

Discussion Paper Series – CRC TR 224

Discussion Paper No. 341
Project C 01

Individual Strategy Choice in the Repeated Prisoner's Dilemma

Teresa Backhaus¹
Yves Breitmoser²

June 2024
(*First version: March 2022*)

¹ University of Bonn, E-mail: tbackhaus@uni-bonn.de
² Bielefeld University, E-mail: yves.breitmoser@uni-bielefeld.de.

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Individual Strategy Choice in the Repeated Prisoner's Dilemma

Teresa Backhaus*
University of Bonn

Yves Breitmoser
Bielefeld University

June 17, 2024

Abstract

Reanalyzing 12 experiments on the repeated prisoner's dilemma (PD), we find strong evidence for players' use of behavior strategies. Starting with unrestricted memory-1 strategies, the most parsimonious non-rejected representation of behavior distinguishes three subject types: defectors, cautious cooperators and strong cooperators. The defectors defect with a high probability in every round. Both cooperating types play semi-grim behavior strategies with different cooperation rates in round 1. This simple three-type mixture fits significantly better than 10^{46} combinations of (generalized) pure strategies from the literature, which we fitted at the treatment level. Semi-grim behavior strategies fit better than all 10^{46} mixtures of (generalized) pure strategies even when we use a constant and pre-defined specification, without using free parameters or any kind of post-hoc econometric magic. Furthermore, the resulting type shares correlate with the treatment parameters in a predictable manner, and the strategies themselves are largely predictable thanks to their approximate invariance, but the strategies cannot be rationalized as responses to expected payoffs.

JEL-Code: C91, C72, C73, D12

Keywords: Repeated game, Behavior, Tit-for-tat, Mixed strategy, Estimation, Laboratory experiment

*We thank Andrea Ariu, Kai Barron, Niels Boissonnet, Friedel Bolle, Fabian Dvorak, Sebastian Fehrler, Drew Fudenberg, Daniel Gietl, Paul Heidhues, Steffen Huck, Gustav Karreskog, Johannes Leutgeb, Vlada Pleshcheva, Sebastian Schweighofer-Kodritsch, Roland Strausz, Robert Stüber, Georg Weizsäcker, Roel van Veldhuizen, and Joachim Winter, as well as the audiences in Schwanenwerder, Florence, Tel Aviv, Konstanz, and Berlin for many helpful comments. An earlier version of the paper was circulated as “God does not play dice, but do we?”. Financial support by the Deutsche Forschungsgemeinschaft (BR 4648/1, CRC TRR 190, and CRC TR 224) is gratefully acknowledged. Teresa Backhaus: Institute for Applied Microeconomics, Adenauerallee 24–42, 53113 Bonn, Germany, email: tbackhaus@uni-bonn.de, phone: +49 228 73-9244. Yves Breitmoser: Universitätsstr. 25, 33615 Bielefeld, Germany, email: yves.breitmoser@uni-bielefeld.de.

1 Introduction

One of the most dynamic research fields over the last two decades has been behavioral game theory, i.e. the econometric and theoretical analysis of laboratory games to align observed behavior with game-theoretical concepts. How should we think of beliefs, utilities, and subjects' choices, and is it possible to explain choices as responses to incentives? There has been substantial progress in aligning observed behavior and theoretical predictions across many classes of games. In generic normal-form games involving dominated strategies, behavior is captured after relaxing rational expectations (Costa-Gomes et al., 2001); in games without dominated strategies, behavior tends to mainly reflect logistic errors in choice (Weizsäcker, 2003; Brunner et al., 2011); and in games involving the distribution of monetary benefits, preference interdependence seems to organize behavior (Fehr and Schmidt, 1999; Charness and Rabin, 2002). Similarly, in single-object auctions, behavior shows to be reasonably consistent with theory after accounting for risk aversion (Bajari and Hortacsu, 2005), biased beliefs (Eyster and Rabin, 2005), or projection (Breitmoser, 2019).

