Online Appendices

Online Appendix A: Sample Instructions and Screens

i. Memory 1 elicitation

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation with cash vouchers, privately at the end of the session. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

Please turn off pagers and cellular phones now. Please close any program you may have open on the computer.

The entire session will take place through computer terminals, and all interaction between you will take place through the computers. Please do not talk or in any way try to communicate with other participants during the session.

We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the session and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

This experiment has three parts; we will start with part 1. Once this part is over, instructions for the other parts will be given to you. Your decisions in this part have no influence on the other parts.

General Instructions: Part 1

1. In this experiment you will be asked to make decisions in several rounds. You will be randomly paired with another person for a sequence of rounds. Each sequence of rounds is referred to as a match.

2. The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. This is as if we would roll a four sided die and end if 1, or 2 comes up and continue if 3, or 4 comes up. So, for instance, if you are in round 2, the probability there will be a third round is 50% and if you are in round 9, the probability there will be another round is also 50%.

3. Once a match ends, you will be randomly paired with another person for a new match.

4. The choices and the payoffs in each round are as follows:

	the other's choice	
your choice	1	2
1	32, 32	12, 50
2	50, 12	25, 25

The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are matched with.

- As you can see, this shows the payoff associated with each choice. Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

That is, if:

You select 1 and the other selects 1, you each make 32.

You select 1 and the other selects 2, you make 12 while the other makes 50.

You select 2 and the other selects 1, you make 50 while the other makes 12.

You select 2 and the other selects 2, you each make 25.

5. Between matches, you will be told the history of your decisions and the decisions of the person you were matched with in the last match.

- The first match to end after 20 minutes of play will mark the end of Part 1. At the end of the experiment you will be paid $0.0045 for every point scored. There is no show-up fee for this experiment.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

General Instructions: Parts 2 and 3

1. The basic structure of these parts is similar to part 1: the payoffs in each round are the same as before. However, now you are also asked to specify a plan of action.

2. A plan of action is specified by answering 5 questions:

- In round 1 select…
- After round 1 if

  I last selected 1 and the other selected 1, then select…

  I last selected 1 and the other selected 2, then select…

  I last selected 2 and the other selected 1, then select…

  I last selected 2 and the other selected 2, then select…

(please pay attention to the fact that these question will appear in varying orders on your screen).

3. However, to start, this plan of action will have no role, that is you will then take decisions, one round at a time, just as in part 1. Those decisions may or may not correspond to what you have selected as your plan of action. Between matches, you will be told the

history of your decisions, the decisions of the person you were matched with, and the decisions your plan of action specified given the actions you took and the person you were matched with took in the last match. Furthermore, in between each match, you have the opportunity to change your plan of action.

4. Part 2 will last for 20 minutes. Exactly when those 20 minutes have passed, part 3 will start. In part 3, the plan of action that you have specified will start taking all decisions for you. As of that point you will have no occasions to revise your plan of action. The transition between parts 2 and 3 can happen midway through a match. Part 3 will last for 15 matches. As with Part 1, at the end of the experiment you will be paid $0.0045 for every point scored in Parts 2 and 3.

- Are there any questions?

Before we start, let me remind you that just as in part 1, in parts 2 and 3:

- The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

ii. Menu Elicitation

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation with cash vouchers, privately at the end of the session. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

Please turn off pagers and cellular phones now. Please close any program you may have open on the computer.

The entire session will take place through computer terminals, and all interaction between you will take place through the computers. Please do not talk or in any way try to communicate with other participants during the session.

We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the session and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

This experiment has three parts; we will start with part 1. Once this part is over, instructions for the other parts will be given to you. Your decisions in this part have no influence on the other parts.

General Instructions: Part 1

1. In this experiment you will be asked to make decisions in several rounds. You will be randomly paired with another person for a sequence of rounds. Each sequence of rounds is referred to as a match.

2. The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. This is as if we would roll a two sided die and end if 1 comes up and continue if 2 comes up. So,

for instance, if you are in round 2, the probability there will be a third round is 50% and if you are in round 9, the probability there will be another round is also 50%.

3. Once a match ends, you will be randomly paired with another person for a new match.

4. The choices and the payoffs in each round are as follows:

	the other's choice	
your choice	1	2
1	48, 48	12, 50
2	50, 12	25, 25

The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are matched with.

- As you can see, this shows the payoff associated with each choice. Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

That is, if:

You select 1 and the other selects 1, you each make 48.

You select 1 and the other selects 2, you make 12 while the other makes 50.

You select 2 and the other selects 1, you make 50 while the other makes 12.

You select 2 and the other selects 2, you each make 25.

5. Between matches, you will be told the history of your decisions and the decisions of the person you were matched with in the last match.

- The first match to end after 20 minutes of play will mark the end of Part 1. At the end of the experiment you will be paid $0.0045 for every point scored. There is no show-up fee for this experiment.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

General Instructions: Parts 2 and 3

1. The basic structure of these parts is similar to part 1: the payoffs in each round are the same as before. However, now you are also asked to specify a plan of action.

2. A plan of action can be specified in two ways.

Either, by answering 5 questions:

- In round 1 select…

- After round 1 if

I last selected 1 and the other selected 1, then select…

I last selected 1 and the other selected 2, then select…

I last selected 2 and the other selected 1, then select…

I last selected 2 and the other selected 2, then select…

(please pay attention to the fact that these question will appear in varying orders on your screen).

Or, by choosing one of the plans of action described on the computer screen.

You will be presented with a list of plans, if you want instead to specify your plan of action by answering the 5 questions, you should select "Build your own plan". If you change your mind and want to return to the previous screen you can do that by pressing "Choose a plan from the menu". You can go back and forth until you confirm your choice by clicking "Submit".

3. However, to start, this plan of action will have no role, that is you will then take decisions, one round at a time, just as in part 1. Those decisions may or may not correspond

to what you have selected as your plan of action. Between matches, you will be told the history of your decisions, the decisions of the person you were matched with, and the decisions your plan of action specified given the actions you took and the person you were matched with took in the last match. Furthermore, in between each match, you have the opportunity to change your plan of action.

4. Part 2 will last for 20 minutes. Exactly when those 20 minutes have passed, part 3 will start. In part 3, the plan of action that you have specified will start taking all decisions for you. As of that point you will have no occasions to revise your plan of action. The transition between parts 2 and 3 can happen midway through a match. Part 3 will last for 15 matches. As with Part 1, at the end of the experiment you will be paid \$0.0045 for every point scored in Parts 2 and 3.

- Are there any questions?

Before we start, let me remind you that just as in part 1, in parts 2 and 3:

- The length of a match is randomly determined. After each round, there is a 50% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

iii. Screen shots from memory 1 elicitation

Decision Screen (Phase 1)

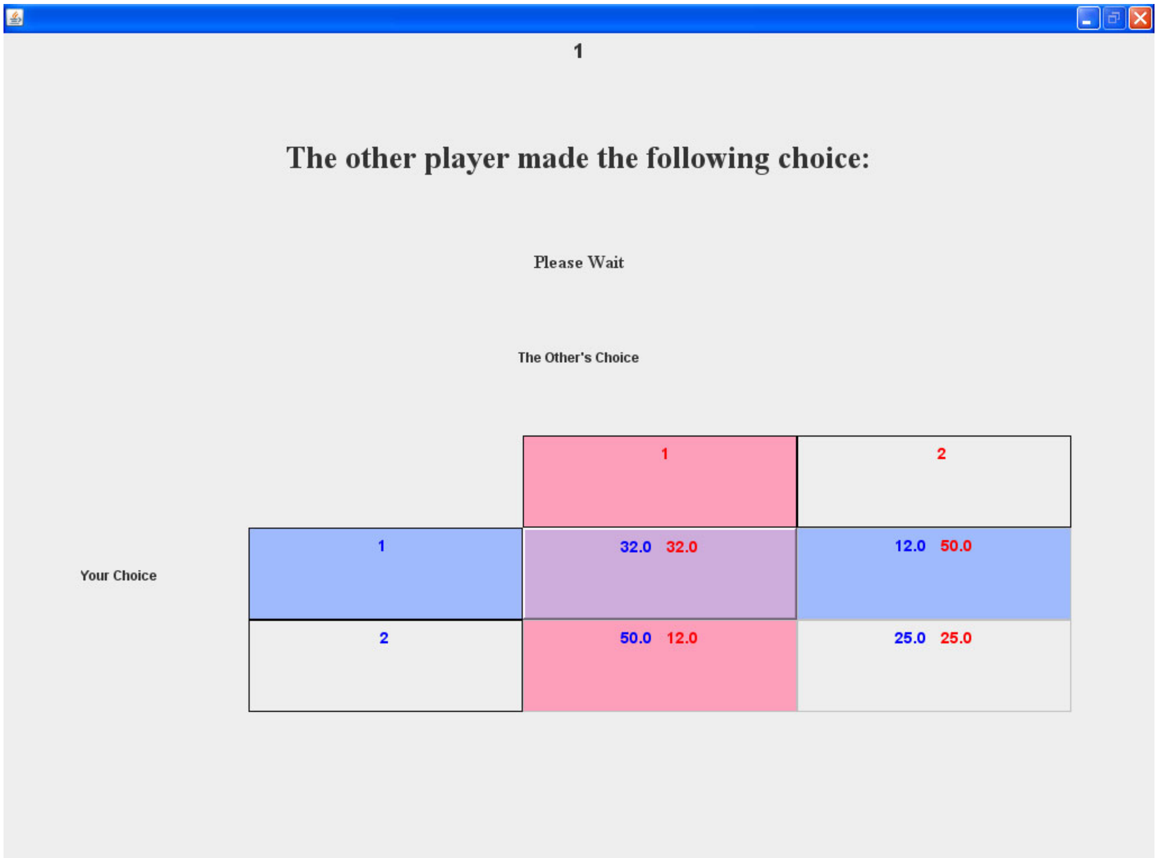

Choices are highlighted after both subjects have made their selection.

Feedback Screen (in between matches, Phase 1)

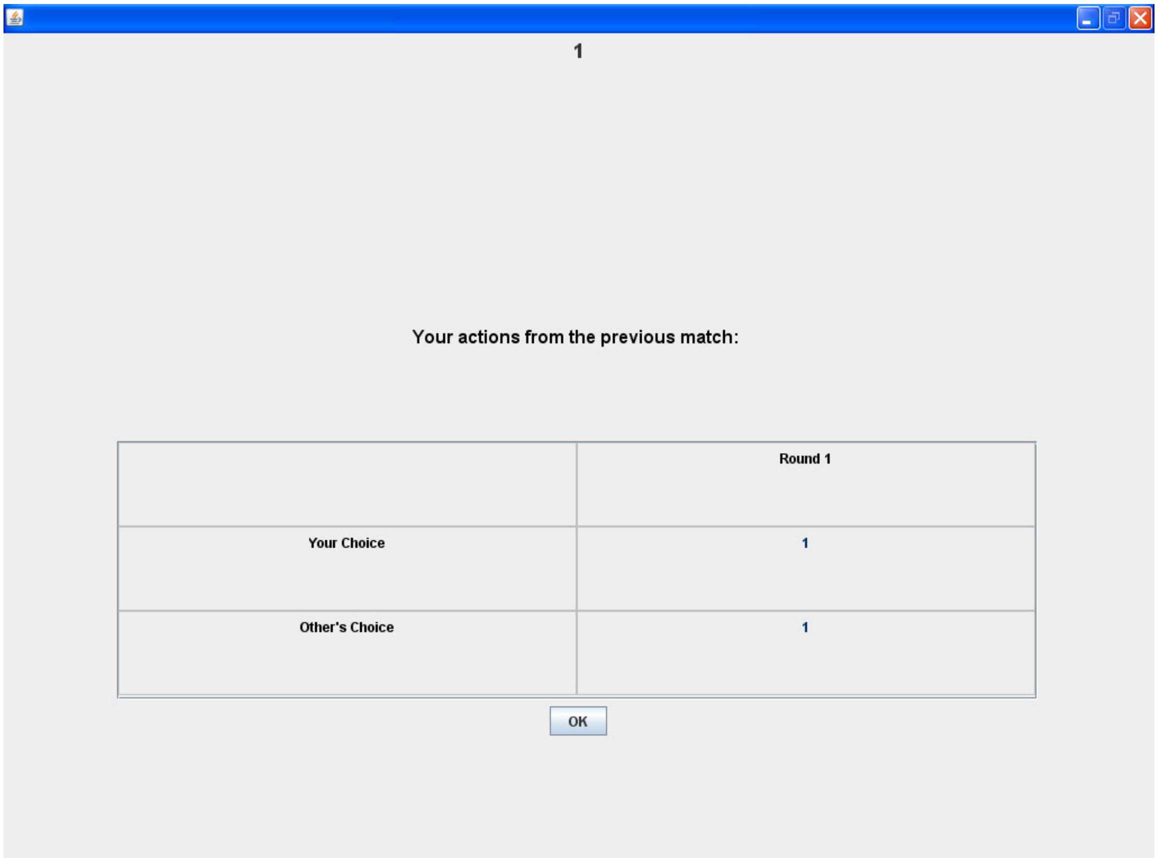

This match only had one round. If there were more, they would all be displayed and a scroll bar would be available if they were too numerous to all fit in one screen.

