

Strategy Choice in The Infinitely Repeated Prisoners' Dilemma

BY PEDRO DAL BÓ AND GUILLAUME R. FRÉCHETTE*

We use a novel experimental design to reliably elicit subjects' strategies in an infinitely repeated prisoners' dilemma experiment with perfect monitoring. We find that three simple strategies represent the majority of the chosen strategies: Always Defect, Tit-For-Tat, and Grim. In addition, we identify how the strategies systematically vary with the parameters of the game. Finally, we use the elicited strategies to test the ability to recover strategies using statistical methods based on observed round by round cooperation choices and find that this can be done fairly well, but only under certain conditions.

* Dal Bó: Department of Economics, Brown University, Providence, RI 02912 (e-mail: pdalbo@brown.edu); Fréchette: Department of Economics, New York University, New York, NY 10012 (e-mail: frechette@nyu.edu). Jeffrey Ely was the coeditor for this article. We are grateful to Jim Andreoni, Jeffrey Ely, Drew Fudenberg, Asen Ivanov, Thomas Palfrey, Philippe Jehiel, anonymous referees, and seminar participants at UCSD, Princeton, GMU, CMU, UQAM, Tilberg, Oxford, Stanford, Paris I, Purdue, Washington St. Louis, Bologna, Harvard, Toronto, Columbia, Boston University, Boston College, HBS, Maastricht University, Bocconi University, CIDE, UBC, Sciences Po, Carlos III, CEMFI, University of Hawaii, UCSB, Chapman, Concordia University, Caltech, and at the ESA conference in Tucson, the Conference on Social Dilemmas, the SAET conference, the Conference on New Directions in Applied Microeconomics, and the Mont Tremblant conference for very useful comments. We thank Emanuel Vespa, Sevgi Yuksel, and Tianzan Pang for their research assistance, CASSEL (UCLA) and SSEL (Caltech) for the Multistage software, as well as Rajeev Advani and Anwar Ruff for its adaptation. Fréchette gratefully acknowledges the support of the NSF via grants SES-0924780, SES-1225779, and SES-1558857 as well as support from the Center for Experimental Social Science (CESS), and the C.V. Starr Center. Dal Bó gratefully acknowledges the support of NSF via grant SES-0720753. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NSF. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

I. Introduction

The theory of infinitely repeated games has been an active area of research in recent decades and is central to many applications.¹ A key insight from this literature is that repeated interactions may allow people to overcome opportunistic behavior. While a series of experiments has supported this theory, less is known about the types of strategies people actually use to overcome opportunistic behavior.²

Learning which strategies people actually use is of interest for many reasons. First, it can help future theoretical work to identify refinements or conditions that lead to these strategies being played. Rubinstein (1998, p.142) argues that folk theorems are statements about payoffs and that more attention could be paid to the strategies supporting those payoffs. Identifying the set of strategies used to support cooperation can provide a tighter test of theory than efforts based on the study of outcomes. For example, it allows us to test whether the strategies used coincide with the ones that theory predicts should be used (i.e., are the strategies used part of a subgame perfect equilibrium?). Second, it can help theorists focus on the most relevant set of strategies. For example, an influential literature in biology considers which strategies are likely to survive evolution. Given the complexities of working with the infinite set of possible strategies, researchers focus on finite subsets of strategies that include those usually studied by theorists (e.g., Imhof, Fudenberg and Nowak 2007). If we can identify the strategies that people use, we can better define the finite set of strategies that should be considered. Third, it can also help identify those environments in which cooperation is more likely to emerge. That is, the results can be used to modify the theoretical conditions needed for cooperation.

Previous papers have estimated the use of strategies from the observed realization of behavior, but there are serious hurdles for identification. First, the set of possible strategies is infinite (uncountable). Second, while a strategy must specify an action after each possible history, for each repeated game, we observe only one actual finite history and not what subjects would have done under other histories. Two different approaches have been used to overcome these hurdles. Both

¹ Within economics, repeated games have been applied to many areas: industrial organization (see Friedman 1971, Green and Porter 1984 and Rotemberg and Saloner 1986), informal contracts (Klein and Leffler 1981), theory of the firm (Baker et al. 2002), public finance (Phelan and Stacchetti 2001) and macroeconomics (Rotemberg and Saloner 1986 and Rotemberg and Woodford 1990) to name just a few.

² Roth and Murnighan (1978) and Murnighan and Roth (1983) were the first papers to induce infinitely repeated games in the lab by considering a random continuation rule. The probability with which the game continues for an additional round induces the discount factor. A large experimental literature now exists on infinitely repeated games and is reviewed in Dal Bó and Fréchette (2018).

methods start by specifying a family of strategies to be considered, but they differ in how, among these, the best-fitting strategies are selected. One approach trades off goodness-of-fit of a set of strategies versus a cost of adding more strategies (see Engle-Warnick and Slonim 2004 and 2006; and Camera, Casari, and Bigoni 2012; for a related but Bayesian approach to this, see Engle-Warnick, McCausland, and Miller 2004). A second approach uses maximum likelihood estimation either to estimate the prevalence of each strategy in the set under consideration (see Dal Bó and Fréchette 2011; Fudenberg, Rand, and Dreber 2012; Camera, Casari, and Bigoni 2012; Fréchette and Yuksel 2013; Jones 2014; Breitmoser 2015; Rand, Fudenberg, and Dreber 2015; Vespa 2015; Vespa and Wilson 2015; Arechar, Dreber, Fudenberg; and Rand 2016, and others) or to estimate the best-fitting strategy while allowing for subject-specific heterogeneity in the transitions across states of the strategy (Aoyagi and Fréchette 2009). In this paper, we propose an alternative approach to studying strategies: the elicitation of strategies (i.e., a modified strategy method, Selten 1967 – at the end of the introduction we discuss related work on elicitation of strategies in repeated games). We ask subjects to design strategies that will play in their place. A major challenge to the use of the strategy method is that it can affect behavior.³ We introduce our modified strategy method in ways that limit the chances that the procedure will affect behavior.

Note that a distinct advantage of strategy elicitation is that it allows us to consider many more strategies than the previously mentioned estimation methods. Consider the following example: a repeated game where the researcher thinks it is likely that subjects would use trigger strategies with a finite number of periods of punishments. For each possible punishment length, estimation requires the addition of one strategy to the set, and this could result in a very large (and difficult to identify) set of possibilities. When using elicitation, this can be included as a single strategy “type,” where subjects simply determine one parameter, namely the number of periods of punishment. This makes elicitation potentially particularly attractive to study environments that are not well understood.

The main findings of the current study fall in two categories: new results about the strategies subjects use in infinitely repeated prisoner’s dilemma games with perfect monitoring and an assessment of the ability to recover strategies from choices econometrically using experimental (as opposed to simulated) data.

³ Hoffman, McCabe, and Smith (1998), Gueth et al. (2001), and Brosig et al. (2003) report evidence that the strategy method may affect behavior. Brandts and Charness (2000) find that it does not. See Brandts and Charness (2011) for a review of the literature.

With respect to strategies, we find that a majority of subjects choose simple strategies that only condition on the outcome from the previous round when allowed complex strategies. In fact, three simple strategies represent the majority of the chosen strategies: Always Defect (AD), Tit-For-Tat (TFT), and the Grim trigger strategy (Grim). This result holds for both short and long expected supergames. We also find that the strategies used to support cooperation depend on the parameters of the game in a systematic way: as the probability of continuation increases subjects tend to choose shorter punishments (i.e. TFT over Grim).⁴ In addition, a low payoff to joint cooperation is associated with a higher preponderance of a variant of TFT that starts by defection: Suspicious-Tit-For-Tat (STFT).

Our elicitation method also provides an opportunity to test the validity of the methodology proposed by Dal Bó and Fréchette (2011), called Strategy Frequency Estimation Method (SFEM), to estimate the prevalence of strategies using observed round by round cooperation decisions. Using the laboratory to generate data that allows evaluating the performance of an estimation method is a recent approach (see Bajari and Hortacsu 2005, Fréchette et al. 2005, and Salz and Vespa 2019). It has the advantage, over simulation methods, that it does not require the researcher to assume the specific way in which the data differs from the assumptions of the model. In a simulation, the potential problems are thought of and generated by the researcher (sample size, incorrect assumptions about the error distribution, etc.).⁵ Here, the laboratory serves as an alternative, “organic,” data generating process. As such, it can reveal problems we would not have considered.⁶ Furthermore, although a simulation can identify situations where the SFEM does not perform well, it does not speak to whether these situations are more or less likely to occur.⁷ For instance, a simulation can identify situations where the SFEM systematically misattributes

⁴ The only prior studies of strategy choice under perfect monitoring over several parameter combinations are Dal Bó and Fréchette (2011) and Fudenberg et al. (2012). Both estimate strategic prevalence from round-by-round behavior; the former considers 6 strategies and the latter 11 in their estimation. Neither article studies how strategy choices vary with the payoff parameters or probability of continuation nor games that are as long, on average, as the ones considered here.

⁵ Even if there is agreement on potential concerns, the exact form they take can be unclear. For instance, many experimenters think session effects are present in certain experiments, yet they rarely specify what form they take. These are many, and could be simple to model or come about in complicated ways from the interactions of subjects (see Fréchette 2011 for a discussion). The performance of an estimation method will be sensitive to the specific way in which these session effects come about. The current approach does not require specifying what, if anything, session effects are.

⁶ This is, for instance, what Fréchette et al. (2005) report. They find that in using coalition governments data to test theories of bargaining, a potential problem that had not been considered is that the specific way people play these games in the laboratory suggests there will be identification problems.

