the understanding that a defect might have been by error and also hoping the other would assume my defection was an error when the other was a cooperator.

b/c=4, E=1/8

- I appreciated their choosing A (I would always assume they chose it). hoped a B was a 1 time thing on their part and so didn't want to sour future relations.
- I selected B at times because I doubted the trustworthiness of the other player. [if?] I decided to trust the other person would consistently select A
- depends on how many times the player has picked B.
- I thought their B may have been the 1/8 that the answer was switched. If so, they might come back to A in this round and we'd be even again. if they kept playing B, I'd play B so I wouldn't lose more points and so they wouldn't get more points
- If I were continuously playing A then they played B out of nowhere I'd play B the next round out of revenge. No need to be greedy. if they were to continue to play B then I would to, that way no one wins or gains.
- I would choose A unless the other person picked B. If the other person picked B, the next round I would also pick B. if the person then picked A I would think that it was an error and then continue to click A. if the person pressed B, i would continue to choose B. if I had accidentally picked B, i would pick A and make it up to the other person. then if they had chosen B i would assume they were checking me and continue to choose A.
- I never chose B if the previous round was both As. B could have been the error so I was willing to give the other person benefit of the doubt.
- if I chose A, and the other person chose B, the next time I chose B as a defense mechanism so I didn't lose anymore points
- in most of the interactions, if the other person did not play A at first, they eventually played A consistently after seeing that I had played A for the start two or more times in a row. Therefore, a mutual strategy seemed to be derived where both players would play A every time in order to maximize the number of units gained in the whole interaction (mutual altruistic response). Additionally, whenever this strategy was employed and the screen showed one time that the other person chose B, i automatically attributed it to error and continued to play A for the following rounds.
- trying to get equilibrium of all-As
- it felt terrible being betrayed after a long "A" streak. assume at first that their B could be a mistake. .
- if both played A, 12 total points earned. That's why I played it.
- keep the other person honest. Compensate for errors.
- when someone, myself or the other got B I stuck with B straight through.
- after I got a B from someone I always gave them the benefit of the doubt that it was due to error. After that if I got another B I would usually give B as well.
- there was a chance it was the computer if they chose A & I chose A, then we could both get six points, which outweighs the risk of losing 2 point, which is something like 7 cents. If they were consistently choosing B (for 3 rounds) then I

would switch to B as well. However, I did so knowing we could both be better off with As. I always started with A for at least the 1st 2 rounds to see how much the other person wanted to cooperate. In rounds that we both began choosing B I never made as much money.

- I always played A unless they had played B.
- as the number of rounds increased, I was more likely to play B in the later rounds (expected number of rounds per interaction=8).
- B was the dominant strategy but gave better payoffs. I always chose B.
- since this game was completely anonymous and there was a chance that my choice would be changed, I felt no incentive to be nice and choose A .
- I started with A to demonstrate good faith and hoped that the other person would too so we could establish picking A. I had originally been skeptical and started with B, but then we both ended up choosing B for each round.
- if you are perceived as trustworthy, the interaction benefits both parties.

b/c=4, E=1/16

- Random chance might have produced their Bs so I want to try to possibly salvage the relationship. I would simply ignore occasional B's in a long string of A's
- The only way to earn money was to work together most of the time. I wasn't sure if B was the players intended decision.
- It could have been an error that B was played.
- Depends on if B was just one occurrence, or if they had played lots of Bs before
- I started with B in several rounds and stuck with it. The last few rounds I started with A.
- I would continue for a string of As, but if the other player played a B 2\3 times recently I would switch to Bs.
- didn't choose B at all during all interactions
- play A. If they go two Bs in a row play B (once or twice) then get back to A's
- if we both chose A, we typically stuck with that for the remainder of tee interaction.
- play A until other person chooses B X2, then play B and see what the outcome is. Play A next round etc.
- I never chose B on purpose, and to avoid getting into a locked system of B's I never played B.
- I chose A to begin with and then about every fourth time chose B to get more earnings. I then went back to A so as to blame the previous round on the 1\16th chance of A change. If the other person played B twice in response, I chose B the rest of the time.
- Always start with B. If other player is B, stay B until other player chooses A. If other player is A, stay A. If switch to B, give benefit of doubt due to 1\16 switch. If repeats, go to B.
- I always put A regardless of any choice given to me.
- I chose A for the whole interaction except for one B or if the other played B many times.

- I chose A until a few rounds passed. I then would throw in a B to make it seem as though the computer randomly chose.
- I played A as much as possible and then switched to B if they picked B twice in a row. Occasionally I'd try picking an A twice if we played for a long time at 0 to see if they wanted to actually switch to A.
- I started by choosing A, and continued to do so. If the other player chose B, I would choose B the next round. Then I would go back to A. I never chose B unprovoked, and if the computer generated that response for me, I would choose A the next round.
- I started by choosing A. Then I chose A always unless the other player chose B twice. If so, I chose B next round. If they chose B in that round, I chose B next. If they chose A in that round, I chose A next.
- I started choosing A until the other person chose B two times. Then I switched to B until the other person chose A. I immediately returned to choosing A as soon as the other person chose A. (A merciful strategy)
- I started play A 3 times to attempt to signal that I wanted to collaborate. If the other did not switch from B, I would change to B to attempt to change his/her play to A.
- If we got on a spree of both A's I'd pick B to get the +8 and then switch back. It eventually didn't really affect my earnings considerably.
- Started choosing B and then switched to choosing A or B once I ended up with 0 continuously.

b/c=4, E=0

- If they were helping me, I was willing to help them, 1\16 of a dollar isn't a lot to lose. If they chose B first, I felt like they were being greedy and I wanted to show them that strategy wouldn't pay off in the long run
- Always B.
- I hate when things don't make sense. To choose B in after we both chose A would hurt me because the other player would stop choosing A and I'd earn no points. Whether my motives are self-serving or altruistic, A is the most logical. If they've previously chosen B many times, I don't like them, and I don't want them to earn any points
- first I chose A. Then I started with B after seeing the first interaction, I would choose either A or B.
- going for B was not helping with long term sustained benefit.
- If the other person did A I would follow with A
- Both players win if both pick A, no one likes negative points
- If it were one of the middle rounds, I would definitely choose A. If one of the later rounds, most likely B.
- Played A every time. Only had 1 person one time choose B. Played A next round to give them a chance. I think if they had chosen B again I might have considered switching to B but my plan was to always play A.
- I chose A to take more money, collectively, from the university I would prefer another student have it that the university.

- Choosing anything other than A 100% of the time is just mean and vindictive. Its just a study why not help people get rich?
- I played A 3 to 4 times after which I only played B
- I started by choosing A and switched to B if the other player played B in the previous round, but would remain on B as long as the other player did too.
- I started by choosing B. then if the other person chose A in the next round I would choose A. If the other person chose B too, then I would continue playing B.