First Screen of Phase 2

The decision screen in Phase 2 is identical to the one in Phase 1.

Feedback Screen (in between matches of Phase 2)

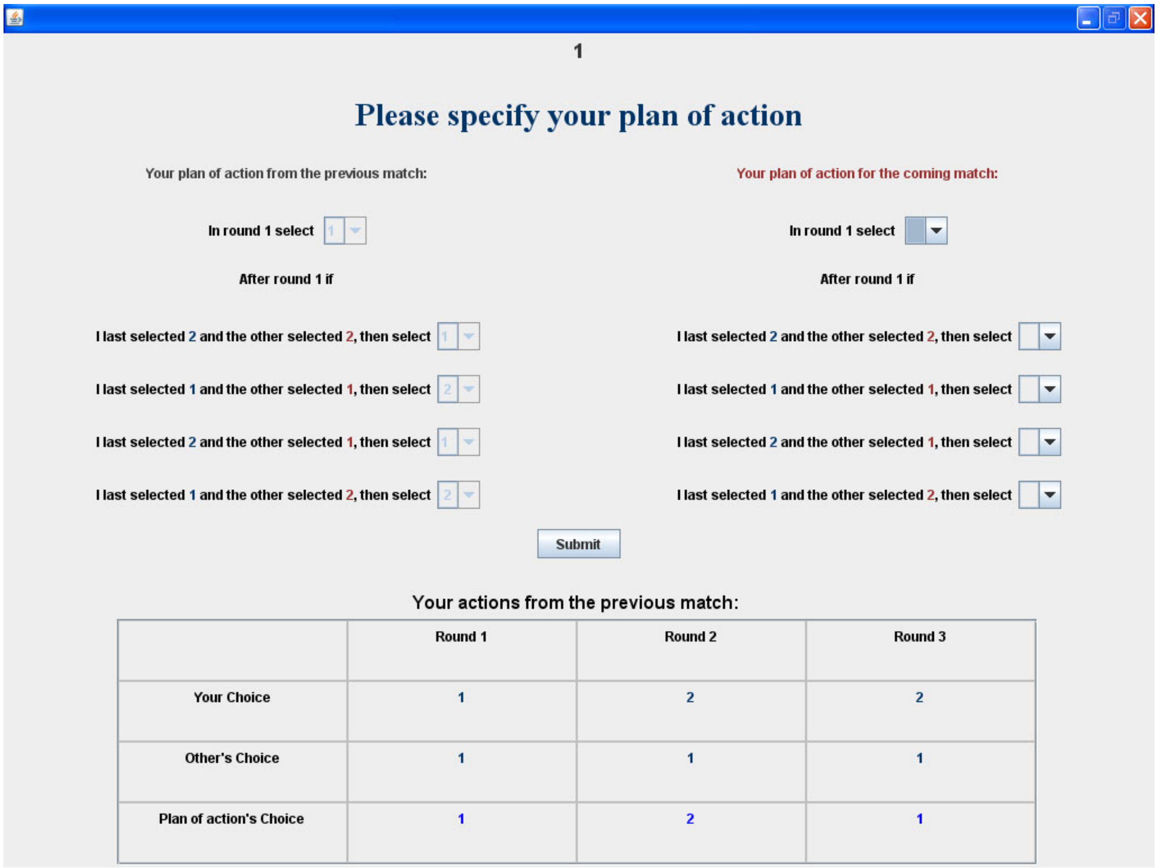

	Round 1	Round 2	Round 3
Your Choice	1	2	2
Other's Choice	1	1	1
Plan of action's Choice	1	2	1

Phase 3 Screen

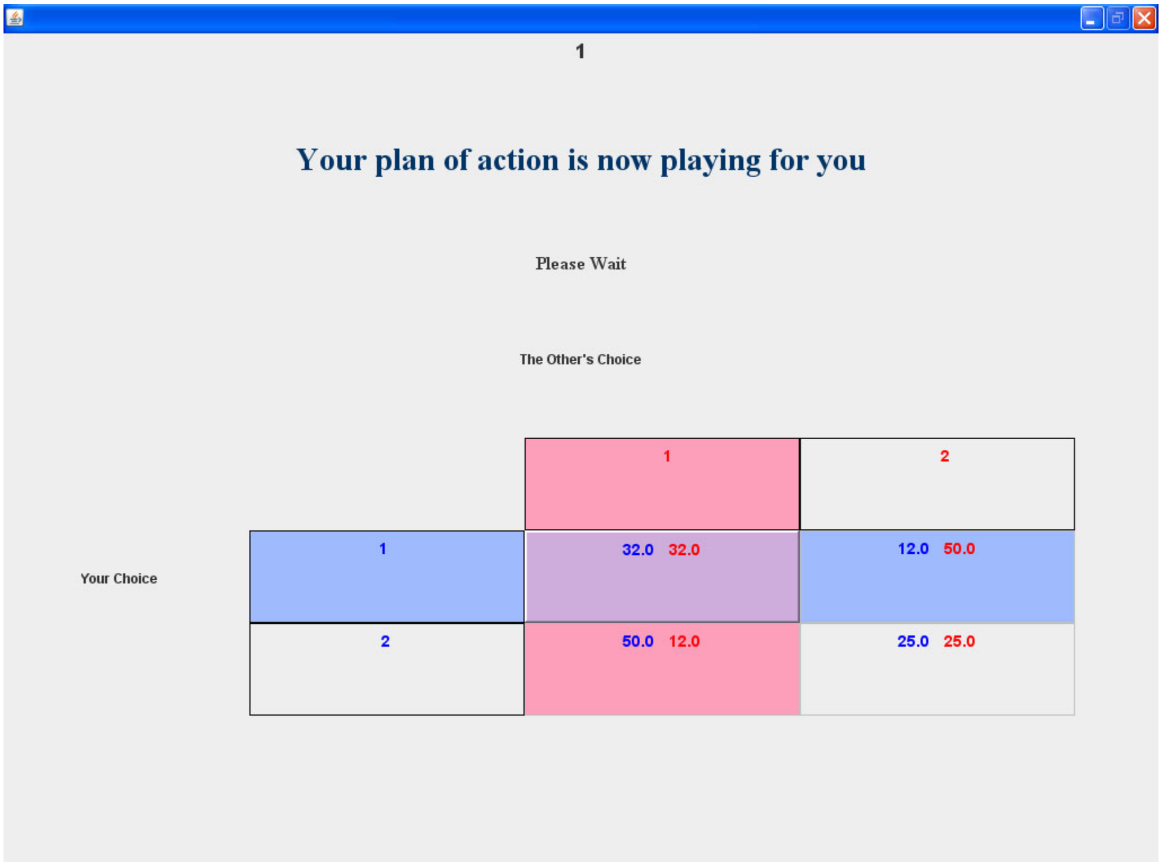

iv. Screen shots from most flexible interface elicitation

Choice Screen With Most Flexible Interface (and Greater Feedback)

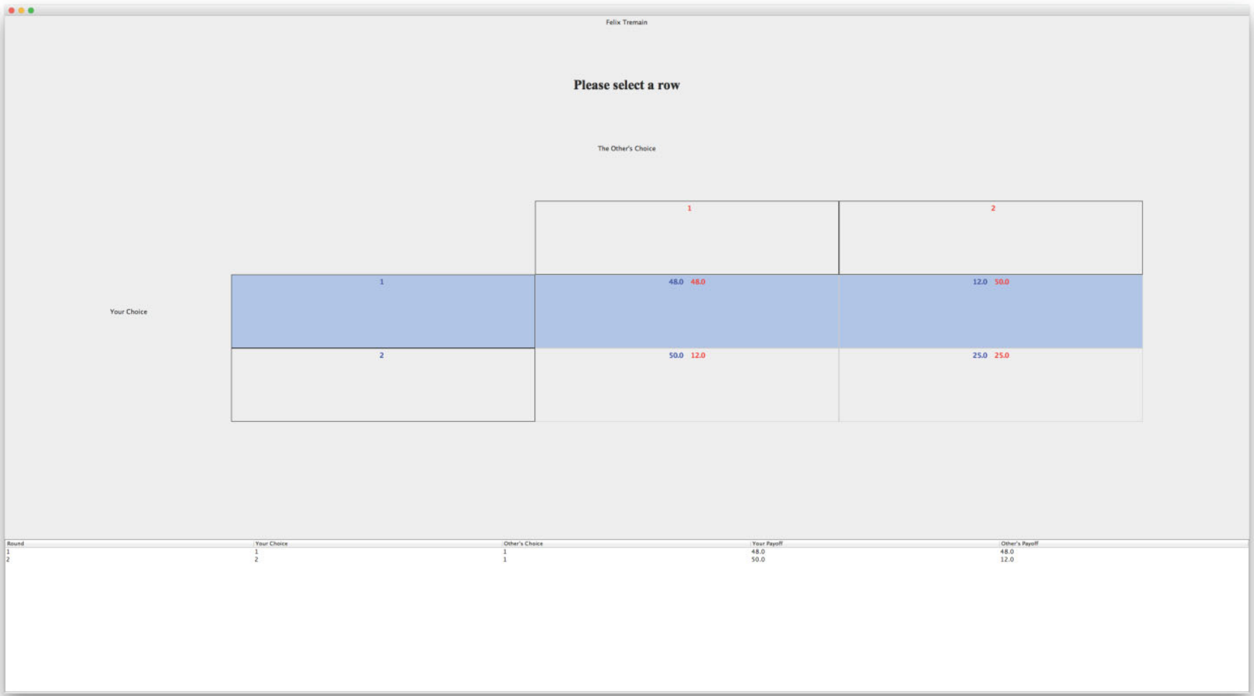

Strategy Choice with Most Flexible Interface

Two examples:

First, this combines the two Screens a user would toggle between when selecting from "Previous Two Round Choices." The top part (menu) is present on all strategy choice screens.