⁷ Bajari and Hortacsu (2005) use the laboratory in this way when they use experimental data to select which of four structural models performs best at recovering the underlying value distribution. They are either establishing the model whose assumptions are “closest” to how people play, or the approach that is most robust to the actual deviations from the assumption. Vespa and Salz (2019) illustrate that a potential problem in doing counterfactual analysis in a repeated game setting (namely changing equilibrium selection as the parameters change) is indeed an actual problem that appears in the laboratory and find that its magnitude is non-negligible in practice.

behavior to the wrong strategies. However, it cannot determine whether such strategies are used and thus to what extent this is an actual problem. Clearly, elicitation is not a substitute for simulations, as they both have advantages, but we believe using the laboratory in this way, as a data generating process to study econometric estimation techniques, has much to offer.

We find that despite the aforementioned hurdles for identification, the estimation methodology SFEM does well under certain conditions. In fact, the qualitative results are often correct, even if games are relatively short, on average, or if some of the strategies are not included in the estimation. However, if a strategy that is excluded is relatively important, then results become misleading as the “closest” strategy is overestimated. In contrast, we find that including in the estimation strategies that are not used by subjects results only in small errors in the estimated prevalence of strategies. Overall, our results show that it is possible to learn about the strategies used by the subjects to support cooperation from their round by round cooperation decisions.

There is related work asking subjects to construct strategies in repeated games. For example, the literature starting with Axelrod (1980b) studied how human generated strategies for infinitely repeated games perform in computer tournaments. That literature focused on the relative performance of different strategies.⁸ Axelrod’s (1980a) first competition was a finitely repeated game. Another study that elicits strategies but focuses on the case of finite repetitions is that of Selten et al. (1997). Bruttel and Kamecke (2012) provide a partial elicitation of infinitely repeated game strategies. They are not interested in strategies, however; their motivation is to find alternative ways of inducing repeated games in the laboratory. Embrey et al. (2016a, 2016b) ask subjects for simple strategies in Cournot and Bertrand games allowing for (costly) deviations. Elicitation is used in this setting as estimation is not practical given the number of choices in their stage games.

Recent work by Romero and Rosokha (2018, 2019a, and 2019b) uses an alternative method to elicit strategies in infinitely repeated games. Their method is more flexible than ours, as it does not restrict the set of strategies but rather lets subjects choose actions as a function of freely defined choice histories. This approach may be particularly useful in settings for which researchers are uncertain of the relevant types of strategies. One disadvantage, however, is that it requires the researcher to impose defaults for certain histories (or in case of conflicts). It is also not easy to

⁸ Most of that literature has moved towards simulations rather than tournaments. Examples of more recent papers using computer simulations are Nowak and Sigmund (1993), who introduce stochastic strategies, and Nowak et al. (1995), who add mutations.

program, and thus less accessible. Moreover, our design includes several features to reduce or eliminate the impact of elicitation on behavior as we will discuss. Romero and Rosokha (2019a) use their method to explore the impact of making strategy changes costly, and Romero and Rosokha (2019b) extend their approach to allow for mixed strategies and confirm the result that most subjects use simple memory-one strategies.

The next section describes the main experimental design, in particular how we elicit strategies. Section III shows that there is no clear evidence that elicitation affected behavior and describes the chosen strategies and how their prevalence is affected by the parameters of the game. Section IV uses the elicited strategies as a benchmark to study the performance of SFEM (the methodology developed in Dal Bó and Fréchette (2011) to estimate the prevalence of strategies using observed round by round cooperation decisions). Section V revisits the impact of elicitation, this time focusing on strategies.

II. Experimental Design

The experimental design is in three phases.⁹ In all phases, subjects participate in randomly terminated repeated prisoners' dilemma games. A repeated game or supergame consists of multiple rounds. After each supergame, subjects are randomly rematched.

In Phase 1 subjects simply play the randomly terminated supergames. Between supergames, they are reminded of the decisions they made in the last supergame and of the choices of the person they were matched with in that supergame. The first supergame to end after 20 minutes of play marks the end of Phase 1.

In Phase 2, subjects are first asked to specify a *plan of action* — i.e., a strategy. The type of plan of actions they can specify depends on the session. In some sessions, subjects specify their plan of action by answering five questions: “In round 1 select $\{1, 2\}$,” and the answer to the four questions covering all permutations of “After round 1 if, I last selected $[1, 2]$ and the other selected $[1, 2]$, then select $\{1, 2\}$.” That is, in these sessions, elicitation is limited to strategies with memory one or less.¹⁰ We call this the *memory-one* elicitation interface. The choices are presented as drop-

⁹ When first reading the instructions, subjects are informed that there are multiple phases, but they are only told the procedures for Phase 1. They receive additional instructions after Phase 1. Instructions and screen-shots are available in the online appendix A.

¹⁰ While this elicitation is based on the outcome of the previous round, subjects can still report the two memory zero pure strategies, AD and AC, by choosing the same action in each of the five questions. For simplicity we will refer to all the strategies that can be elicited with these five questions as memory-one.

down menus, and the order in which the four questions appear after round 1 is randomized. In other sessions, subjects have a larger set of strategies to choose from. In addition to the memory-one strategies, subjects can choose from a *menu* of strategies. The menu includes some of the strategies they can build under memory-one (for example AD), but it also includes some additional strategies allowing for more lenient and forgiving behavior. The complete list of strategies in the menu elicitation can be found in online appendix C. The order strategies appear on the screen is randomized. We call this the *menu* elicitation interface.

After having specified their plan of action, subjects then play the supergame just as in Phase 1, making decisions in every round. At this point, the plan of action they specified is irrelevant. After the first supergame in Phase 2, they are shown the decisions they made in this supergame, the decisions the person they were matched with made, and the decisions that would have been made according to the plan of action they specified, had it been playing instead. They are then asked to specify a plan of action for the next supergame. This process (specify a plan; play a supergame round by round; receive feedback; and specify a plan) is repeated for 20 minutes. After 20 minutes of play in Phase 2, the plan of action takes over for the subjects, finishes the ongoing supergame, and plays an additional 14 supergames; this is Phase 3. That is, in Phase 3 the computers follow the plans of action elicited in the final supergame of Phase 2 and subjects make no choices.¹¹

The stage game is as in Table 1. Each subject is exposed to only one treatment (between-subjects design). The main treatment variables are the payoff from mutual cooperation R , the probability of continuation δ , and the elicitation method. In the main treatments, the payoff from mutual cooperation R takes values 32 or 48, δ takes values $\frac{1}{2}$, $\frac{3}{4}$, or $\frac{9}{10}$, and the elicitation interface is memory-one or menu.¹² Not all combinations of parameters and elicitation methods were used.¹³ Three sessions are conducted per treatment with the memory-one interface, and two sessions for treatments with the menu interface.¹⁴ Payments are based on the sum of the points accumulated in the three phases converted to dollars, with 100 points equaling \$0.45. Given those parameters,

¹¹ Thus, the data from Phase 2 is composed of round-by-round cooperate-defect choices over multiple supergames, and strategy choices (which we will sometimes refer to as the elicited strategies)—one for each supergame.

¹² Appendix D presents results that show the robustness of the results in section III in several dimensions: even longer games ($\delta = 95/100$, $R = 32$), as well as richer feedback and an even more flexible elicitation interface for the case of $\delta = 9/10$ and $R = 32$.

¹³ The treatment $\delta = \frac{1}{2}$, $R = 32$ is only implemented with the memory-one interface, while all other main treatments ($\delta = 1/2$, $R = 48$, $\delta = 3/4$, $R = 32$, $\delta = 3/4$, $R = 48$, and $\delta = 9/10$, $R = 32$) are implemented using both memory-one and menu interfaces.

¹⁴ The original set of treatments involved the 2×2 factorial of $\delta = \frac{1}{2}$ and $\frac{3}{4}$ and $R = 32$ and 48 with the memory-one interface. These treatments correspond to the ones in Dal Bó and Fréchette (2011). Following a suggestion by James Andreoni we added $\delta = 9/10$, $R = 32$. Later, based on reactions to our initial results, we added sessions using the menu interface for all treatments except $\delta = \frac{1}{2}$, $R = 32$ as that treatment was unlikely to yield interesting results. Finally we added the treatments discussed in Appendix D to further stress test the robustness of our findings.

cooperation can be supported as part of a subgame perfect equilibrium (henceforth SPE) in all treatments, except for the one where $\delta = 1/2$ and $R = 32$.

TABLE 1—STAGE GAME PAYOFFS

	C	D
C	R, R	12, 50
D	50, 12	25, 25

The main design considerations are the following. First, one concern is that asking subjects for a strategy might affect their behavior. For this reason, the design includes Phase 1, where subjects have time to learn what they want to do in this environment before the elicitation of strategies starts. An additional concern is that subjects may not think in terms of strategies and may not know how to describe their plan. To address that concern, the design gives feedback about what both subjects did and what the plan of action would have done during Phase 2. This gives subjects an opportunity to determine the situations in which the specified plan of action is not doing what they actually want to do. Finally, subjects are incentivized to choose the plan of action that is closest to what they want to do since whenever they specify a plan of action, this may be the supergame where Phase 2 ends and where the plan of action takes over.