One class of games that has experienced less progress in aligning behavior and predictions is the large class of repeated games, including even supposedly simple instances such as the repeated prisoner's dilemma. Repeated games are the main approach toward modeling long-run interactions, in particular to study cooperation and defection, and they have been a core object of game-theoretic analyses at least since the folk theorem for repeated games with discounting (Fudenberg and Maskin, 1986). Regarding individual strategies used in experiments, however, there is no consensus on what subjects actually do¹—not even, whether they play pure, mixed or behavior strategies.

In this paper, we propose a new approach towards the analysis of strategies in repeated games, and apply it to the repeated prisoner's dilemma in order to answer three questions: Which qualitative strategies do subjects actually play? Is there heterogeneity across subjects? And, are the particular strategies played and the shares of them predictable across conditions?

Regarding the first question, much of the existing literature restricts attention to strategies that are pure (with trembles), but recent evidence suggests that behavior strategies might better explain behavior (reviewed below). Regarding the second and third question, existing evidence suggests that the type shares playing specific cooperative strategies fluctuate without obvious patterns between treatments, which is puzzling but may reflect inadequate constraints in the strategy estimation that we shall relax in our analysis.

¹Dal Bó and Fréchette (2018) summarize determinants of first-round and aggregate cooperation rates but there exists no consensus on general or specific properties of individual strategies.

More generally, the existing literature distinguishes four approaches towards analyzing strategy choice in repeated games: indirect elicitation approaches (e.g. Engle-Warnick and Slonim, 2006; Bruttel and Kamecke, 2012), direct elicitation approaches (e.g. Dal Bó and Fréchette, 2019; Romero and Rosokha, 2023), model-free clustering (e.g. Heller and Tubul, 2023), and model-based clustering by means of mixture models (e.g. Dal Bó and Fréchette, 2011; Fudenberg et al., 2012). All existing approaches seem to have advantages and disadvantages. Indirect elicitation approaches restrict attention to pure strategies, which is a restriction that more recent approaches seek to avoid. Direct strategy elicitation asks subjects to define repeated game strategies (automata) in a way following Axelrod (1980a,b), and these automata are then matched against the automata programmed by say 20 other players. The overall payoff accumulated in all these say 20 supergames determines the experimental pay-out of subjects. Elicitation is appealing due to an apparent transparency and direct identification of the strategies used. However, they entail a potential bias towards simple strategies, for which expected payoffs may be easier to compute, and they have the important disadvantage of disturbing the incentives of players by offering hedging opportunities: As players are committed to their automaton for many supergames, the single automaton choice determines the aggregate payoff of hundreds of round-by-round decisions. This enables hedging over hundreds of lotteries, as discussed in more detail below, and induces an incentive to play other strategies than in hot play where choices are made one-by-one. To our knowledge, it is currently not possible to quantify the behavioral bias implied by the incentive change in programming automata, amongst others because it depends on the beliefs about the automata of others. This implies that we cannot say how close the programmed automata are to the strategies actually used in hot play.

Thirdly, model-free clustering approaches estimate individual strategies from hot play (e.g. by estimating individual cooperation strategies in each memory-1 state) and then estimates clusters of strategies to obtain say five representative strategies that characterize a population (e.g. by K -means clustering). Analyzing hot play, model-free clustering does not bias incentives and requires relatively few assumptions such as pre-selection of candidate strategies (beyond the restriction to memory-1, for example). However, model-free clustering cannot correct for measurement errors regarding individual strategies, it cannot control for differences in individual sample sizes per state for different subjects (who enter different states different numbers of time), as discussed shortly, and it requires large data sets to estimate the parameters of say five representative strategies without further assumptions. Experimental data sets are sufficiently large for this approach only after pooling observations from several treatments and experiments, which is adequate under the assumption that the

strategies are the same across treatments, but this assumption has been refuted in prior work (e.g. Fudenberg et al., 2012). As a result, aggregating observations from several treatments, clustering will tend to overestimate the plurality of strategies, as strategies from different treatments are observed, and considering the aforementioned measurement errors, the plurality will likely be overestimated even further. As an example of the latter, consider a fair coin that is tossed four times, by 100 subjects, and let us cluster subjects with respect to their relative frequency of tossing heads. Note, that individual strategies are often estimated based on few (e.g. 4) observations for some states. We ran such simulations, and in all our simulation runs, we estimated at least 3 clusters (based on optimum average silhouette width in k -means clustering), one with subjects hitting heads 50%, and the other two with approximately 15% and 85% probability. In strategy estimation, this issue is prevalent when estimating cooperation probabilities following rounds where one subject cooperated and the other one defected, as such states are observed comparably rarely.²