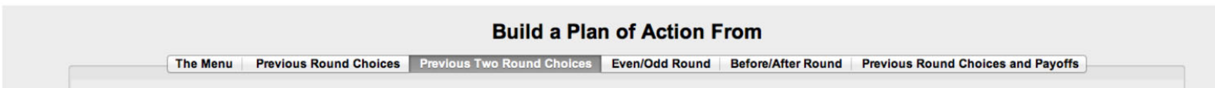

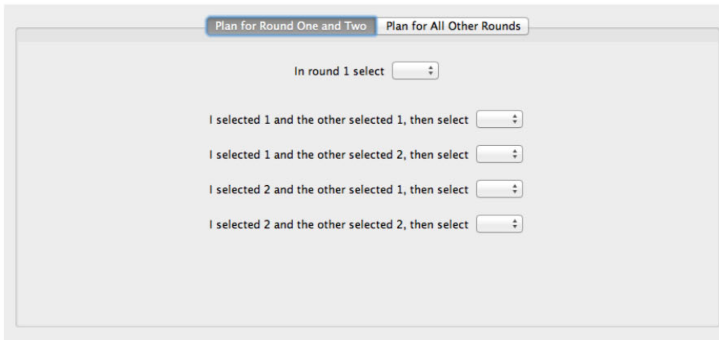

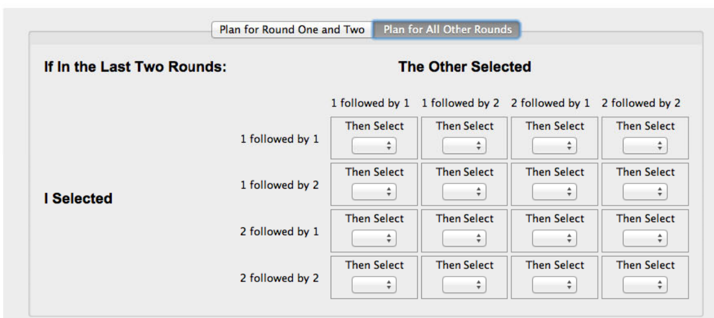

Second, the screen when selecting from "Previous Round Choices and Payoffs."

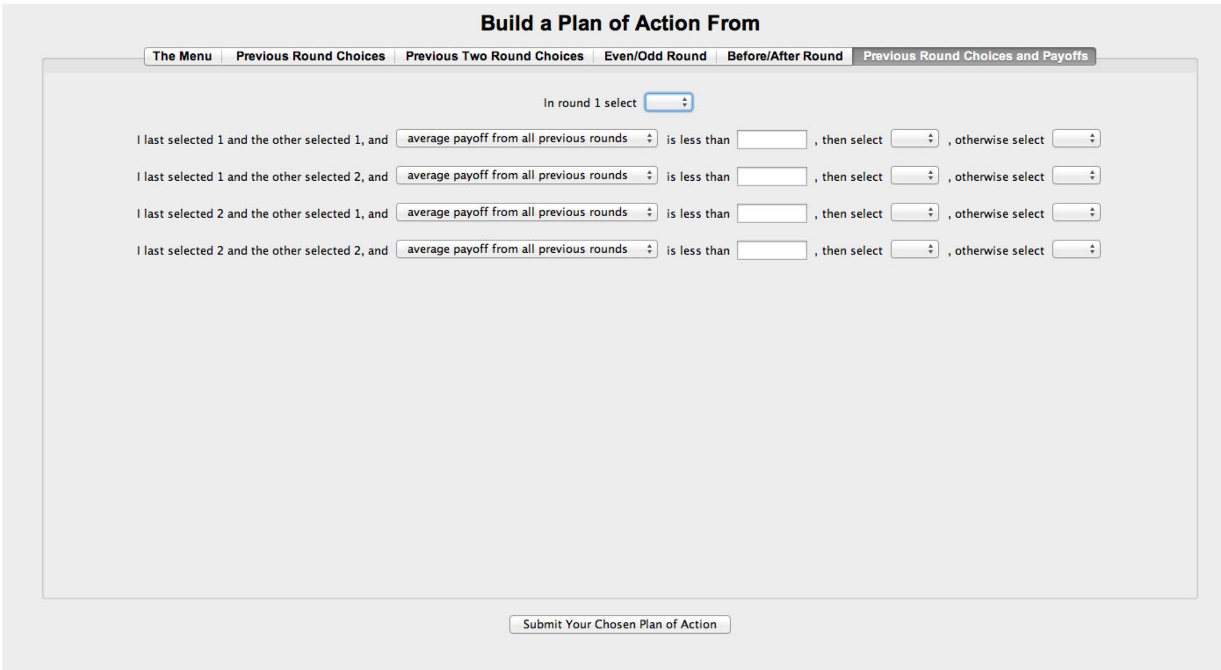

Online Appendix B: Additional Tables and Figures

Table A1: Session Characteristics (Main Treatments in Gray)

Panel A: Memory-One Interface		**δ = ½**		**δ = ¾**		**δ = 9/10**	
Session	**Payoff from cooperation**	**32**	**48**	**32**	**48**	**32**	
	Number of subjects	18	16	14	18	18	
	Number of Games	86	63	50	44	25	
1	**Phase 1**	37	25	16	14	5	
	Phase 2	35	23	20	16	7	
	Phase 3	*14*	15	*14*	*14*	*13*	
	Number of subjects	16	18	16	12	18	
	Number of Games	82	92	60	54	31	
2	**Phase 1**	35	46	25	24	8	
	Phase 2	32	33	21	17	10	
	Phase 3	15	*13*	*14*	*13*	*13*	
	Number of subjects	16	22	14	16	14	
	Number of Games	72	70	43	55	38	
3	**Phase 1**	24	28	17	23	12	
	Phase 2	33	28	11	18	13	
	Phase 3	15	14	15	14	*13*	
Total Number of subjects per treatment		50	56	44	46	50	
Panel B: Menu Interface		**δ = ½**		**δ = ¾**		**δ = 9/10**	**δ = 95/100**
Session	**Payoff from cooperation**		**48**	**32**	**48**	**32**	
	Number of subjects		18	12	16	14	20
	Number of Games		77	52	37	30	26
1	**Phase 1**		37	22	10	12	7
	Phase 2		25	15	13	4	4
	Phase 3		15	15	*14*	*14*	15
	Number of subjects		16	16	14	20	16
	Number of Games		82	45	49	25	23
2	**Phase 1**		37	15	21	6	5
	Phase 2		30	16	14	5	4
	Phase 3		15	*14*	*14*	*14*	*14*
Total Number of subjects per treatment			34	28	30	34	36
Panel C: Expanded Interface						**δ = 9/10**	
Session	**Payoff from cooperation**					**32**	
	Number of subjects					22	
	Number of Games					47	
1	**Phase 1**					18	
	Phase 2					15	
	Phase 3					*14*	
	Number of subjects					16	
	Number of Games					41	
2	**Phase 1**					21	
	Phase 2					6	
	Phase 3					*14*	
Total Number of subjects per treatment						38	

Notes: Italics indicate a phase that started midway through a match.

Table A2: Cooperation Percentage by Phase

	All Supergames			
	Phase 1		Phase 2	
R	First Round	All Rounds	First Round	All Rounds
32	11.86%	9.10%	8.80%	6.93%
48	62.35%	56.15%	67.47%	58.57%
32	24.53%	19.64%	28.37%	21.10%
48	68.21%	55.85%	83.05%	76.71%
32	42.29%	40.86%	61.76%	52.66%

	Last in Phase			
	Phase 1		Phase 2	
R	First Round	All Rounds	First Round	All Rounds
32	4.00%	2.44%	6.00%	5.88%
48	47.78%	35.77%	67.78%	50.72%
32	27.78%	22.10%	26.39%	20.59%
48	67.11%	57.05%	81.58%	84.48%
32	57.14%	51.25%	61.90%	56.47%

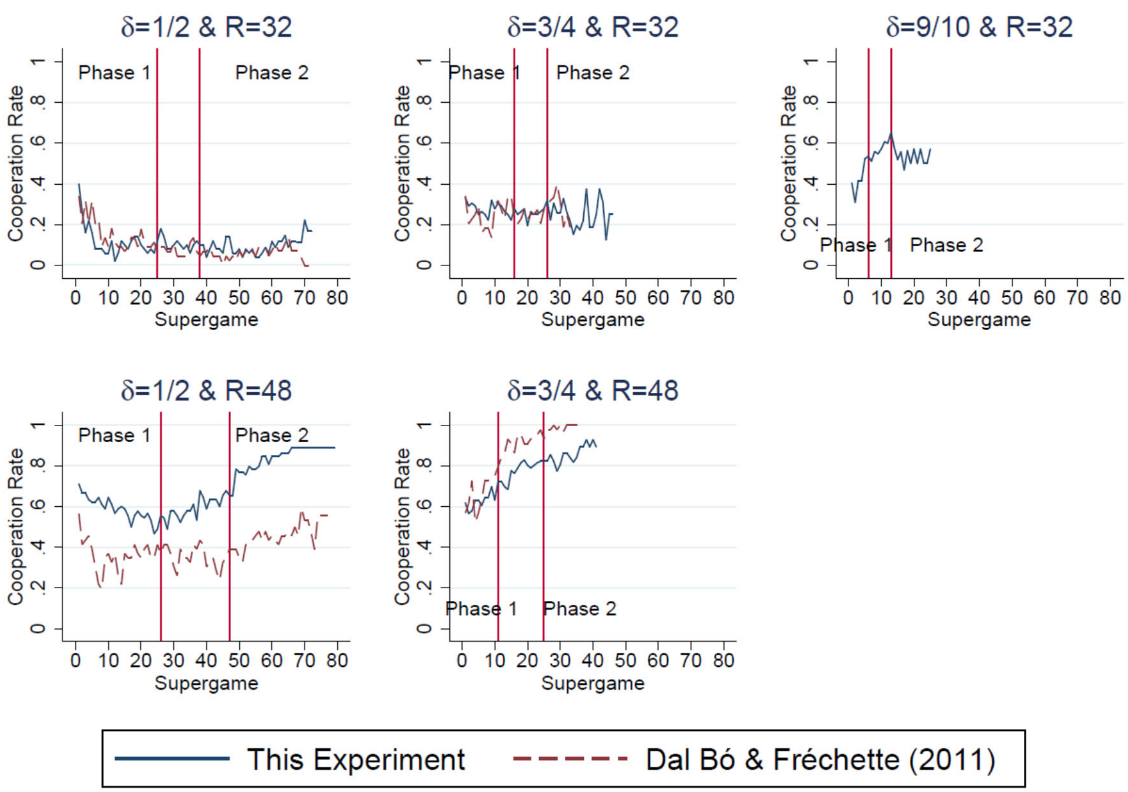

Note: Vertical lines denote the first and last supergame for the start of Phase 2 among sessions for each treatment in this experiment.

Figure A1: Evolution of Cooperation compared with Dal Bó and Fréchette (2011)

Table A3: Does Elicitation Affect Behavior? Regression Results

	δ = ½		δ = ¾	
	R		R	
	32	48	32	48
Panel A: All Rounds				
Elicitation Coefficient	0.0314	0.0391	0.0429	-0.0370
S.E.	(0.0163)	(0.0624)	(0.0238)	(0.0572)
Observations	388	521	259	258
Panel B: Round 1				
Elicitation Coefficient	0.0413*	0.0351	0.0131	-0.017
S.E.	(0.0170)	(0.0656)	(0.0336)	(0.0480)
Observations	388	521	259	258
Panel C: Round 2 after CC history				
Elicitation Coefficient	-0.178	0.0135	-0.0967	-0.0427*
S.E.	(0.2000)	(0.0344)	(0.1010)	(0.0183)
Observations	12	202	71	183
Panel D: Round 2 after CD history				
Elicitation Coefficient	0.0737	0.0819	-0.183	0.0455
S.E.	(0.1400)	(0.0765)	(0.1490)	(0.0573)
Observations	127	245	189	157
Panel E: Round 2 after DC history				
Elicitation Coefficient	0.0102	-0.0458	0.103	0.012
S.E.	(0.0851)	(0.0389)	(0.0951)	(0.1400)
Observations	127	245	189	157
Panel F: Round 2 after DD history				
Elicitation Coefficient	0.00699	0.0117	0.0241	-0.0297
S.E.	(0.0086)	(0.0151)	(0.0352)	(0.0634)
Observations	186	171	186	57

Note: * significantly different from 0 at 10%; ** at 5%; *** at 1%. Results from OLS regression from the specification described in section III.a. The data includes the treatments from these experiments and Dal Bó and Fréchette (2011).