A final concern is whether the possible plans of action that subjects can specify are sufficient to express the strategies they want. Note that despite its simplicity, even the memory-one elicitation allows subjects to specify 32 strategies. These strategies include those that are independent of the past, such as AD and Always Cooperate (AC), and memory-one strategies for which behavior depends only on the past round, such TFT.¹⁵ The memory-one elicitation also allows subjects to construct plans of action that behave exactly as some strategies with memory greater than one. For example, the plan of action that starts by cooperating and will keep cooperating if there was mutual cooperation in the previous round functions exactly as the strategy Grim in this setup without mistakes in implementation.¹⁶ The strategy Grim was also directly available in the menu sessions. In conclusion, many of the strategies most often mentioned in the literature can be specified by the memory-one interface. Moreover, the menu interface allows subjects to construct many strategies with memory greater than one.

¹⁵ TFT starts by cooperating, and in subsequent rounds, matches what the other subject did in the previous round.

¹⁶ Since Grim can be implemented by this memory-one plan of action, we will call both “Grim.”

III. Experimental Results

III.a Data Summary

The main treatments involved a total of 372 NYU undergraduates who participated in 23 sessions, with an average of 16.51 subjects per session, a maximum of 22 and a minimum of 12. The subjects earned an average of \$27.10, with a maximum of \$43.45 and a minimum of \$12.26. The average number of rounds per supergame was 1.92, 3.66, and 9.42, in treatments with $\delta=1/2$, $\delta=3/4$, and $\delta=9/10$, respectively; and the maximum was 9, 22, and 44, respectively. The number of supergames per session (including Phase 3) ranged between 25 and 92. Table A1 in the online appendix provides details on each session.

III.b The Impact of the Elicitation of Strategies on Behavior

Before studying the subjects' strategy choices, we briefly establish whether the elicitation of strategies affects behavior. We start by studying differences between recommended actions by the plan of action and actual actions during phase 2, and how subjects change their plans of action.

In phase 2, only 7.7 percent of decisions by the subjects disagree with what their plan of action would have chosen in that round. This number goes from 11.6 percent in the first supergame to 6.9 percent in the last one, suggesting that most subjects are learning to express their desired strategies. Consistent with this idea, Figure 1 shows that the ratio of subjects who change their strategy from supergame to supergame decreases with experience. While 37 percent of subjects change their plan of action between the first and the second supergame of phase 2, only 14 percent do so between the last two supergames in phase 2.

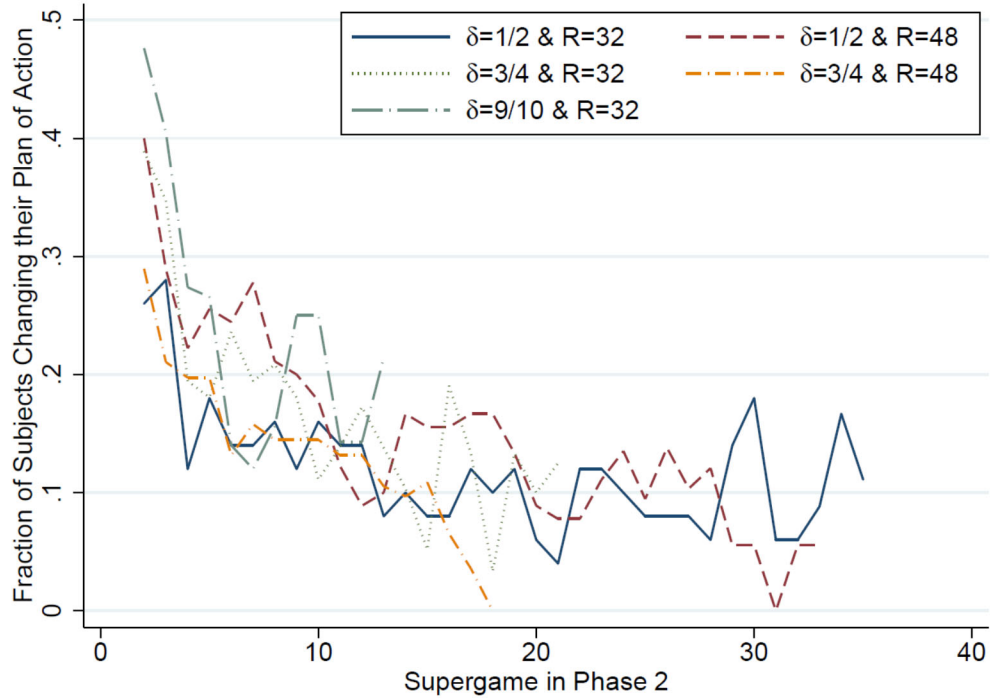


FIGURE 1. EVOLUTION OF THE FRACTION OF SUBJECTS CHANGING THEIR PLAN OF ACTION (PHASE 2)

The first two columns in Table 2 show that subjects are much more likely to change a plan of action if that plan was inconsistent with their actions in the previous supergame. The last column shows that the majority (59 percent) of these changes resolve the previous inconsistency.

TABLE 2—ON CHANGES OF STRATEGIES IN PHASE 2

δ	R	Percent of Supergames in which the Strategy Is Changed		Percent of New Strategies that Solve Previous Disagreement
		Disagreement between Actions and Strategy in Prior Supergame		
		No	Yes	
$\frac{1}{2}$	32	10.08	27.33	74.55
	48	14.19	35.65	59.63
$\frac{3}{4}$	32	12.85	41.51	60.43
	48	13.06	38.10	48.61
9/10	32	17.00	58.11	58.81
Overall		12.99	39.60	59.48

Given that not all subjects come up with their desired plan of action immediately, could the strategy elicitation affect subjects' behavior?

To answer this question, we use data from Dal Bó and Fréchette (2011) —where the experimental design is identical to the one in this paper with the exception that strategies are not elicited—as a benchmark to compare with behavior in sessions featuring elicitation.¹⁷ The comparison of behavior in sessions with and without elicitation needs to overcome three difficulties: first, elicitation did not start in the same supergame in all sessions; second, behavior evolves across supergames as subjects gain experience (see Figure A1 in online appendix); and, third, there are random variations in the cooperativeness of subjects across sessions. To overcome these difficulties, we will compare behavior when strategies are being elicited (phase 2) with behavior when strategies are not being elicited (phase 1 of this experiment and all supergames in Dal Bó and Fréchette 2011) after controlling for experience variables and average cooperativeness in each session (this is analogous to a difference-in-difference estimation). To do this we perform the following regression analysis at the treatment level:

$$C_{sm} = \alpha m + \beta Length_{sm-1} + \gamma Elicitation_{sm} + FE_s + \varepsilon_{sm},$$

where C_{sm} denotes the rate of cooperation in supergame m in session s (this is the actual behavior and not the one suggested by the chosen strategy in phase 2), the variable m denotes the supergame number (to control for the evolution of behavior), $Length_{sm-1}$ is the number of rounds in the previous supergame (which has been shown to affect cooperation), $Elicitation_{sm}$ is a dummy equal to one if strategies were elicited in that supergame, FE_s are fixed effects at the session level (to control for differences in cooperative tendencies across sessions already present in Phase 1), and ε_{sm} is an error term which may be correlated within a session. The unit of observation is a supergame in a session. We consider, alternatively, all rounds, round 1 only, and round 2 after each of the possible histories. Studying behavior for each of the possible round 2 histories is important as it allows us to study whether elicitation affects contingent behavior (at least for round 2, for which we have a large enough number of observations). The regression results are presented in Table A3 in the online appendix, which shows that the estimated coefficient for *Elicitation* is sometimes positive and sometimes negative, reaching statistical significance at the 10 percent level

¹⁷ Procedures for Phase 1 are almost identical to procedures in Dal Bó and Fréchette (2011). Besides issues of timing (Phase 1 is shorter than the entire experiment in Dal Bó and Fréchette 2011), the only difference is that subjects were not reminded of the choices they, and the other player, had made in the previous supergame between supergames in Dal Bó and Fréchette (2011). Remember that the elicitation instructions were not shown to subjects until the start of phase 2, so any differences in behavior in phase 1 cannot be attributed to the elicitation or the expectation of future elicitation.

in only two of 24 cases.¹⁸ Overall, there is little evidence that elicitation affects behavior. An additional piece of evidence suggesting that elicitation did not affect behavior is that the order of the five questions in the memory-one interface did not affect choices (the fact that behavior is not affected by the specific details of the elicitation suggests that elicitation overall does not have an effect either). In summary, we have not found evidence that the proposed method for the elicitation of strategies systematically affects behavior. We come back to the topic of the effect of the elicitation on behavior in section V.

III.c Is Memory One Enough?

Before analyzing the strategies subjects use, we discuss whether the memory-one interface is sufficient for subjects to express their preferred strategy in this environment. Note that in their estimation of strategies, Dal Bó and Fréchette (2011) essentially assumed that the relevant strategies were memory-one (all but one of the strategies considered have memory one, and the one that did not turned out not to be relevant). By contrast, Fudenberg et al. (2012) allow for four memory-2 strategies and two memory-3 strategies among the eleven strategies they consider in their estimation and find that a majority of strategies have memory greater than one under imperfect public monitoring but that is not the case under perfect monitoring. Their estimated prevalence of memory-one strategies under perfect monitoring is between 40 and 100 percent depending on the treatment. Hence, whether or not memory-one strategies are sufficient to describe play under perfect monitoring is an open question, which the elicitation with an infinite number of possible strategies can address more directly and comprehensively than estimation with a small number of strategies.