Table A4: Does Elicitation Affect Behavior? Regression Results (Only Data from this Experiment)

	δ = ½		δ = ¾	
	R		R	
	32	48	32	48
Panel A: All Rounds				
Elicitation Coefficient	0.0265	0.0659	0.0668**	0.0999
S.E.	(0.0262)	(0.0643)	(0.0192)	(0.0735)
Observations	193	307	173	165
Panel B: Round 1				
Elicitation Coefficient	0.0300	0.0920	0.0546*	0.0793
S.E.	(0.0232)	(0.0628)	(0.0232)	(0.0658)
Observations	193	307	173	165
Panel C: Round 2 after CC history				
Elicitation Coefficient	-0.594***	-0.0502	-0.0574	-0.0197
S.E.	(8.6e-16)	(0.0393)	(0.119)	(0.0130)
Observations	8	130	48	105
Panel D: Round 2 after CD history				
Elicitation Coefficient	0.0246	0.102	-0.176	0.0600
S.E.	(0.245)	(0.0606)	(0.195)	(0.111)
Observations	60	148	121	101
Panel E: Round 2 after DC history				
Elicitation Coefficient	0.0811	-0.0266	0.0591	0.0284
S.E.	(0.138)	(0.0541)	(0.143)	(0.150)
Observations	60	148	121	101
Panel F: Round 2 after DD history				
Elicitation Coefficient	0.00970	0.0153	0.0227	0.00336
S.E.	(0.00753)	(0.0226)	(0.0565)	(0.0485)
Observations	91	77	118	43

Note: * significantly different from 0 at 10%; ** at 5%; *** at 1%. Results from OLS regression from the specification described in section III.a. The data only includes the treatments from these experiments.

Table A5: Distribution of Elicited Strategies - Memory 1 Elicitation

Strategy*						δ = ½		δ = ¾		δ = 9/10
						R		R		R
					AKA	32	48	32	48	32
C	C	C	C	C	AC		3.57%	2.27%	**8.7%**	2%
C	C	C	C	D	AC'			2.27%		4%
C	C	C	D	C				2.27%	2.17%	2%
C	C	C	D	D	TFT	6%	12.5%	11.36%	32.61%	42%
C	C	D	D	C	WSLS		**1.79%**		**2.17%**	2%
C	C	D	D	D	Grim	6%	**35.71%**	**4.55%**	**39.13%**	**12%**
C	D	C	D	D		2%				
C	D	D	D	C			1.79%			
C	D	D	D	D			1.79%	4.55%		
D	C	C	C	D					2.17%	
D	C	C	D	C				2.27%		
D	C	C	D	D	STFT	12%	5.36%	18.18%		10%
D	C	D	C	C						2%
D	C	D	C	D	AD'					2%
D	C	D	D	C				2.27%		
D	C	D	D	D	AD'	10%	**3.57%**	**9.09%**	**2.17%**	**6%**
D	D	C	D	C					2.17%	2%
D	D	C	D	D		4%				
D	D	D	C	D	AD'	2%				
D	D	D	D	C			1.79%			
D	D	D	D	D	AD	**58%**	**32.14%**	**40.91%**	**8.7%**	**14%**
Number of Observations:						50	56	44	46	50

Note: Only last elicited strategy from Phase 2 is considered for each subject. AC' (AD') denotes that a strategy will behave as AC (AD) in every history it will reach if choices are perfectly implemented. Subgame perfect strategies are denoted in bold, and NE that are not SPE are underlined.

* The letters in the strategy names denote the recommended action after each possible contingency: initial round, CC, DC, CD and DD, where the second letter designates the other's choice.

Table A6: Distribution of Elicited Strategies - Menu Elicitation

		δ = ½	δ = ¾		δ = 9/10
Memory 1 Strategies *		R	R		R
	AKA	48	32	48	32
C C C D D	TFT	2.94%		13.33%	11.76%
C C D D C	WSLS				2.94%
C C D D D	GRIM	**23.53%**	7.14%	**3.33%**	
C D C D D					2.94%
C D D D D		2.94%			
D C C C D					2.94%
D C D D D	AD'		**7.14%**		
D D D D D	AD		**10.71%**		**2.94%**
Total		29.41%	25%	16.67%	23.53%
Menu Strategies					
AD		**14.71%**	**35.71%**	**23.33%**	**14.71%**
CD-1 (CDDDD)			3.57%		
CD-3				3.33%	
DWSLS			3.57%	**3.33%**	
GRIM		**23.53%**	**21.43%**	**26.67%**	**20.59%**
GRIM-1 (GRIM)		**2.94%**		**3.33%**	
GRIM-2				3.33%	5.88%
GRIM-4			3.57%		
STFT					11.76%
TFT		23.53%	7.14%	10%	11.76%
TF1T (TFT)		2.94%			
TF2T				3.33%	2.94%
T1					2.94%
T2		**2.94%**			2.94%
T10					**2.94%**
T12				**3.33%**	
WSLS				**3.33%**	
Total		70.59%	75%	83.33%	76.47%
Number of Observations:		34	28	30	34

Note: Only last elicited strategy from Phase 2 is considered for each subject. Menu strategies in italics denote that they can be generated using the original strategies. AC' (AD') denotes that a strategy will behave as AC (AD) in every history it will reach if choices are perfectly implemented. Subgame perfect strategies are denoted in bold and only NE are underlined,

* The letters in the strategy names denote the recommended action after each possible contingency: initial round, CC, DC, CD and DD, where the second letter designates the other's choice.

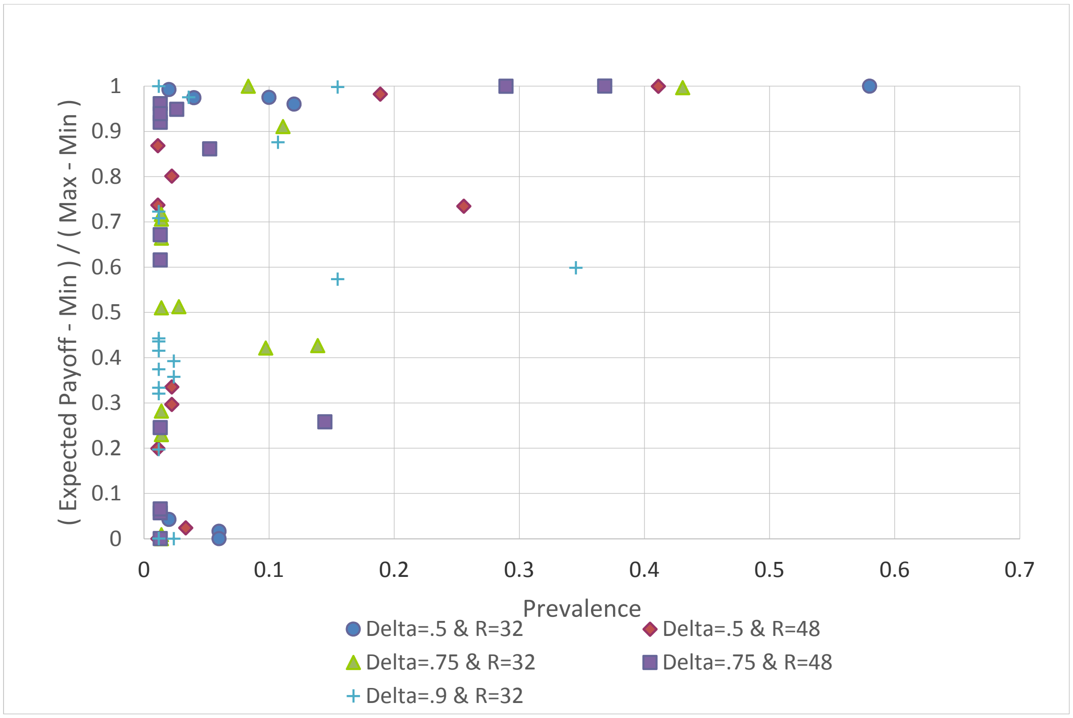

Figure A2: Normalized Payoffs of Strategies When Playing the Population as a Function of Their Popularity

Table A7: Difference between Elicited and Estimated Strategy Prevalence Restricted Samples (Table 4 with Standard Errors)

		AC	TFT	Grim	STF	AD	γ	β	Number of Obs.
δ = ½, R = 32	Elicited			0.08	0.19	0.73			
	Estimation			0.00	0.24	0.76°	0.26	0.98	1563
				(0.00)		(0.11)***	(0.03)***		
δ = ½, R = 48	Elicited	0.04	0.22	0.42	0.04	0.29			
	Estimated	0.06	0.13	0.48	0.06	0.27	0.30	0.96	3247
		(0.03)*	(0.10)	(0.14)***		(0.10)***	(0.04)***		
δ = 3/4, R = 32	Elicited	0.02	0.12	0.13	0.13	0.59			
	Estimated	0.00°	0.13	0.10	0.28	0.48	0.36	0.94	1917
		(0.02)	(0.08)	(0.10)		(0.16)***	(0.04)***		
δ = 3/4, R = 48	Elicited	0.09	0.36	0.41		0.14			
	Estimated	0.17	0.13°°°	0.56°		0.15	0.23	0.99	2320
		(0.06)***		(0.14)***		(0.06)**	(0.04)***		
δ = 9/10, R = 32	Elicited	0.02	0.50	0.22	0.07	0.18			
	Estimated	0.09	0.50	0.11	0.07	0.23	0.32	0.96	2152
		(0.07)	(0.09)***	(0.11)		(0.07)***	(0.04)***		

Bootstrapped standard errors in parenthesis. The parameter γ was used in estimation with $\beta = \frac{1}{1+e^{-\frac{1}{\gamma}}}$.

° Significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.
* Significantly different from 0 at 10%; ** at 5%; *** at 1%.

Table A8: Difference between Elicited and Estimated Strategy Prevalence

Strategy Set from DF 2011 (Table 6 with Standard Errors)

		AC	TFT	WSLS	Grim	STFT	AD	T2	Other	γ	β	Number Of Obs.
δ = ½, R = 32	Elicited	0.00	0.05	0.00	0.06	0.14	0.53	0.00	0.36			
	Estimated	0.00	0.02	0.00	0.02°°°		0.96°°°	0.00		0.31	0.96	2146
		(0.00)	(0.03)	(0.00)	(0.02)		(0.04)***			(0.04)***		
δ = ½, R = 48	Elicited	0.03	0.19	0.03	0.35	0.03	0.25	0.00	0.15			
	Estimated	0.06	0.06°	0.00°°	0.55		0.31	0.02		0.35	0.95	3684
		(0.05)	(0.07)	(0.02)	(0.13)***		(0.13)**			(0.05)***		
δ = 3/4, R = 32	Elicited	0.02	0.10	0.01	0.10	0.10	0.47	0.00	0.30			
	Estimated	0.04	0.12	0.00°°°	0.10		0.74°°	0.00		0.46	0.90	2420
		(0.02)*	(0.08)	(0.02)	(0.09)		(0.09)***			(0.05)***		
δ = 3/4, R = 48	Elicited	0.08	0.30	0.03	0.35	0.00	0.12	0.00	0.14			
	Estimated	0.18°	0.11	0.00	0.57°		0.14	0.00		0.28	0.97	2772
		(0.06)***	(0.12)	(0.01)	(0.13)***		(0.05)***			(0.04)***		
δ = 9/10, R = 32	Elicited	0.02	0.39	0.01	0.17	0.05	0.14	0.00	0.28			
	Estimated	0.08	0.52°°°	0.00°°°	0.16		0.24	0.00		0.37	0.94	2810
		(0.06)	(0.08)***	(0.00)	(0.08)*		(0.06)***			(0.05)***		

Bootstrapped standard errors in parenthesis. The parameter γ was used in estimation with $\beta = \frac{1}{1+e^{-\frac{1}{\gamma}}}$.

° Significantly different from the elicited frequency at 10%; °° at 5%; °°° at 1%.

* Significantly different from 0 at 10%; ** at 5%; *** at 1%.

Table A9: Cooperation Rates before Elicitation Starts in this Paper and Dal Bó and Fréchette (2011) – Round 1

	δ = ½		δ = ¾	
	R		R	
	32	48	32	48
Panel A: All Supergames in Phase 1				
Dal Bó and Fréchette (2011)	0.14	0.37	0.25	0.75
This paper	0.12	0.62	0.25	0.68
p-value of difference	0.748	0.015**	0.925	0.221
Panel B: Last Supergame in Phase 1				
Dal Bó and Fréchette (2011)	0.14	0.39	0.25	0.89
This paper	0.04	0.48	0.28	0.67
p-value of difference	0.104	0.574	0.738	0.027**

Note: As there were not phases in Dal Bó and Fréchette (2011), for comparison we denote as belonging to phase 1 the supergames in the first half of each session.

Table A10: Estimates From DF 2011 vs This Paper

Phase 1 Data

		AC	TFT	WSLS	Grim	STFT	AD	β	Obs.
δ = ½, R = 32	DF 2011	0.36	0.00	0.22	0.02	0.02	0.37	0.94	
		(0.09)	(0.02)		(0.03)	(0.03)	(0.08)		
	This Paper	0.33	0.02	0.28	0.00	0.02	0.35	0.96	4246
		(0.08)	(0.03)		(0.00)	(0.04)	(0.08)		
P-value	of equality	0.83	0.59	0.79	0.47	0.98	0.83		
	of joint equality			0.87					
δ = ½, R = 48	DF 2011	0.19	0.19	0.09	0.12	0.13	0.19	0.85	
		(0.04)	(0.07)		(0.05)	(0.07)	(0.04)		
	This Paper	0.33	0.10	0.15	0.02	0.20	0.00	0.90	6490
		(0.11)	(0.11)		(0.03)	(0.10)	(0.06)		
P-value	of equality	0.24	0.16	0.24	0.06*	0.57	0.01***		
	of joint equality			0.00***					
δ = 3/4, R = 32	DF 2011	0.25	0.08	0.24	0.00	0.16	0.27	0.91	
		(0.09)	(0.06)		(0.01)	(0.09)	(0.08)		
	This Paper	0.51	0.08	0.28	0.01	0.12	0.00	0.94	6044
		(0.18)	(0.08)		(0.03)	(0.09)	(0.11)		
P-value	of equality	0.21	1.00	0.81	0.64	0.75	0.06*		
	of joint equality			0.47					
δ = 3/4, R = 48	DF 2011	0.03	0.54	0.03	0.05	0.31	0.03	0.93	
		(0.02)	(0.17)		(0.08)	(0.11)	(0.02)		
	This Paper	0.19	0.29	0.08	0.05	0.38	0.00	0.91	5822
		(0.07)	(0.09)		(0.04)	(0.08)	(0.03)		
P-value	of equality	0.03**	0.19	0.30	0.97	0.60	0.35		
	of joint equality			0.20					

Note: "p-value of equality" reports the p-value for the test that the fraction of a strategy is the same in the DF 2011 data and in this data. In the case of WSLS, this is obtained as the test that the sum of all the other coefficients is equal across the two samples. "p-value of joint equality" tests that the estimates for AC, TFT, Grim, STFT, and AD, are respectively equal across samples.

Online Appendix C: Options in the sessions with a menu of strategies.

(What is in parentheses was not presented to the subjects):

- Select 1 in every round. (AC)
- Select 2 in every round. (AD)
- Select 1 for X rounds, then select 2 until the end. (CD-X)
- Select 1 X% of the time and 2 1-X% of the time. (RANDOM-X)
- In round 1 select 1. After round 1: if both always selected 1 in the previous rounds, then select 1; otherwise select 2. (GRIM)
- In round 1 select [1 or 2]. After round 1: if the other selected 1 in the previous round, then select 1; if the other selected 2 in the previous round, then select 2. (TFT or STFT)
- In Round 1 select [1 or 2]. After round 1: if both made the same choice (both selected 1 or both selected 2) in the previous round, then select 1; otherwise select 2. (WSLS or D WSLS)
- In round 1 select 1. After round 1: if in X consecutive rounds either the other or myself selected 2, then select 2; otherwise select 1. (GRIM-X)
- In round 1 select 1. After round 1: Select 2 if other selected 2 in all of the previous X [select number] rounds; otherwise select 1. (TFXT)
- Start by selecting 1 and do so until the other or myself selects 2, in that case select 2 for X rounds. After that go back to the start. (T-X)
- Build your own. (This offers the same option as in the memory-1 treatment.)

When [1 or 2] was an option, it was presented as a drop-down menu, and when X needed to be specified, subjects could enter a number in the appropriate box.

Online Appendix D: More Robustness Checks

In this section, we return to focusing on the strategy choices and explore the robustness of our results in a few dimensions: with respect to the average length of a game, with respect to feedback, and by expending the interface to be even more flexible.[43]

i. Even Longer Supergames: δ = 95/100

In this subsection we study whether most subjects choose memory-one strategies even for very high probabilities of continuation. The longer expected length of the supergames could increase the incentives to develop more complicated strategies. We present here results from two sessions with menu elicitation interface but with $\delta = 95/100$ (and R=32). The increased number of rounds per supergame (22.7 rounds per supergame on average) resulted in fewer supergames being played than in the other treatments.[44]

Figure A2 shows the round one cooperation rate for sessions with R = 32 and $\delta = 9/10$ with menu interface and for the new sessions with R = 32 and $\delta = 95/100$. As can be seen, cooperation rates are not different at the beginning of the sessions but a gap emerges as subjects gain experience. We report cooperation rates for this treatment since unlike the other combinations of R and δ that had been studied in Dal Bó and Fréchette (2011), no experiment has been reported on supergames with $\delta = 95/100$. Taking all supergames, the difference in cooperation rates is statistically significant when all rounds are considered (56% with $\delta = 95/100$ versus 42% with with $\delta = 9/10$); but not for round one only (60% with $\delta = 95/100$ versus 51% with $\delta = 9/10$).

[43] We conducted four additional sessions with a total of 74 subjects.
[44] In order for the number of supergames played not to decrease too much, we increased the pace of the third phase, where subjects do not make any decisions themselves.

Figure A2: Evolution of Cooperation in Menu Interface for δ = 9/10 and 95/100 (first rounds)

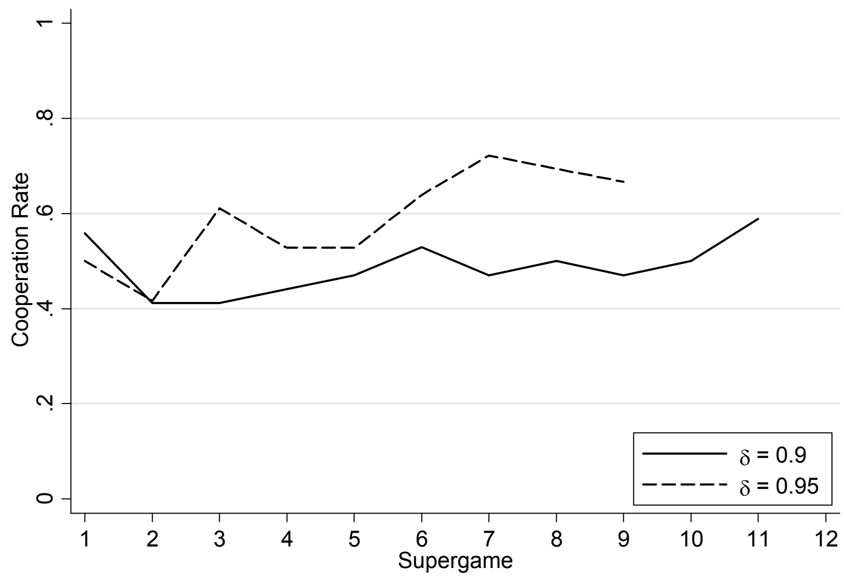

Note: For a given treatment only includes matches that have data in all sessions.

Table A11 reports the strategy choices in the last supergame of Phase 2 for the sessions with δ = 95/100 (for comparison we also report the same information for R=32 and δ = 9/10 from the menu interface sessions). As observed in the sessions with menu interface, the vast majority of strategies (75%) are memory-one strategies for the treatment with δ = 95/100. However, comparing to the treatment with δ = 9/10, where 85% of strategies are memory-one strategies, there is a decrease in the popularity of such strategies but this differences is marginally not statistically significant (p-value = 0.11). The main source of the difference between these two treatments is a decrease in the use of Grim (about 15 percentage points, p-value =0.14) mostly replaced by the use of Grim-2 (an increase of about eight percentage points, p-value = 0.01). The increase of the prevalence of a lenient strategy as δ increases is consistent with the idea that subjects may be less inclined to start long punishments in games which are expected to last for a greater number of rounds.

Table A11: Distribution of Chosen Strategies in Menu Sessions with δ=9/10 and δ=95/100 (last supergame)

	δ = 9/10	δ = 95/100
	R	R
Strategy:	32	32
AC	0.00%	0.00%
TFT	23.53%	25.00%
Grim	20.59%	5.56%
Grim-2	5.88%	13.89%
WSLS	2.94%	11.11%
STFT	11.76%	2.78%
AD	17.65%	19.44%
Other	17.65%	22.22%
Memory 1 or equivalent strategies in Menu Interface	85.29%	75.00%
SPE	41.18%	27.78%
NE	64.71%	52.78%
TFT/(TFT+Grim)	53.33%	81.82%
Observations:	34	36

As can be seen in Figure A3, fixing the returns from mutual cooperation R, increasing δ leads to a decrease in the use of memory-one strategies in all cases. It remains the case, however, that memory-one strategies account for the majority of the data, and that strategies that condition on more than the last rounds do not go back very far. For instance, in the R = 32 and δ = 95/100 treatment, almost 92% of strategies condition on at most the last two rounds and the remaining 8% condition on three rounds in the past. These findings suggest that for a large subset of the possible continuation values, the most important class of strategies are memory-one strategies.

Figure A3: Percentage of Memory-One Strategies in Menu Sessions

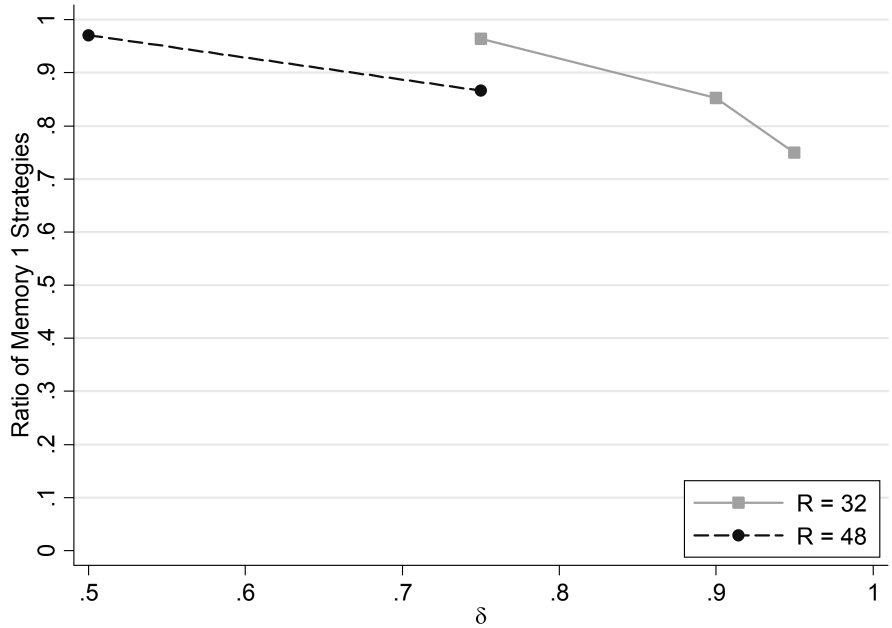

The key observations from Section III are robust to increasing the probability of continuation to δ = 95/100. First, it is still the case that AD, TFT and Grim represent the majority of the strategies chosen by the subjects (see Table A11). Second, a comparison with the treatment with δ = 9/10 shows again that increases in the probability of continuation tend to decrease the prevalence of Grim relative to TFT.[45] The difference with respect to δ = ¾ is significant at the 5% level, but the difference with respect to *δ = 9/10* is not significant at the 10% level. The p-value for the joint test however is 0.002. Without clustering the corresponding p-values are 0.004, 0.14, and 0.02. Third, the popularity of STFT keeps on decreasing with δ, and the difference between *δ = 9/10* and *δ = 95/100* is significant at the 1% level (5% level without clustering). Finally, it is still the case with *δ = 95/100* that many subjects (72%) chose strategies that are not part of a symmetric pure-strategy subgame perfect equilibrium, but a majority of strategies (53%) correspond to symmetric pure-strategy Nash equilibria (see Table A11).