The results from the menu interface suggest that indeed memory one is sufficient for most subjects. Although 76 percent of the subjects chose a strategy from the menu and not from the memory-one interface to specify strategies when both are available to them, most of the menu strategies selected are equivalent to those available in the memory-one set (for example, AD, Grim, and TFT). The main result from the menu interface sessions is that 91 percent of the final strategies could be defined using the memory-one interface. More specifically, 97 percent of the strategies

¹⁸ A similar conclusion is reached if the same analysis is done without including the data from Dal Bó and Fréchette (2011). In this case, for identification, we exploit the fact that elicitation (phase 2) started at different supergames across sessions.

in the last supergame for $\delta=1/2$ and $R=32$ can be defined using the memory-one set of strategies; 96 percent and 87 percent when $\delta=3/4$ and $R=32$ or 48, respectively; and 85 percent for $\delta=9/10$ and $R=32$. Furthermore, in all treatments, the most popular strategy that cannot be expressed using the memory-one interface ranks among the least popular strategies. In only one treatment ($\delta=9/10$ and $R=32$), where it ranks fifth, it is not the least popular.¹⁹

In addition to the above, which directly shows that when given the choice, subjects use memory-one (or less) strategies, we note the following. In the sessions with the memory-one interface, 95 percent of actions by the last supergame of phase two correspond to what the elicited strategy would do. This corresponds to 88 percent of subjects playing exactly as their elicited strategy would in all rounds of the last supergame of phase 2. This is another way of seeing that memory-one strategies are, for the most part, adequate to represent how subjects play in these games.

III.d Description of strategies

We start by defining and discussing some key strategies, after which we show their popularity by the end of Phase 2 for each treatment. The simplest strategies are Always Cooperate (AC) and Always Defect (AD). Another much studied strategy is Tit-For-Tat (TFT), which starts by cooperating, and in subsequent rounds matches what the other subject did in the previous round. The Grim trigger strategy (Grim) also starts by cooperating, cooperates as long as both players have cooperated in the last round, and defects otherwise.²⁰ A strategy that is often discussed in the literature is Win-Stay-Lose-Shift (WSLS, also known as Perfect TFT, or Pavlov). It starts by cooperating, and cooperates whenever both players made the same choice last round, and defects otherwise. WSLS is considered desirable because when it plays itself, it would not defect forever after a deviation. Finally, another strategy of interest is suspicious Tit-For-Tat (STFT). It starts by defecting and, from then on, matches what the other subject did in the previous round.

What can be expected? Both subjects choosing AD is a SPE in every treatment, and previous experiments on infinitely repeated games have observed situations in which some subjects defect. Hence, it seems plausible to expect that AD will be selected. What about strategies that support

¹⁹ In online Appendix D we show that these results are robust to considering higher δ ($\delta=0.95$), and elicitation with a larger set of strategies and more complete feedback.

²⁰ More precisely, Grim is the strategy that starts by cooperating and cooperates unless there has ever been a defection in the past. It is not possible to exactly construct this strategy using the memory-one interface. However, it is possible to construct a memory-one machine that will be equivalent to Grim in how it plays against any other strategy.

cooperation? TFT is a likely candidate given that it was the winner in Axelrod's (1980b) tournament. It is also a very intuitive strategy to specify. However, TFT is not subgame perfect in general (and in all the treatments we consider). Thus, TFT may not be popular, while WSLS, which has performed well in simulations and is a SPE for high δ , could. However, WSLS cannot support cooperation for as many values of δ as Grim. The Grim strategy, in that sense, is more robust, and it can support cooperation for much lower values of δ . However, once it starts defecting, Grim never stops. Thus, in that sense, it is not forgiving, which may make it less appealing. These different tensions make it unclear which strategy one should expect to see most.

Table 3 shows the distribution of chosen strategies at the end of Phase 2 by treatment with a focus on the strategies that reach more than 10 percent prevalence in at least one treatment: TFT, Grim, STFT and AD.²¹ To those, we add the two commonly studied strategies: AC and WSLS.

TABLE 3—DISTRIBUTION OF CHOSEN STRATEGY IN PERCENTAGE

(ELICITED STRATEGIES IN THE LAST SUPERGAME OF PHASE 2)

Strategy:	$\delta = \frac{1}{2}$		$\delta = \frac{3}{4}$		$\delta = 9/10$
	R		R		R
	32	48	32	48	32
AC	0.00	2.22	1.39	5.26	1.19
TFT	6.00	18.89	9.72	28.95	34.52
Grim	6.00	41.11	13.89	36.84	15.48
WSLS	0.00	1.11	0.00	2.63	2.38
STFT	12.00	3.33	11.11	0.00	10.71
AD	58.00	25.56	43.06	14.47	15.48
Other	18.00	7.78	20.83	11.84	20.24
Memory 1 or equivalent strategies in Menu Interface		97.06	96.43	86.67	85.29
SPE	58.00	71.11	65.28	56.58	35.71
NE	86.00	90.00	65.28	86.84	71.43
TFT/(TFT+Grim)	50.00	31.48	41.18	44.00	69.05
Observations:	50	90	72	76	84

²¹ Tables A5 and A6 in the online appendix show the full distribution of strategies.

Four main results emerge from the direct observation of the choice of strategies. First, the most popular strategies across treatments are TFT, Grim, and AD.²² These three strategies on their own correspond to more than 65 percent of the data in each treatment and as much as 80 percent in two treatments. In contrast, WLS has a very limited presence. These results are largely consistent with previous estimation-based results for perfect monitoring environments.²³ Second, the two most popular strategies to support cooperation, Grim and TFT, vary in systematic ways across treatments.²⁴ In the treatments where cooperation can be supported in equilibrium (that is, all treatments but $\delta = 1/2$ and $R=32$), the ratio of TFT to TFT + Grim increases as the continuation probability increases. The difference, going from $\delta = 3/4$ to $\delta = 9/10$ when $R = 32$, is significant at the 10 percent level, but the difference is not significant when going from $\delta = 1/2$ to $\delta = 3/4$ when $R = 48$; and the p-value of the joint test is 0.16 (without clustering, p-values are 0.05, 0.19, and 0.06 respectively).²⁵ This suggests that how subjects choose to support cooperation depends on the particular parameters of the game in a systematic way—that is, as the expected length of the game increases, subjects choose to use shorter punishments.

We think that there are two possible reasons behind this difference: 1) Greater δ makes it possible to support cooperation with shorter periods of punishment. To see this, consider supporting cooperation using a trigger strategy with finite punishment. As δ increases, the value of cooperation is higher as compared to the temptation of deviating and a shorter punishment is sufficient to maintain incentives. 2) Another potential reason has to do with robustness (see for example the work of Fudenberg and Maskin 1990 and 1993, and Dal Bó and Pujals 2012). Higher

²² One may wonder if the popularity of Grim for some parameter combinations may be an attempt to commit to strong punishments even if subjects are unlikely to follow through with them. We can use the analysis of disagreement between round-by-round actions and the strategy recommendation to shed light on this issue. Consider the following histories: a subject first cooperate, while his partner defects, then in round two he defects while the other cooperates. Grim indicates that this subject should continue defecting in round 3, and those who select Grim fail to do so in only 17% of supergames. As a point of comparison, faced with the same situation, subjects who select TFT fail to choose cooperate (as TFT dictates) in 16% of cases. It does not appear less likely that a subject follows through with a “harsh” punishment strategy.

²³ See Dal Bó and Fréchette (2011), Fudenberg et al. (2012), Fréchette and Yuksel (2016), and Dal Bó and Fréchette (2018). The exception is Breitmoser (2015), who argues that a majority of subjects follow a strategy (semi-Grim) that randomizes with a higher probability of cooperation after mutual cooperation than under mutual defection, and in all other cases cooperates with the same intermediate probability. However, as discussed in Dal Bó and Fréchette (2018), if subjects followed a strategy that includes random behavior, there should be an amount of variability that is not observed in the data. In fact, most of the behavior can be perfectly accounted for by the simple pure strategies in Table 3. We find that this is the case even in very long supergames. In long supergames, the probability that a strategy involving randomizations would behave exactly like one of the simple strategies is negligible. Moreover, the difference in behavior documented in footnote 22 directly refutes semi-Grim. If subjects played semi-Grim, they should all cooperate at the same rate in those histories.

²⁴ While Dal Bó and Fréchette (2011) and Fudenberg et al. (2012) estimated the prevalence of strategies for several combinations of δ and R , they did not study how the prevalence of Grim and TFT changed with these parameters.

²⁵ These statistical tests are assessed using linear probability models with indicator variables for each treatment. Standard errors are clustered at the level of the session. All tests are two-sided. However, it seems unlikely that strategy choice is correlated within session since this requires that subjects can determine the strategies of the other players simply by observing their choices. This is a challenging task even for the econometrician, who has much more data than the subjects. For this reason, we also provide results that do not cluster standard errors when comparing strategy propensities across treatments. In the case of these tests, the independent variable is a dummy for TFT, and the observations are limited to subjects who chose TFT or Grim as a strategy.

δ implies that reverting to defection forever will be more costly since the expected number of future rounds increases. Although not relevant on the equilibrium path, if agents consider the possibility that others may make mistakes, or tremble, then such considerations are relevant.