[45] The statistical tests in this case are performed using only the menu sessions to keep the interface comparable across treatments.

ii. Additional Feedback and Even More Flexible Interface

To further study the robustness of the results to the way in which strategies are elicited, in this section we present results from a treatment offering even more flexibility in the elicitation of strategies than was offered in the previous treatments. In addition to the strategies in the memory-one and menu interfaces, this richer interface offered strategies that condition on past payoffs, on whether the rounds are odd or even, and the combination of past choices going back two rounds. Furthermore these conditional statements could be a function of a cutoff round, allowing for different strategies early versus late in a supergame. This is done for the case of $R = 32$ and $\delta = 9/10$. The screen shots of this interface are in the online Appendix A.iv. Given the many changes between the other interfaces and this one, we also shortened the timing between rounds and supergames as we expected that subjects would take longer to make choices. To further *stress test* our procedure, we also changed the within supergame feedback. In the previous treatments, subjects are told after every round what both players did, their payoffs, and they are reminded of the complete history between supergames. In these sessions, subjects see more, namely, all past choices and stage payoffs of both players in the current supergame when making a choice (in every round). We call this the *expanded* elicitation interface.

Table A12 reports the strategy choices in the last supergame of Phase 2 under the column "Expanded," (it also reports the strategy choices under the menu sessions with $R = 32$ and $\delta = 9/10$). The most important result can be found in the last row: Memory-one or Equivalent Strategies (those with memory one or less) are 84.21% of all the chosen strategies. This number was 85.29% in the Menu sessions with $R = 32$ and $\delta = 9/10$ (this difference is not statistically significant – p-value = 0.876). Hence, adding a richer set of alternatives and given more feedback did not change the fact that a vast majority of selected strategies have memory one. Moreover, it is still the case that AD, TFT and Grim represent a majority of the chosen strategies.[46] STFT is about as popular in both

[46] When entering their strategies as a set of conditional statements going back two rounds, subjects had to enter their choice for round one, their strategy for round two (which could condition on at most one round), and their strategy for rounds three and above. The code mistakenly saved the information of the strategy for rounds three and above over the strategy for round two. Hence our definition of strategies for the six subjects who chose to define strategies based on the "two round history" interface is based on their reports

cases. And as before, the fact that many subjects chose TFT shows that some subjects chose strategies that are not part of a subgame perfect equilibrium. However, a majority of subject chose strategies that are part of a Nash equilibrium.

Table A12: Distribution of Chosen Strategies in Menu and Expanded Sessions with δ=9/10 (last supergame)

Strategy:	Menu	Expanded
AC	0.00%	2.63%
TFT	23.53%	21.05%
Grim	20.59%	7.89%
WSLS	2.94%	0.00%
STFT	11.76%	15.79%
AD	17.65%	28.95%
Other	23.53%	23.68%
Memory 1 or equivalent strategies in Menu Interface	85.29%	84.21%
TFT/(TFT+Grim)	53.33%	72.73%
Observations:	34	38

for the first period and periods 3 and above. Four of these six subjects reported a strategy consistent with TFT.

Table A13: Distribution of Elicited Strategies in Menu Sessions with δ = 9/10 and 95/100 (Last Supergame)

Original Strategies					AKA	δ = 9/10 R = 32	δ = 95/100 R = 32
C	C	C	D	C			2.78
C	C	C	D	D	TFT	11.76	8.33
C	C	D	C	D			2.78
C	C	D	D	C	WSLS	2.94	2.78
C	D	C	D	D		2.94	2.78
D	C	C	C	D		2.94	
D	C	C	D	D	STFT		2.78
D	C	D	D	D	AD'		**2.78**
D	D	D	D	D	AD	**2.94**	**2.78**
Total						23.53	27.78

Menu Strategies	δ = 9/10	δ = 95/100
AD	**14.71**	**16.67**
GRIM	**20.59**	**5.56**
GRIM-2	5.88	13.89
GRIM-3		5.56
STFT	11.76	
TFT	11.76	13.89
TF1T (TFT)		2.78
TF2T	2.94	2.78
TF3T		2.78
T1	2.94	
T2	2.94	
T10	**2.94**	
WSLS		8.33
Total	76.47	72.22
Number of Observations:	34	36

Note: Only last elicited strategy from Phase 2 is considered for each subject. Menu strategies in italic denote that they can be generated using the original strategies. AC' (AD') denotes that a strategy will behave as AC (AD) in every history it will reach if choices are perfectly implemented. Sub-game perfect strategies are denoted in bold, and only NE are underlined.

Table A14: Distribution of Elicited Strategies in All Robustness Sessions with δ = 9/10 (Last Supergame)

						Menu	Expanded
Original Strategies					AKA		
C	C	C	D	D	TFT	11.76	10.53
C	C	D	D	D	Grim		2.63
C	D	D	D	D			
C	C	D	D	C	WSLS	2.94	
C	D	C	D	D		2.94	
D	C	C	C	D		2.94	
D	C	C	D	D	STFT		2.63
D	C	D	D	D	AD'		5.26
D	D	C	D	D			2.63
D	D	D	D	D	AD	2.94	
Total						23.53	23.68

Menu Strategies	Menu	Expanded
AC		2.63
AD	14.71	28.95
GRIM	20.59	5.26
GRIM-2	5.88	
GRIM-3		2.63
TFT	11.76	
TF2T	2.94	
STFT	11.76	13.16
T1	2.94	
T2	2.94	
T10	2.94	
Before/After		2.63
Payoff&History		2.63
Two Rounds History		15.79
mem 1 for rounds>2 (=TFT)		10.53
Even AD & Odd Grim		
Odd AD & Even Grim		2.63
Even & Odd Two Rounds History		
AD on Even		
Odd does not condition on even and AD on Even		
Total	76.47	76.32
Memory 1 or Equivalent Strategies	85.29	84.21
Number of Observations:	34	38

Online Appendix E: Simulations

This appendix uses simulations to explore the robustness of the methods used in this paper. The first three sections study the ability of the SFEM to identify the prevalence of strategies. These start with a stripped-down case, where there is a single subject playing against an opponent who randomly chooses and where few strategies are considered. We then move to cases that are similar to the ones considered in this paper in terms of the strategies considered and the sizes of the data sets. This is followed by considering whether SFEM can detect strategies with memory greater than one in sample sizes of the kind considered in this paper. The last two sections study whether the observed consistency of behavior with three simple strategies would still hold if subjects used more complicated strategies.

i. Even with data from only one subject, the SFEM provides consistent estimates when there is enough variation in the history nodes reached

We present here simple simulations that indicate that the SFEM indeed recovers the fraction of play of a strategy in the population in a "best-case" scenario. Namely, we consider the case of subjects playing against random opponents: opponents that cooperate or defect each half of the time. This is a best-case scenario as it guarantees that as the sample becomes large enough, all nodes of the game tree are explored. The results presented in Figure A4 show the estimated fraction of subjects playing Grim (ϕ^{Grim}) as a function of the number of supergames in the sample, for three continuation probabilities: ½, ¾, and 9/10. Each point in a figure is the average estimate from 10,000 simulated samples. The estimation always includes three strategies (the three most popular in our sample): AD, Grim, and TFT. All simulations assume $\beta = 0.9$ in the generation of the data. This parameter was estimated together with the strategy prevalence ones in the estimation stage of the simulations. The top left panel reports the estimates of the prevalence of Grim when the data consists of a single subject who plays Grim (hence $\phi^{Grim} = 1$). The top right panel considers the case of a single subject who plays Grim in half of the supergames and AD in the other half; while the bottom left panel considers two subjects, one playing Grim and the other AD, and the bottom right corner increasers the sample size further by having 20 subjects play Grim and 20 playing AD. Thus, in these three cases $\phi^{Grim} = \frac{1}{2}$. A summary of the simulation design is provided in Table A15.

Table A15: Summary Information For Simulations of Figure A4

	Top Left	Top Right	Bottom Left	Bottom Right
No. Subjects	1	1	2	40
True Parameters	$\beta = 0.9$, $\phi_{AD} = 0$, $\phi_{Grim} = 1$, $\phi_{TFT} = 0$	$\beta = 0.9$, $\phi_{AD} = \frac{1}{2}$, $\phi_{Grim} = \frac{1}{2}$, $\phi_{TFT} = 0$		
Estimated Parameters	$\hat{\beta}$, $\hat{\phi}_{AD}$, $\hat{\phi}_{Grim}$, $\hat{\phi}_{TFT}$			
Opponent	Selects C and D randomly (50% chance of each).			
Supergame Length	Determined randomly according to δ = ½, ¾, and 9/10.			
No. Simulations	10,000 per δ.			

As the figure indicates, when there is substantial variation (an opponent playing randomly in this case), one can get very close to the true ϕ even with data from a single subject. Not surprisingly, as the sample size increase, because the continuation probability is higher, or there are more supergames, or there are more subjects, the estimates approach the true value.

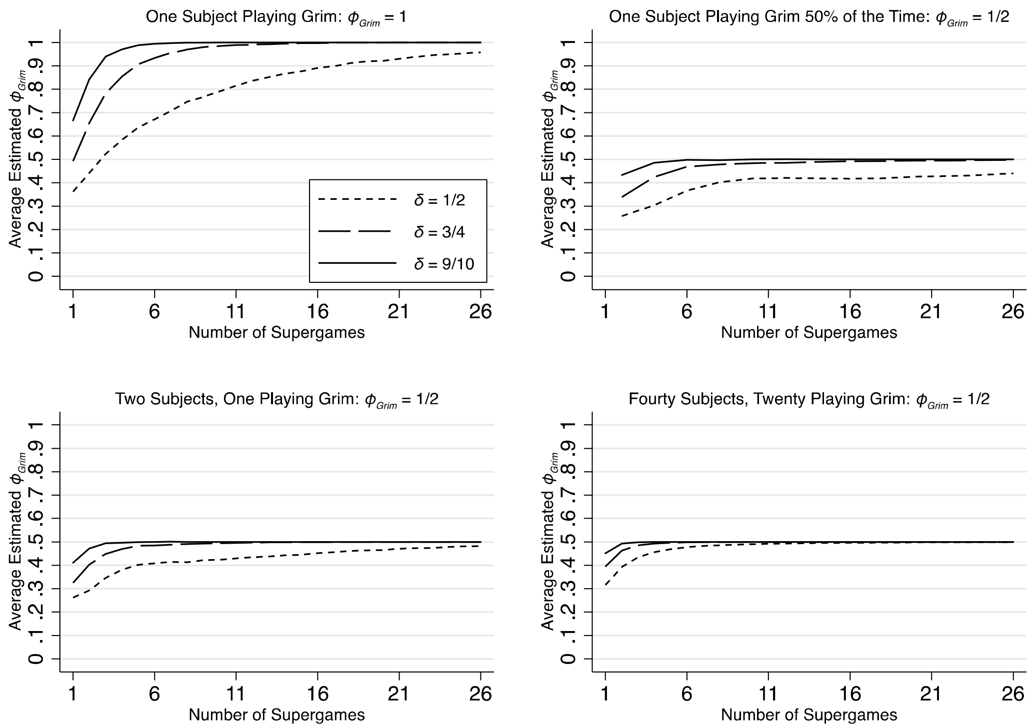

Figure A4: Average Estimate Using Simulated Data Against Random Choices

ii. The SFEM provides precise estimates using the sample size and strategy sets from this paper

In this section we provide simulation results showing that the SFEM's are close to the true value of the parameters for the number of observations and set of strategies considered in this paper. Each simulated experiment consists of a data generating stage and an estimation stage. In the data generating stage we generate a thousand simulated experiments for each treatment. Each simulated experiment has the same number of sessions, (average) subjects per session, and (average) supergames per subject as is in our data set. Subjects are randomly matched and the length of each supergame is randomly determined according to the continuation probability δ. In each simulated experiment, we assume that a subject independently chooses a strategy following the estimated proportions shown in Table 5 (rounded to the second decimal). For example, in treatment (δ=1/2, R=32) a subject chooses TFT, Grim, STFT and AD with probabilities 0.02, 0.02, 0.24 and 0.72 respectively.[47] A subject intends to follow the same strategy in all supergames but may deviate in each round with probability 1-β where β is equal to the estimated value in Table 5. Once the behavior has been generated, we estimate the prevalence of strategies and the mistake probability using the SFEM. The set of strategies we consider is AC, TFT, WSLS, Grim, STFT and AD, which is the set of strategies used in Table 5. A summary of the simulation design is provided in Table A16.