The third main result about strategy choice refers to STFT. When the payoff to joint cooperation (R) is low, STFT is observed, and its popularity decreases as R increases. Of the two possible comparisons, one is significant at the 5 percent level, when $\delta = 3/4$, and one at the 10 percent level, when $\delta = 1/2$; the joint test is significant at the 5 percent level (p-values 0.01, 0.05, and 0.005 respectively without clustering). Note that STFT also decreases with δ , but these differences are not statistically significant. Also note that the treatments with $R = 32$ are the ones with the lowest cooperation rates. Hence, the higher frequency of STFT could be coming from subjects who have a “preference” for cooperating, but do not want to be taken advantage of, a type of strategic conditional cooperator.²⁶

Finally, as shown by the fact that many subjects chose TFT or STFT, we find that subjects are willing to choose strategies that are not part of a symmetric pure-strategy subgame perfect equilibria.²⁷

Hence, the results confirm some of the earlier observations about the use of simple strategies under perfect monitoring (Dal Bó and Fréchette 2018) but using direct elicitation, which allows for a much larger set of possible strategies than used in the estimation of strategies in previous papers. The fact that subjects still chose simple strategies when many more complicated strategies are available provides clear evidence that a vast majority of subjects prefer to use simple strategies under perfect monitoring. In addition, the current experiment extends this result to longer supergames than studied in previous papers. And although we do not expand on this, the cooperate-defect choice data confirm previous results on how changing R and δ impact cooperation rates and also extends those to longer games. Finally, the elicitation of strategies provides the first direct evidence of a systematic relationship between the environment (payoffs and continuation probability) and the strategies subjects use to support cooperation. We now consider how our

²⁶ We note that this observation, as well as the ones above relating to the effect of increasing δ on strategies, are speculative. Our experiment is not designed to tease apart the motives behind these changes.

²⁷ It could be argued that what is important is whether the distribution of strategies is a Nash equilibrium. Figure A2 in the Appendix reports the expected payoff of each strategy elicited when playing the population of strategies. Two observations can be made; the most popular strategies are, for the most part, among the ones that perform best. However, there is a minority of strategies that earn significantly lower payoffs. In that sense, behavior is not in line with the distribution of strategies being a Nash equilibrium.

experimental design allows us to understand better our ability to recover strategies from round-by-round choices.

IV. Does econometric estimation recover the same strategies?

Our direct observation of subjects' choice of strategies provides a benchmark to evaluate whether the strategies used by subjects can be recovered econometrically from observed behavior. We can compare the estimated prevalence of strategies using the cooperate-defect choices from Phase 2 to the prevalence of strategies that subjects actually chose (i.e. the elicited strategies).²⁸

We study the performance of the estimation procedure proposed in Dal Bó and Fréchette (2011) and also used in Fudenberg et al. (2012), Camera et al. (2012), Fréchette et al. (2013), Fréchette and Yuksel (2013), Jones (2014), Breitmoser (2015), Rand et al. (2015), Bigoni et al. (2015), Vespa (2015), Vespa and Wilson (2015), Fréchette et al. (2016), Arechar et al. (2016), to name a few. Let us refer to this approach as the Strategy Frequency Estimation Method (SFEM). The SFEM assumes that subject i chooses strategy k with probability ϕ^k for all supergames. In each round the subject plays according to the chosen strategy with probability $\beta \in \left(\frac{1}{2}, 1\right)$ and makes a mistake with probability $(1-\beta)$. We can derive the likelihood function as follows. By s_{imr}^k , denote the choice that subject i would make in round r of supergame m if he followed strategy k . s_{imr}^k is a function of past choices by both players and is coded as 1 for cooperate and 0 for defect. The choice that subject i actually made in that round and supergame is denoted by c_{imr} (also coded 1 for cooperate and 0 for defect), and the indicator function taking value 1 when the two are the same and 0 otherwise is $I_{imr}^k = 1\{c_{imr} = s_{imr}^k(\cdot)\}$. The likelihood that the observed choices were generated by strategy k is $\Pr_i(s^k) = \prod_{M_i} \prod_{R_{im}} (\beta)^{I_{imr}^k} (1-\beta)^{1-I_{imr}^k}$ for a given subject and strategy (where M and R represent the sets of supergames and rounds) where β is a parameter to be estimated. When β is close to $1/2$, choices are almost random, and when it is close to 1, choices are almost perfectly predicted.²⁹

²⁸ Readers should be careful not to confuse our use of *elicited strategies* (that are directly elicited from subjects) and *estimated strategies* (that are inferred—estimated using round-by-round cooperate-defect choices).

²⁹ Note that we assume that the error probability is the same for all strategies. Fudenberg et al. (2012) report that allowing the error probability to differ across strategies does not affect their estimates much. We also explored such a case and found no important changes.

From this, we obtain a loglikelihood $\sum_I \ln \left(\sum_K \phi^k \Pr_i(s^k) \right)$ where K represents the set of strategies we consider, labeled s^1 to s^K , and ϕ^k are the parameters of interest—namely, the proportion of the data attributed to strategy s^k . One can think of this in a large sample as giving the probability of observing each strategy. The SFEM estimate is the vector of strategy prevalence that maximizes this loglikelihood function.³⁰

To get a better sense of this approach, think about the case where only one strategy is considered; then, the only parameter to be estimated is β . Suppose that the estimate is 0.8 . In other words, when the strategy suggests cooperation [defection], cooperation [defection] is predicted to occur 80 percent of the time. The quality of the fit can be compared to chance, which would predict cooperation with a 50 percent probability (since there are only two choices). Now, consider the case of more than one strategy. Imagine that the strategies included are: AD, TFT, and Grim, and that the estimates of their proportion is one third for each of them, and β is still 0.8 . This implies that in round 1 of a supergame, the estimated model predicts a 60 percent cooperation rate. Now, suppose that we look at a specific subject who first cooperated, but he was matched with someone that defected; then, in round 2, the estimated model predicts that he will cooperate with a 20 percent probability. If the person they are matched with cooperated in round 2, the estimated model's prediction for round 3 would now be a 40 percent chance of cooperation.³¹

This model is estimated on Phase 2 data (from both memory-one and menu sessions—note that there are no menu sessions for $\delta = 1/2$ and $R = 32$ and, thus, a smaller sample) after dropping the first one-quarter of the data (as a function of the total number of rounds played). The first quarter of Phase 2 is dropped to limit the importance of supergames where behavior is still changing. The variance-covariance matrix is obtained by bootstrap.³²

The comparison of the estimated prevalence of strategies with the strategies actually chosen allows us to study the performance of the SFEM. When evaluating the performance, our main

³⁰ Some conditions are required for the vector of strategy prevalence to be identified. It has to be the case that the data includes behavior for histories in which two of the considered strategies differ, as otherwise it would be impossible to distinguish between them. Fudenberg et al. (2012) and online Appendix E provide support for the SFEM's ability to uncover the true distribution of strategies based on Monte Carlo simulations.

³¹ In round 1, AD predicts cooperation with probability 0.2, while the other two strategies with probability 0.8. In round 2, all three strategies predict defection with probability 0.8. In round 3, only TFT has cooperation as more likely than defection.

³² 1000 samples are drawn. A session is selected, then subjects within that session are drawn, finally supergames for each subject are randomly chosen. We use bootstrapping as it provides a simple way to compute the variance-covariance matrix while respecting the structure of the data generating process.

concern will be whether the SFEM identifies the key strategies, the ones that are most prevalent, and whether it can establish their relative importance.³³

We start by studying the performance of the SFEM when the set of strategies the estimation allows for is the “right” one. To do this, the sample is restricted to the cases where the elicited strategies are among the strategies presented in Table 3: AC, AD, TFT, Grim, WSLS, or STFT. In addition, we first estimate the model with all six strategies and then re-estimate it keeping only the strategies for which the estimated proportion is not 0.³⁴ These results are presented in Table 4 (WSLS is not included in Table 4 because its frequency is always estimated to be 0). First, notice the high value of β , which indicates a very good fit. Qualitatively, the SFEM performs well at identifying the most important strategies. For instance, the most prevalent strategy is correctly identified in all treatments. In two of the five treatments, it correctly identifies the three most popular strategies and their order of popularity (the treatments with $\delta = 1/2$) while in a third treatment ($\delta = 9/10$) it identifies the three most popular strategies but mistakes the ranking of the second and third strategy. The table also reports whether the hypothesis that the estimated proportions equal the frequency elicited in the data can be rejected.³⁵ As can be seen, there are three cases (out of 22) for which the estimated and elicited frequencies are statistically different. Note, also, that although they are statistically different, the differences tend to be small in magnitude. In only one of these cases is the difference meaningful in terms of size (the proportion of TFT in the treatment with $\delta = 3/4$ and $R = 48$). When considering the joint hypothesis that all the estimated frequencies are equal to the elicited ones, that treatment is the only one for which the hypothesis of equality can be rejected.³⁶

³³ From a practical point of view, our approach here takes the elicited strategies as the benchmark to assess estimation results. However, one could view these results as providing evidence that the elicitation method “works,” in other words taking the estimation results as the point of comparison. We think one can reasonably take both of these positions, and we suspect different readers will feel differently about this. Our position (more broadly) is that both approaches accumulate evidence about what strategies people use. The fact that both methods point in the same direction makes it more likely that this is indeed the case. We do not take for granted that either approach is “correct,” but rather they provide different sources of evidence.

³⁴ Estimates close to or at the boundary of the parameter space cause difficulties for the estimation of the variance (see, for instance, Andrews 1998). Since this is a challenging issue with no straightforward solution and is not important for the point being made here, we simply re-estimate the model ignoring the strategies that have 0 weight in this case.

³⁵ To perform these hypothesis tests, we do a bootstrapped Wald test to take into account the variability in both the SFEM estimator and the variance in the elicited strategies. More specifically, when we randomly select the sample on which to perform the estimation, we also select the associated elicited strategies. The bootstrapped Wald statistic is the average of the Wald statistic for each bootstrapped sample testing that the estimator and the elicited frequencies are equal.