Table A16: Summary Information For Simulations of Tables A12

	$\delta = \frac{1}{2}$ $R = 32$	$\delta = \frac{1}{2}$ $R = 48$	$\delta = \frac{3}{4}$ $R = 32$	$\delta = \frac{3}{4}$ $R = 48$	$\delta = \frac{9}{10}$ $R = 32$
No. Subjects	16	18	12	5	5
No. Sessions	3	5	5	14	14
No. Supergames	23	20	10	10	4
True Parameters	Reported in Tables A12 (Based on Estimated from Table 5).				
Estimated Parameters	$\widehat{\beta}$, $\widehat{\phi}_{AC}$, $\widehat{\phi}_{AD}$, $\widehat{\phi}_{Grim}$, $\widehat{\phi}_{TFT}$, $\widehat{\phi}_{STFT}$, $\widehat{\phi}_{WSLS}$				
Opponent	Randomly matched each supergame to simulated subject within simulated session.				
Supergame Length	Determined randomly according to appropriate δ.				
No. Simulations	1000				

There are several sources of variation in the estimates from simulation to simulation. First, the realized distribution of strategies in one simulation may be different from the expected proportions as each subject independently draws a strategy. Second, the length of the supergames may differ from simulation to simulation as the length is randomly drawn. Third, the matching of subjects may differ from simulation to simulation as it is random. And fourth, conditional on an action, the behavior of subjects may differ from simulation to simulation due to the probability of not following the recommendation of the strategy (1-β).

The simulation results are presented in Table A17. That table shows the strategy prevalence and β parameters that were assumed to generate the data (Truth), and the mean, median and standard deviation of the SFEM estimates. In all five treatments, the

[47] For the treatment (δ=1/2, R=48) the estimates shown in Table 5 do not add up to one due to rounding. Hence, we adjust up the prevalence of Grim.

mean and median estimates are close to the true value of the parameters. the estimates are fairly precise with the average standard deviation around 0.032.

Table A17: Simulation Results – SFEM Estimates Compared with Parameters

	Prevalence of Strategies and Probability of Correct Action						
	AC	TFT	WSLS	Grim	STFT	AD	β
Panel A: δ = ½, R = 32							
Truth	0.000	0.020	0.000	0.020	0.240	0.720	0.970
Mean estimate	0.000	0.020	0.000	0.020	0.240	0.719	0.970
Median estimate	0.000	0.021	0.000	0.021	0.237	0.719	0.970
SD of estimate	0.000	0.021	0.001	0.022	0.080	0.083	0.004
Panel B: δ = ½, R = 48							
Truth	0.070	0.060	0.000	0.560	0.050	0.260	0.950
Mean estimate	0.069	0.058	0.001	0.561	0.050	0.261	0.950
Median estimate	0.067	0.055	0.000	0.562	0.045	0.256	0.950
SD of estimate	0.027	0.036	0.004	0.059	0.024	0.046	0.004
Panel C: δ = 3/4, R = 32							
Truth	0.040	0.090	0.000	0.100	0.290	0.480	0.920
Mean estimate	0.040	0.090	0.000	0.100	0.289	0.481	0.920
Median estimate	0.033	0.085	0.000	0.099	0.286	0.480	0.920
SD of estimate	0.025	0.038	0.001	0.040	0.063	0.066	0.006
Panel D: δ = 3/4, R = 48							
Truth	0.180	0.110	0.000	0.570	0.010	0.130	0.970
Mean estimate	0.178	0.108	0.001	0.573	0.011	0.130	0.970
Median estimate	0.175	0.108	0.000	0.578	0.014	0.129	0.970
SD of estimate	0.048	0.050	0.003	0.066	0.012	0.041	0.003
Panel E: δ = 3/4, R = 48							
Truth	0.080	0.420	0.000	0.180	0.110	0.210	0.940
Mean estimate	0.080	0.421	0.000	0.181	0.108	0.209	0.940
Median estimate	0.075	0.421	0.000	0.181	0.106	0.209	0.940
SD of estimate	0.033	0.061	0.002	0.052	0.037	0.048	0.005

The precision of the estimates can also be seen in Figure A5 which shows the true value of the parameters on the horizontal axis and the 25^{th} percentile, median, and 75^{th} percentile of the estimates on the vertical axis. The median estimate are, in all cases, close to the 45 degree line, showing that they are close to the true value of the parameter. Also note that the 25^{th} and 75^{th} percentiles are also relatively close to the 45 degree line (the difference between these percentiles and the median is around 2 percentage points).

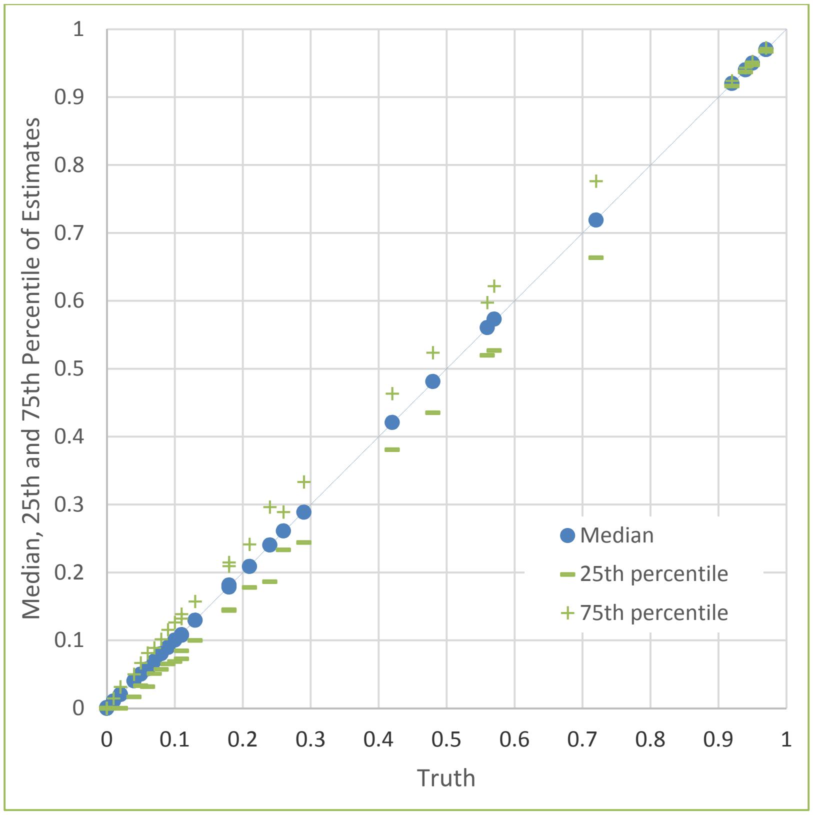

Figure A5: Median, 25th and 75th Percentile of Estimates Compared with True Values

In conclusion, these simulations suggest that the SFEM can recover the prevalence of the different strategies given the number of observations, level of mistakes, and set of strategies considered in this paper.

iii. How much of the data can be explained by strategies with memory more than one?

In this section we study what fraction of round by round cooperation data can be accounted for by strategies with memory one or less versus higher memory. To tackle this question, we need to narrow our focus on a subset of strategies with memory more than one. We focus on the strategies that Fudenberg et al. (2012) identify in statistically significant proportions in randomly terminated PD games with imperfect public monitoring. This set includes several strategies with memory greater than one. For each of our treatments, we estimate the SFEM proportions for this alternative set of strategies to see what fraction of strategies with memory two or more are estimated to fit the data.

The estimation results are presented in Table A18. As can be seen, the fraction of the data estimated to be consistent with higher memory strategies is very low. In two treatments it is close to 20%, in two other treatments it is almost completely absent, and in one it is 7%. Hence, in no treatment, does the majority of the data corresponds to the strategies with memory greater than one.

Table A18: Estimated Strategy Prevalence, Fudenberg et al. (2012) Specification

Full Sample

	δ = ½, R = 32	**δ = ½, R = 48**	**δ = 3/4, R = 32**	**δ = 3/4, R = 48**	**δ = 9/10, R = 32**
AD	0.72***	0.26***	0.48***	0.13***	0.21***
	(0.12)	(0.09)	(0.09)	(0.05)	(0.06)
AC	0.00	0.07	0.01	0.12*	0.00
	(0.00)	(0.04)	(0.02)	(0.06)	(0.01)
Grim	0.02	0.55***	0.10	0.46***	0.09
	(0.02)	(0.13)	(0.10)	(0.14)	(0.07)
TFT	0.02	0.00	0.09**	0.09	0.38***
	(0.03)	(0.02)	(0.04)	(0.09)	(0.10)
STFT	0.24**	0.05	0.29***	0.01	0.10*
	(0.10)	(0.05)	(0.10)	(0.02)	(0.06)
TF2T	0.00	0.00	0.03	0.03	0.06
	(0.00)	(0.01)	(0.03)	(0.05)	(0.07)
TF3T	0.00	0.00	0.00	0.03	0.04
	(0.00)	(0.02)	(0.01)	(0.04)	(0.04)
2TFT	0.00	0.07	0.00	0.00	0.00
	(0.00)	(0.08)	(0.03)	(0.06)	(0.04)
2TF2T	0.00	0.00	0.00	0.03	0.03
	(0.00)	(0.02)	(0.01)	(0.05)	(0.04)
Grim2	0.00	0.00	0.00	0.09	0.08
	(0.00)	(0.04)	(0.01)	(0.06)	(0.06)
Grim3	0.00	0.00	0.00	0.00	0.00
γ	0.30***	0.35***	0.39***	0.26***	0.35***
	(0.03)	(0.05)	(0.04)	(0.04)	(0.04)
Memory ≤ 1	**1.00**	**0.93**	**0.97**	**0.82**	**0.79**
Memory > 1	**0.00**	**0.07**	**0.03**	**0.18**	**0.21**

Estimated frequency significantly different from zero: * at 10%; ** at 5%; *** at 1%.