³⁶ To keep Table 4 easy to read, we do not include the standard errors, but these are available in the appendix and include tests of equality to 0 to show that the estimates have a reasonable amount of precision.

TABLE 4—ELICITED VS ESTIMATED STRATEGY PREVALENCE, RESTRICTED SAMPLES

		AC	TFT	Grim	STFT	AD	β	Joint Test ^a
$\delta = 1/2, R = 32$	Elicited			0.082	0.187	0.730		
	Estimated			0.000	0.243	0.757 [°]	0.98	0.154
$\delta = 1/2, R = 48$	Elicited	0.036	0.217	0.412	0.039	0.295		
	Estimated	0.059	0.132	0.476	0.059	0.274	0.96	0.520
$\delta = 3/4, R = 32$	Elicited	0.023	0.122	0.131	0.130	0.593		
	Estimated	0.000 [°]	0.133	0.103	0.281	0.483	0.94	0.369
$\delta = 3/4, R = 48$	Elicited	0.091	0.356	0.414		0.139		
	Estimated	0.165	0.130 ^{°°}	0.557 [°]		0.148	0.99	0.044
$\delta = 9/10, R = 32$	Elicited	0.022	0.504	0.220	0.071	0.183		
	Estimated	0.086	0.499	0.115	0.073	0.227	0.96	0.179

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

^a Reports the p-value of the joint hypothesis test that each estimated proportion equals the elicited proportion.

The exercise of Table 4 is meant mainly to facilitate hypothesis testing between the estimated proportions and the elicited proportions; however, in practice, the sample will not be restricted and certain strategies that subjects use may not be included in the SFEM. Hence, a natural question is whether the results would be qualitatively misleading once all the data are included. These results are presented in Table 5.

The larger the frequency of “Other” strategies, the more difficult it has to be for the estimates to be good. However, in this case, the estimates are still fairly close to the elicited frequencies. The estimates pick up the most popular strategies and, in most cases, the estimates are not statistically different from the elicited strategies. In particular, the strategy estimated to be the most popular is always the one that is most popular according to the elicitation. It is also interesting to consider which strategies are statistically different from 0, since this would typically be the standard used to determine if a strategy is important or not. In all but one treatment (namely, $\delta = 3/4$ and $R = 48$), the two most popular strategies are identified as statistically different from 0. What seems to be more difficult is to identify the relative importance of TFT and Grim when cooperation rates are high. This is not particularly surprising since higher rates of cooperation result in fewer

observations differentiating between those two strategies. However, in treatments where substantially more subjects choose Grim over TFT, Grim is estimated to be more popular than TFT and vice-versa. Note that, in the treatment with $\delta = 1/2$ and $R = 32$, the difference between the elicited and estimated for AD can easily be understood. As can be seen in Tables A5 and A6, strategies that here fall under “other” are almost all variants of AD that behave exactly as AD in our data.

TABLE 5—ELICITED AND ESTIMATED STRATEGY PREVALENCE, FULL SAMPLE

		AC	TFT	WSLS	Grim	STFT	AD	Other
$\delta = 1/2, R = 32$	Elicited	0.00	0.04	0.00	0.06	0.14	0.53	0.22
	Estimated	0.00 (0.00)	0.02 (0.03)	0.00	0.02 ^{°°°} (0.02)	0.24 (0.10)	0.72 ^{°°} (0.11)	
$\delta = 1/2, R = 48$	Elicited	0.03	0.19	0.03	0.35	0.03	0.25	0.12
	Estimated	0.07 (0.04)	0.06 [°] (0.07)	0.00	0.55 (0.13)	0.05 (0.05)	0.26 (0.10)	
$\delta = 3/4, R = 32$	Elicited	0.02	0.10	0.01	0.10	0.10	0.47	0.20
	Estimated	0.04 (0.02)	0.09 (0.05)	0.00	0.10 (0.09)	0.29 [°] (0.10)	0.48 (0.09)	
$\delta = 3/4, R = 48$	Elicited	0.08	0.30	0.03	0.35	0.00	0.12	0.14
	Estimated	0.18 [°] (0.06)	0.11 (0.12)	0.00 ^{°°°}	0.57 [°] (0.14)	0.01 (0.02)	0.13 (0.05)	
$\delta = 9/10, R = 32$	Elicited	0.02	0.39	0.01	0.17	0.05	0.14	0.22
	Estimated	0.08 (0.07)	0.42 (0.10)	0.00 ^{°°°}	0.18 (0.10)	0.11 (0.06)	0.21 (0.06)	

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

Also note that the strategies that represent less than 10 percent of the data are rarely identified as statistically significant. Hence, and to some extent not surprisingly, strategies that are present, but that are not very popular, are difficult to detect. A more interesting observation has to do with the ability to identify strategies in short games. Ex ante, it would have seemed reasonable to think that long supergames are required to identify strategies, since one needs many transitions to uncover the strategy. With $\delta = 1/2$, half the games last only 1 round, and, thus, strategy estimation may seem hopeless. However, there is no clear evidence from Table 5 that proportions are estimated any better in the longer games. A reason for this might be that, in the treatments we study, most subjects use memory-one strategies; thus, one does not need many interactions to uncover the strategies being used.³⁷ Another reason may be that when δ is low, we have many supergames.

³⁷However, Appendix E.iii. shows that even strategies with longer memories can be identified in treatments with low δ .

In the estimation above, although some of the elicited strategies do not correspond to the strategies allowed by the estimation, the most popular ones are all allowed for. Without prior knowledge, this may be difficult to guarantee. For instance, in Dal Bó and Fréchette (2011), six strategies were considered: AC, AD, Grim, TFT, WSLS, and T2: a trigger strategy that starts by cooperating and defects for two rounds following a defection of the other before returning to cooperation. Thus, considering this set of possible strategies allows us to study the effect of not having a prevalent strategy in the set of possible ones (STFT), while having one that is not used (T2).

Results for that specification are presented in Table 6 (table with standard errors in the online appendix). As can be seen by comparing the results in Tables 5 and 6, in the two treatments where STFT is more important, not including it (in addition to other missing strategies) in the estimation leads to a statistically significant overestimation of the fraction of AD. However, having T2 in the set of possible strategies, instead, has no effect on the estimation as it is correctly estimated that no one chooses it.

TABLE 6—ELICITED VS ESTIMATED STRATEGY PREVALENCE, STRATEGY SET FROM DAL BÓ AND FRÉCHETTE (2011)

		AC	TFT	WSLS	Grim	STFT	AD	T2
$\delta = 1/2, R = 32$	Elicited	0.00	0.05	0.00	0.06	0.14	0.53	0.00
	Estimation	0.00	0.02	0.00	0.02 ^{°°}		0.96 ^{°°°}	0.00
$\delta = 1/2, R = 48$	Elicited	0.03	0.19	0.03	0.35	0.03	0.25	0.00
	Estimation	0.06	0.06 [°]	0.00 ^{°°}	0.55		0.31	0.02
$\delta = 3/4, R = 32$	Elicited	0.02	0.10	0.01	0.10	0.10	0.47	0.00
	Estimation	0.04	0.12	0.00 ^{°°°}	0.10		0.74 ^{°°}	0.00
$\delta = 3/4, R = 48$	Elicited	0.08	0.30	0.03	0.35	0.00	0.12	0.00
	Estimation	0.18 [°]	0.11	0.00	0.57 [°]		0.14	0.00
$\delta = 9/10, R = 32$	Elicited	0.02	0.39	0.01	0.17	0.05	0.14	0.00
	Estimation	0.08	0.52 ^{°°°}	0.00 ^{°°°}	0.16		0.24	0.00

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

The prevalence of elicited strategies that were not included in the estimation was 0.22, 0.12, 0.2, 0.14, and 0.22 for each treatment respectively.

One might be interested in understanding what happens when the set of strategies considered is enlarged to include many strategies that are not used by the subjects.³⁸ To see most directly the effect of having strategies that are not used, we restrict the data sample to that of subjects whose elicited strategy is among AD, AC, Grim, TFT, and STFT (as done for the analysis shown in Table

³⁸ Although not directly motivated by this question, online Appendix E.iii. explores this question via simulations and suggests that the estimator still performs well when multiple strategies are included but not relevant.

4), but include in the estimation additional strategies that are not used by these subjects. The set of strategies considered for estimation includes the five strategies used by the subjects in the data sample plus six additional strategies (TF2T, TF3T, 2TFT, 2TF2T, and Grim 2).³⁹ This is the set of strategies used for estimation in Fudenberg et al. (2012). Table 7 shows the estimates from this specification.

The estimated prevalence of the additional strategies is never statistically significant and the estimated prevalence is mainly small. For two treatments ($\delta = 1/2$, $R = 32$ and $\delta = 3/4$, $R = 32$) none of the additional strategies are given any positive weight and the effect on the recovered prevalence's of the strategies actually present in the sample is either minimal or nonexistent. For the other three treatments, the prevalence incorrectly attributed to the additional strategies is modest, being at most 15 percent. However, in these three treatments, the impact on the estimated prevalence of the five relevant strategies is not distributed evenly across them; rather it tends to impact one of the cooperative strategies, either TFT or Grim depending on the case, given that the additional strategies are cooperative ones as well.