One potential concern is that, in our simple environment, the SFEM might have difficulty distinguishing play that comes from higher memory strategies from choices of simple strategies. To investigate this possibility, we study the ability of the SFEM to estimate the proportion of memory-one strategies by simulating experiments as in the previous section. The main difference is that here we focus on two treatments: (δ=1/2, R=48) and (δ=9/10, R=32), and that we assume that the true prevalence of strategies includes strategies with memory greater than one.[48]

As in the previous section, each simulated experiment consists of a data generating stage and an estimation stage. In the data generating stage we generate a thousand simulated experiments for each treatment. Each simulated experiment has the same number of sessions, (average) subjects per session, and (average) supergames per subject as is in our data set for the treatments (δ=1/2, R=48) and (δ=9/10, R=32). Subjects are randomly matched and the length of each supergame is randomly determined according to the continuation probability δ. In each simulated experiment, we assume that a subject independently chooses a strategy following the estimated proportions in Fudenberg et al. (2012) for their treatments b/c = 2, b/c = 2.5, and b/c = 4 (thus three distinct simulations for each case considered).[49] A subject intends to follow the same strategy in all supergames but it may deviate in each round with probability 1-β where β is equal to the corresponding estimated values in Fudenberg et al. (2012). Once the behavior has been generated, we estimate the prevalence of strategies and the mistake probability using SFEM for each simulated data set. The set of strategies we consider is, again, as in in Fudenberg et al. (2012): AD, AC, Grim, TFT, STFT, TF2T, TF3T, 2TFT, 2TF2T, Grim2, Grim3. A summary of the simulation design is provided in Table A19.

---

[48] We focus on these because they cover the two extremes that one may think could lead to difficulties. When δ=1/2 one might worry that, because the average supergame is short, it could be difficult to identify strategies; on the other hand there are many supergames per subjects and thus some of them will be longer. When δ=9/10 supergames are on average longer, but there are only a few of them per subject.
[49] In that paper, they estimate that only a minority of strategies has memory one and, hence, the prevalence of AD, Grim and TFT is much smaller than in our data. We focus on three of their four main treatments: b/c = 2, b/c = 2.5, and b/c = 4. In the fourth one cooperation cannot be supported in equilibrium. The relevant results can be found in their Table 4. Since they have three treatments that can support cooperation, and each of these is estimated to have a different distribution of strategies, we will perform three simulations for each of our two δ and R combinations. In each of the three distributions, they report only between 33% and 40% of the data is a memory-one strategy (AD, AC, TFT, Grim, or STFT) and the other strategies are: TF2T, TF3T, 2TFT, 2TF2T, Grim2, and Grim3, which accounts for 60% to 67% of their data.

Table A19: Summary Information For Simulations of Figures A4 and A5

	Figure A6			Figure A7		
	b/c = 2	b/c = 2.5	b/c = 4	b/c = 2	b/c = 2.5	b/c = 4
No. Subjects	18	18	18	14	14	14
No. Sessions	5	5	5	5	5	5
No. Supergames	20	20	20	4	4	4
True Parameters	Based on Estimates from Table 4 of Fudenberg et al. (2012).					
Estimated Parameters	$\widehat{\beta}$, $\hat{\phi}_{AC}$, $\hat{\phi}_{AD}$, $\hat{\phi}_{Grim}$, $\hat{\phi}_{TFT}$, $\hat{\phi}_{STFT}$, $\hat{\phi}_{TF2T}$, $\hat{\phi}_{TF3T}$, $\hat{\phi}_{2TFT}$, $\hat{\phi}_{2TF2T}$, $\hat{\phi}_{Grim2}$, $\hat{\phi}_{Grim3}$					
Opponent	Randomly matched each supergame to simulated subject within simulated session.					
Supergame Length	Determined randomly according to $\delta = \frac{1}{2}$.			Determined randomly according to $\delta = \frac{9}{10}$.		
No. Simulations	1000					

Figures A4 and A5 below report the sum of the estimates for strategies with memory greater than one. More specifically, it draws the kernel density estimate of that sum. The figures also indicate via vertical lines the actual sum in the data generating process (i.e. the estimates of Fudenberg et al. 2012) and the sum of the estimates from our data set (the dashed vertical line). The figures clearly show that the SFEM recovers higher frequencies of strategies with memory two or above than found in our data set. In fact, for the two treatments we perform the simulations in, there is not a single simulation (out of the 6000 simulated experiments) where the estimates of strategies with memory two and above are as low as in the estimates for our data. Note also that the estimates correspond reasonably well to the data generating process. These results show that our finding that most subjects use simple strategies cannot be due to a systematic bias arising from the SFEM.

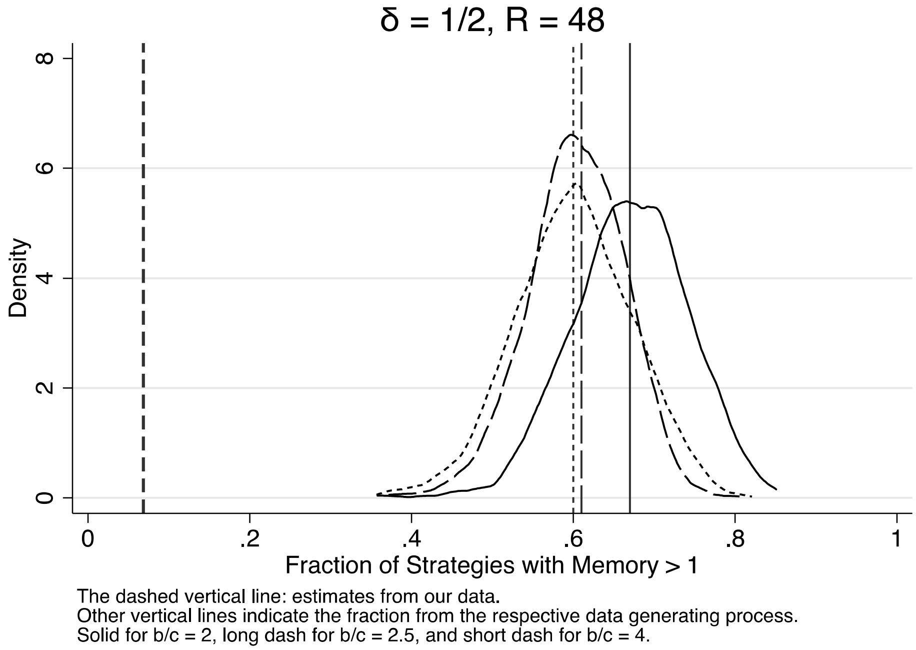

Figure A6: SFEM Estimates of Memory Greater than One if Population Follows Proportions in Fudenberg et al. (2012)

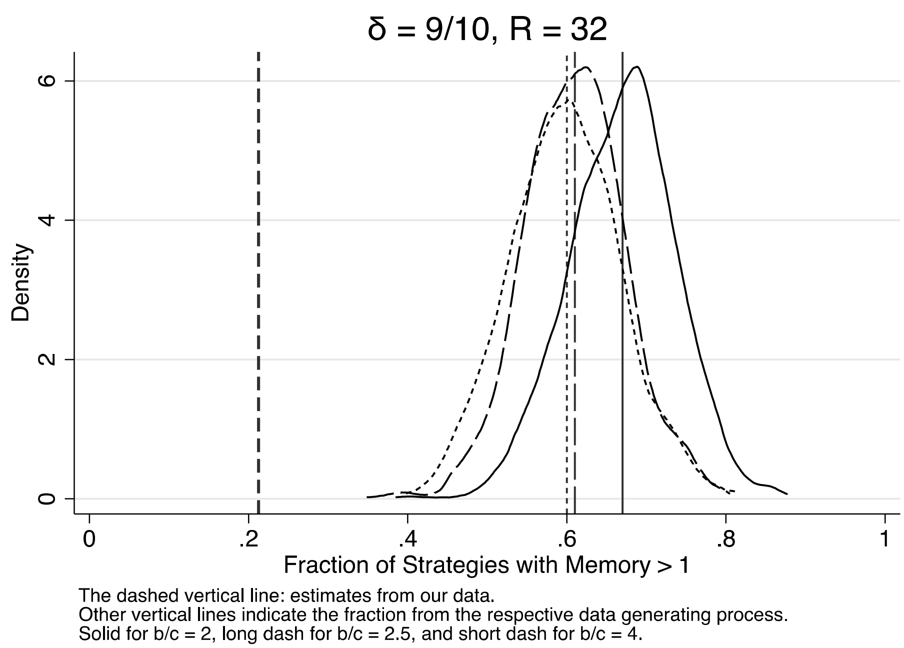

Figure A7: SFEM Estimates of Memory Greater than One if Population Follows Proportions in Fudenberg et al. (2012)

iv. A majority of behavior is consistent with AD, Grim, or TFT and simulations show results would be different if subjects used strategies with longer memory

Both elicitation and SFEM results show that, in the environments studied in this paper, a majority of subjects chose one of three strategies: AD, Grim or TFT. Support for the idea that this results is not due to limitations of the elicitation or SFEM can be found in the fact that, in all of our treatments, the vast majority of choices can be perfectly accounted for by those three strategies. That is, the entire sequence of choices made by a subject in a supergame is typically consistent with those three strategies, as shown in Table A20.

Table A20: Percentage of Behavior Consistent with AD, Grim or TFT in Supergames with More Than One Round

δ	R	
	32	48
1/2	95%	92%
3/4	80%	91%
9/10	75%	

But is it the case that those three strategies would be consistent with a majority of behavior even if subjects used other, possibly more complex strategies? Is it possible that the simplicity of our environment leads mechanically to this result?

Table A21: Percentage of Behavior Consistent with AD, Grim or TFT in Supergames with More Than One Round

		Treatment	
		δ=1/2, R=48	**δ=9/10, R=32**
Our Experimental Data		**92%**	**75%**
Simulated Data Based On:	b/c = 2	65%	30%
	b/c = 2.5	66%	31%
	b/c = 4	67%	35%
Fraction of Experiments with Fit ≥ Experimental Data	b/c = 2	0%	0%
	b/c = 2.5	0%	0%
	b/c = 4	0%	0%

To address this question, we simulate behavior arising from more complicated strategies and repeat the exercise above. Clearly, one could consider very different alternative strategies in these simulations. One approach would be to randomly select strategies of memory more than one, and less than some number. However, this would often select strategies that seem strange, or unlikely. Instead, we will use results from the strategy estimation of previous experiments that were done in an environment where subjects have been shown to use more complex strategies. More specifically, as in Appendix E.iii. we will use results from Fudenberg et al. (2012) which studies randomly terminated PD with imperfect public monitoring (see footnote 38 for more details).

Using these simulated experiments, we study the percentage of simulated behavior which is consistent with AD, Grim, or TFT and whether this differs from what we find in our experiments. We generate 1000 simulated experiments of the same size as the one we conducted for each of the three possible distributions of strategies and for each of the two treatments.[50] All the numbers provided are for supergames with more than one round (since in supergames with only one round choices are always consistent with these strategies).

Table A21 makes it clear that the degree with which choices in our experiment line up with AD, Grim and TFT would not happen if the data was generated with the more complex strategies identified in Fudenberg et al. (2012). In fact, out of the 6,000 simulated experiments reported above, not a single one generates sequences of choices consistent with those three strategies in rates comparable to what is found in our data.

One criticism of the exercise performed above is that, although the distribution of strategies may be sensible, the degree of noise (β) could exaggerate the level of randomness relevant for our current setting. One could make the case that in a more complicated game, such as one with imperfect public monitoring, it is likely that subjects exhibit more randomness. Hence, we repeat the exercise above but this time using the β reported in our paper for the relevant treatments, which are indeed higher than those of Fudenberg et al. (2012).[51]

[50] See Table A17 for a summary of the structure of the simulations.
[51] These can be found in Table 5 were we report our estimates when using the entire data set.

Table A22: Percentage of Behavior Consistent with AD, Grim or TFT in Supergames with More Than One Round (alternative β)

		Treatment	
		δ=1/2, R=48	**δ=9/10, R=32**
Our Experimental Data		**92%**	**75%**
Simulated Data Based On:	b/c = 2	77%	45%
	b/c = 2.5	77%	46%
	b/c = 4	74%	43%
Fraction of Experiments with Fit ≥ Experimental Data	b/c = 2	0%	0%
	b/c = 2.5	0%	0%
	b/c = 4	0%	0%

These results are reported in Table A22. As expected there is an increase in the frequency with which choices are now consistent with one of the simple strategies. However, the performance is still far from what is observed in our data.

Hence, although the environment under study is simple, the results of these simulations suggest that behavior would be noticeably different if a large fraction of subjects used more complicated strategies.