TABLE 7—ESTIMATED STRATEGY PREVALENCE, FUDENBERG ET AL. (2012) SPECIFICATION
RESTRICTED SAMPLES (OF TABLE 4)

	$\delta = 1/2, R = 32$	$\delta = 1/2, R = 48$	$\delta = 3/4, R = 32$	$\delta = 3/4, R = 48$	$\delta = 9/10, R = 32$
AD	0.77 (0.12)	0.27 (0.09)	0.48 (0.10)	0.15 (0.06)	0.23 (0.07)
AC	0.00 (0.00)	0.06 (0.03)	0.00 (0.01)	0.12 (0.06)	0.00 (0.02)
Grim	0.00 (0.00)	0.46 (0.15)	0.10 (0.09)	0.52 (0.15)	0.02 (0.06)
TFT	0.02 (0.01)	0.00 (0.04)	0.13 (0.05)	0.09 (0.09)	0.53 (0.10)
STFT	0.21 (0.12)	0.06 (0.06)	0.28 (0.11)	0.00 (0.00)	0.07 (0.05)
TF2T	0.00 (0.00)	0.00 (0.01)	0.00 (0.01)	0.02 (0.03)	0.00 (0.04)
TF3T	0.00 (0.00)	0.00 (0.03)	0.00 (0.01)	0.03 (0.03)	0.03 (0.06)
2TFT	0.00 (0.00)	0.14 (0.11)	0.00 (0.02)	0.00 (0.11)	0.00 (0.08)
2TF2T	0.00 (0.00)	0.01 (0.02)	0.00 (0.00)	0.02 (0.03)	0.07 (0.06)
Grim2	0.00 (0.00)	0.00 (0.04)	0.00 (0.01)	0.05 (0.06)	0.05 (0.06)
Grim3	0.00	0.00	0.00	0.00	0.00
β	0.98	0.96	0.94	0.99	0.96

³⁹ See online appendix C or Fudenberg et al. (2012) for the definition of these strategies.

Instead of considering the restricted sample, the set of strategies can be enlarged as above, but using the entire sample. In that case however, the set of strategies considered for estimation is not necessarily larger than the set of strategies used by the subjects, since a few of the elicited strategies are not included in the estimation. Table A18 of Appendix E.iii. presents those results. The specifications of this table and of the one that produces Table 5 overlap for five strategies (including the most popular ones): AD, AC, Grim, TFT, and STFT. Comparing estimates for these five strategies, only nine of the 25 estimates (across the five treatments) change. Of those nine that change, six actually move closer to the elicited value.

Hence, overall it appears that having too large a set of strategies for estimation can affect the results, but these changes are for the most part modest. It seems reasonable to argue that these effects are smaller than the large impact documented in Table 6 from not including an important strategy. This being said, as the number of strategies included increases, identification may (in practice) become more challenging.⁴⁰

In summary, our comparison of the SFEM estimated prevalence of strategies and the elicited ones, show that the SFEM will work well if the set of considered strategies includes the relevant ones used by the subjects, but the results will be misleading if even only one important strategy is not included. Given that our previous analysis finds that estimates from the SFEM are fairly robust to including strategies that are not used, erring on the side of caution by including more, rather than less, strategies seems appropriate.

V. More Evidence on the Effect of Elicitation on Behavior

Having studied how well the SFEM performs, now we use the SFEM to study whether the elicitation of strategies affects behavior in terms of the strategies used, complementing the analysis in Section III. We do this by comparing the SFEM estimated strategy prevalence in this paper, while subjects were being asked to construct strategies at the same time as they were making round by round decisions in phase 2, to those from Dal Bó and Fréchette (2011) where there is no strategy elicitation. These comparisons are reported in Table 8.

The specification used is the same as in Table 4, and it includes all treatments that were included in both Dal Bó and Fréchette (2011) and the current experiments. Beside the estimates and standard

⁴⁰ Arechar et al. (2016) suggests a procedure for considering extremely large strategy sets.

errors, the table also reports the results of the tests of equality for each strategy frequency estimate across samples, as well as a test considering all equalities jointly.

First note that many pairs are very close in magnitude: more than half of the estimates are within five percentage points of one another and 17 of the 24 pairs of estimates are within 10 percentage points.

TABLE 8—ESTIMATES FROM DF 2011 VS THIS PAPER

		AC	TFT	WSLS	Grim	STFT	AD	β
$\delta = 1/2, R = 32$	DF 2011	0.00 (0.00)	0.05 (0.05)	0.00	0.00 (0.00)	0.23 (0.14)	0.72 (0.17)	0.97
	This Paper	0.00 (0.00)	0.02 (0.03)	0.00	0.02 (0.02)	0.24 (0.11)	0.72 (0.12)	0.97
	P-value of equality of joint equality	1.00	0.66	1.00	0.41	0.97	0.99	1.00
$\delta = 1/2, R = 48$	DF 2011	0.07 (0.05)	0.33 (0.10)	0.02	0.02 (0.06)	0.15 (0.09)	0.41 (0.07)	0.94
	This Paper	0.07 (0.04)	0.06 (0.07)	0.00	0.55 (0.13)	0.05 (0.05)	0.26 (0.10)	0.95
	P-value of equality of joint equality	0.96	0.03	0.63	0.00	0.34	0.63	0.01
$\delta = 3/4, R = 32$	DF 2011	0.02 (0.02)	0.25 (0.10)	0.00	0.00 (0.02)	0.34 (0.11)	0.39 (0.11)	0.92
	This Paper	0.04 (0.02)	0.09 (0.06)	0.00	0.10 (0.09)	0.29 (0.10)	0.48 (0.10)	0.92
	P-value of equality of joint equality	0.60	0.18	1.00	0.28	0.72	0.55	0.69
$\delta = 3/4, R = 48$	DF 2011	0.05 (0.10)	0.42 (0.18)	0.00	0.53 (0.19)	0.00 (0.01)	0.00 (0.00)	0.96
	This Paper	0.18 (0.06)	0.11 (0.12)	0.00	0.57 (0.13)	0.01 (0.02)	0.13 (0.04)	0.97
	P-value of equality of joint equality	0.26	0.14	1.00	0.86	0.51	0.01	0.03

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.

The number of observations per treatment was 4174, 6012, 4546, and 4724 for each treatment respectively.

Second, most pairs are not statistically distinguishable. Out of the 48 parameter estimates (24 pairs) of strategy frequencies, only three pairs are statistically different, and these come from the two treatments in which cooperation rates were already different before the strategy elicitation phase (see Table A9 in the online appendix), and thus these differences are most likely due to sampling variation.

In the two treatments with statistically significant differences, not only were there cooperation rate differences before phase 2, but there were also differences in strategies before phase 2. Table

A10 in the online appendix presents the SFEM results for the last three quarters of phase 1 in the current experiment and for the Dal Bó and Fréchette (2011) data.⁴¹ There are already some differences in strategy choices in phase 1 for the two treatments where some estimates are statistically different in phase 2.⁴²

Note that even if the data generating process were identical in both cases (the 2011 experiments and the current ones), one would expect that some of these comparisons would be statistically different. That is why, to be exact, these hypothesis tests would require an adjustment, for instance a Bonferroni correction. Given that this would only result in even fewer statistical differences, we omit such an analysis.

In summary, the comparison of SFEM results from sessions with elicitation and sessions without elicitation finds no significant differences in strategies in two treatments. In the other two treatments there are already differences before the elicitation of behavior starts; suggesting that the elicitation itself has little or no effect on behavior, consistent with the analysis in section III.b.

VI. Conclusions

Several identification hurdles make it difficult to infer strategies from observed behavior in infinitely repeated games. We overcome these hurdles by asking subjects to design strategies that will play in their place. The elicitation method we propose has a few potentially important features. 1) It first introduces the game without strategy elicitation, allowing subjects to “learn” what they would do without confusion, interference, or influence of a different presentation. 2) The elicitation method is then introduced in a way such that subjects have opportunities to learn, through feedback, how to specify a strategy that corresponds to the way they play. 3) Finally, the incentives to specify the strategy correctly become greater as subjects gain experience, since the strategy they specify is more likely to be implemented. Although we do not investigate which of these features is important, they are present here and may be the reason why we find no clear evidence that such strategy elicitation affects behavior in significant ways, supporting the validity of this method.

⁴¹ Phase 1 for the Dal Bó and Fréchette (2011) data is defined as the first half of the experiment (which is about the same number of supergames as the phase 1 for the current study). We use the last three quarter of the supergames to keep the sample size similar to that in Table 8.

⁴² We note in passing that there are large differences in strategies used between phase 1 and phase 2, suggesting important changes over time. These are apparent from cooperation rates but understanding the specific patterns of strategic learning over time would be an interesting question to pursue.

We find that the strategies chosen by the subjects include some commonly mentioned strategies, such as AD, TFT and Grim. Most subjects seem to rely on strategies that have, at most, memory one. This result holds even in the treatments with high probability of continuation. However, this result probably does not hold in general, as there is evidence that subjects use more complex strategies in other environments such as those with imperfect monitoring (see Fudenberg et al. 2012). Understanding the features of the environment that lead people to use simple or more-complex strategies is an interesting question for future work.

The results also reveal systematic patterns in how the environment, i.e. the specific game parameters and continuation probability, affect strategy choices. As cooperation becomes “more valuable,” either because the payoff to joint cooperation increases or the continuation probability increases, subjects are more likely to use strategies that can support some cooperation. Moreover, as the continuation probability increases, TFT becomes more popular relative to Grim, suggesting that the choice of the particular cooperative strategy depends on the parameters of the game. Also, at low levels of joint cooperation payoff, STFT is more likely to be present, but its popularity decreases as the continuation probability increases. These new results could offer a starting point for modeling strategy selection.

The elicited frequencies of strategies offer a benchmark with which we compare the SFEM estimated frequencies based on the cooperate-defect decisions. We find that the SFEM performs reasonably well in terms of identifying the important strategies. This is encouraging, as many experiments do not elicit strategies directly, yet inferring strategies may prove important to understanding behavior. We also highlight the limits of such an estimation method. In particular we find that the SFEM can lead to misleading results if one of the key strategies is omitted, as it shifts most of the weight onto the “closest” strategy. It also has difficulty identifying strategies that are present, but only in small numbers. One surprising result is that, at least in the environment studied here, estimation results are as reliable for treatments with supergames that are short on average as they are for treatments with supergames that are long on average. Moreover, results from the estimation are still qualitatively in line with the elicitation—even when omitting up to 20 percent of the strategies used, if none of them on their own represents a large fraction. Overall, our results show that it is possible to learn about the strategies used by the subjects to support cooperation from their round by round cooperation decisions.

REFERENCES

- Andrews, D. W. K.** 1998. "Hypothesis testing with a restricted parameter space." *Journal of Econometrics*, 84: 155-199.
- Aoyagi, M., and G. R. Fréchet.** 2009. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." *Journal of Economic Theory*, 144: 1135-1165.
- Aoyagi, M., V. Bhaskar, and G. R. Fréchet.** 2016. "The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private." Mimeo.
- Arechar, A., A. Dreber, D. Fudenberg, and D. Rand.** 2016. "I'm just a soul whose intentions are good: The role of communication in noisy repeated games." Mimeo.
- Axelrod, R.** 1980a. "Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution*, 24: 3-25.
- Axelrod, R.** 1980b. "More Effective Choice in the Prisoner's Dilemma." *Journal of Conflict Resolution*, 24: 379-403.
- Bajari, P. and A. Hortacısu.** 2005. "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data." *Journal of Political Economy*, 113: 703-741.
- Baker, G., R. Gibbons and K. J. Murphy.** 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics*, 117: 39-84.
- Bigoni, M., M. Casari, A. Skrzypacz, and G. Spagnolo.** 2015. "Time horizon and cooperation in continuous time." *Econometrica*, 83: 587-616.
- Brandts, J., and G. Charness.** 2000. "Hot vs. Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*, 2: 227-238.
- Brandts, J., and G. Charness. 2011. "The strategy versus the direct-response method: a first survey of experimental comparisons." *Experimental Economics*, 14: 375-398.
- Breitmoser, Y.** 2015. "Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma." *American Economic Review*, 105: 234-248.
- Brosig, J., J. Weimann, and C.-L. Yang** 2003. "The Hot Versus Cold Effect in a Simple Bargaining Experiment." *Experimental Economics*, 6: 75-90.
- Bruttel, L. V. and U. Kamecke.** 2012. "Infinity in the lab: How do people play repeated games?" *Theory and Decision*, 72: 205-219.
- Camera, G., M. Casari, and M. Bigoni.** 2012. "Cooperative Strategies in Anonymous

- Economies: An Experiment.” *Games and Economic Behavior*, 75: 570-586.
- Dal Bó, P. and G. R. Fréchette.** 2011. “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence.” *American Economic Review*, 101: 411-429.
- Dal Bó, P. and G. R. Fréchette.** 2018. “On the Determinants of Cooperation in Infinitely Repeated Games: A Survey,” *Journal of Economic Literature*, 56(1): 60-114.
- Dal Bó, P. and E.R. Pujals.** 2012. “The Evolutionary Robustness of Forgiveness and Cooperation.” Mimeo.
- Dreber, A., D. Fudenberg, and D. G. Rand.** 2014. “Who Cooperates in Repeated Games: The role of altruism, inequality aversion, and demographics,” *Journal of Economic Behavior & Organization*, 98: 41-55.
- Dreber, A., D. G. Rand, D. Fudenberg, and M. A. Nowak.** 2008. “Winners don’t punish.” *Nature*, 452: 348-351.
- Embrey, M., F. Mengel, and R. Peeters.** 2016a. “Strategy Revision Opportunities and Collusion.” Mimeo.
- Embrey, M., F. Mengel, and R. Peeters.** 2016b. “Eliciting strategies in indefinitely repeated games of strategic substitutes and complements.” Mimeo.
- Engle-Warnick, J., W. J. McCausland, and J. H. Miller.** 2004. “The Ghost in the Machine: Inferring Machine-Based Strategies from Observed Behavior.” Mimeo.
- Engle-Warnick, J., and R. L. Slonim.** 2004. “The Evolution of Strategies in a Trust Game.” *Journal of Economic Behavior and Organization*, 55: 553-573.
- Engle-Warnick, J., and R. L. Slonim.** 2006. “Inferring Repeated-Game Strategies From Actions: Evidence From Trust Game Experiments.” *Economic Theory*, 54: 95-114.
- Fréchette, G. R.** 2011. “Session-Effects in the Laboratory.” *Experimental Economics*, 15: 485-498.
- Fréchette, G. R., Matthew Emrey, and Ennio Stacchetti.** 2013. “An Experimental Study of Imperfect Public Monitoring: Efficiency versus Renegotiation-Proofness.” Mimeo.
- Fréchette, G. R., J. H. Kagel, and M. Morelli.** 2005. “Behavioral Identification in Coalitional Bargaining: An Experimental Analysis of Demand Bargaining and Alternating Offers.” *Econometrica*, 73: 1893-1938.
- Fréchette, G. R. and S. Yuksel.** 2016. “Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination.” *Experimental Economics*.

- Forthcoming.
- Friedman, J. W.** 1991. "A Non-Cooperative Equilibrium for Supergames." *Review of Economic Studies*, 38: 1-12.
- Fudenberg, D., and E. Maskin.** 1990. "Evolution of cooperation in noisy repeated game," *American Economic Review*, 90(2): 274-279.
- Fudenberg, D. and E. Maskin.** 1993. "Evolution and Repeated Games." Mimeo.
- Fudenberg, D., and D. K. Levine.** 1997. "Measuring Players' Losses in Experimental Games." *Quarterly Journal of Economics*, 112: 507-536.
- Fudenberg, D., D. G. Rand, and A. Dreber.** 2012. "Slow to Anger and Fast to Forget: Cooperation in an Uncertain World." *American Economic Review*, 102: 720-749.
- Green, E. J. and R. H. Porter.** 1984. "Noncooperative Collusion Under Imperfect Price Information." *Econometrica*, 52: 87-100.
- Gueth, W., S. Huck, and A. Rapoport.** 1998. "The Limitations of the Positional Order Effect: Can it Support Silent Threats and Non-Equilibrium Behavior?" *Journal of Economic Behavior and Organization*, 34: 313-325.
- Hoffman, E., K. A. McCabe, and V.L. Smith.** 1998. "Behavioral Foundations of Reciprocity: Experimental Economics and Evolutionary Psychology." *Economic Inquiry*, 36: 335-352.
- Imhof, L. A., D. Fudenberg, and M. A. Nowak.** 2007. "Tit-for-Tat or Win-Stay, Lose-Shift?" *Journal of Theoretical Biology*, 247: 574-580.
- Jones, M. T.** 2014. "Strategic complexity and cooperation: An experimental study." *Journal of Economic Behavior & Organization* 106: 352-366.
- Klein, B. and K. B. Leffler.** 1981. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy*, 89: 615-641.
- Murnighan, J. K., and A. E. Roth.** 1983. "Expecting Continued Play in Prisoner's Dilemma Games." *Journal of Conflict Resolution*, 27: 279-300.
- Phelan, C. and E. Stacchetti.** 2001. "Sequential Equilibria in a Ramsey Tax Model." *Econometrica*, 69: 1491-1518.
- Rand, D. G., D. Fudenberg, and A. Dreber.** 2013. "It's the thought that counts: The role of intentions in reciprocal altruism." Mimeo.
- Romero, J. and Y. Rosokha.** 2018. "Constructing Strategies in the Indefinitely Repeated Prisoner's Dilemma Game." *European Economic Review*, 104: 185-219.

- Romero, J. and Y. Rosokha.** 2019a. "The Evolution of Cooperation: The Role of Costly Strategy Adjustments." *American Economic Journal: Microeconomics* Forthcoming.
- Romero, J. and Y. Rosokha.** 2019b. "Mixed Strategies in the Indefinitely Repeated Prisoner's Dilemma." Mimeo.
- Rotemberg, J. J. and G. Saloner.** 1986. "A Supergame-Theoretic Model of Price Wars During Booms." *American Economic Review*, 76: 390-407.
- Rotemberg, J. J. and M. Woodford.** "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy*, 100: 1153-1207.
- Roth, A. E. and J. K. Murnighan.** 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology*, 17: 189-197.
- Rubinstein, A.** 1998. *Modeling Bounded Rationality*. Cambridge: MIT Press.
- Salz, T. and E. Vespa.** 2016. "Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation." Mimeo.
- Selten, R., M. Mitzkewitz and G. R. Uhlich. (1997).** "Duopoly Strategies Programmed by Experienced Players." *Econometrica*, 65: 517-555.
- Selten, R.** (1967), "Die Strategiemethode zur Erforschung des eingeschränkt rationale Verhaltens im Rahmen eines Oligopolexperiments," in H. Sauer mann (ed.), *Beiträge zur experimentellen Wirtschaftsforschung*, Tübingen: Mohr.
- Vespa, E.** 2015. "An experimental investigation of strategies in the dynamic common pool game." Mimeo.
- Vespa, E., and A. Wilson.** 2015. "Dynamic games and markov perfection: Putting the "conditional" in conditional cooperation." Mimeo.