Finally, model-based clustering, usually by finite-mixture analyses (McLachlan and Peel, 2004), starts with a model of the data generating process, including certain models for clusters (strategy prototypes), and attempts to optimize the fit between the data and the model taking the actual number of observations per subject and state into account. In the coin tossing examples, this approach can use the information that each subject tosses their coin four times, which is not used in model-free clustering, and it estimates the different types of coins in the population. In our simulations, based on our implementation, this approach has identified the single cluster of a 50-50 coin, and in other simulations discussed below, it reliably identifies the qualitative strategies in simulated data sets of standard size in experimental repeated games. Modeling the data-generating process implies that the number of observations per states are accounted for in the likelihood function, avoiding the concerns about measurement errors emerging in direct estimation of individual strategies, the restriction to certain strategy prototypes reduces the number of free parameters substantially and enables estimation treatment-by-treatment, and the analysis of hot-play enables an unbiased analysis of behavior. The disadvantage is, as mentioned, the underlying restriction to a set of strategy prototypes. This restriction serves to reduce the number of free parameters in the estimation, but can of course bias the results. Consider a possible case that the set of candidate strategies has two elements, e.g. always defect and tit-for-tat, then all subjects will inevitably be classified as playing one of the strategies—even in the extreme case that none of them actually played either strategy. In this paper, we present a novel approach towards

²These comments notwithstanding, Heller and Tubul (2023) estimate five strategy clusters, which are subjects playing either always defect or a strategy that is closer to a semi-grim strategy as estimated below than to any pure strategy. From this perspective, our results are fairly compatible.

resolving this concern, based on a transparent and evidence-backed pre-selection of candidate strategies and data mining to estimate an upper bound for the goodness-of-fit attainable by certain classes of strategies, in order to be able to reach reliable conclusions. This way, we maintain the advantages of model-based clustering while minimizing the disadvantages, in a re-analysis of a large data set comprising 12 experiments to robustly estimate the strategies that subjects play and study how they align with expected payoffs.

Our main results can be summarized as follows. First, on a data set comprising 145,000 decisions from 12 experiments, we use data-mining techniques to obtain an *upper bound* for the goodness-of-fit that could be obtained assuming all subjects play versions of pure strategies. We relax many assumptions made in the literature, grant many degrees of freedom “for free”, allow for systematic randomization in round 1, and allow for either no switching, random switching, or Markov switching of strategies between supergames. This way, we combat the aforementioned pre-selection bias and determine the model with the best-fitting pure strategies treatment-by-treatment, out of 10^{46} combinations of pure-strategy mixtures across treatments.

Second, we estimate a *lower bound* for the goodness-of-fit that can be reached when we relax the restriction to pure strategies and allow cooperating subjects play behavior strategies. To this end, we evaluate the simple population mixture where subjects play either always defect or one of two generic memory-1 behavior strategies (without restrictions to pure or otherwise known strategies). It provides a lower bound for the adequacy of behavior strategies, as it represents just one of many possible mixtures involving behavior strategies, implying that the best possible mixture will be different. However, this simplistic mixture still is highly significantly above the upper bound for the pure strategies. The result answers our first question regarding the qualitative strategies: Some subjects certainly play behavior strategies and strategy choice in the repeated PD can be understood only if the empirical analysis allows for subjects playing behavior strategies.

Regarding the second question, heterogeneity, we employ both a top-down and a bottom-up approach toward model selection. In the top-down approach, we start with a very general specification (three unrestricted behavior strategies plus always defect per treatment) and incrementally restrict the model until further restrictions are rejected at a statistically significant level.³ Independently in both, the first and the second halves of sessions, this top-down approach toward heterogeneity converges to a model containing three subject types across all

³Our unrestricted estimation is initially focused on memory-1 strategies, following a number of results in the literature on behavior in the repeated PD with perfect information (Dal Bó and Fréchette, 2018; Breitmöser, 2015), but the focus on memory-1 is corroborated by a robustness check (to memory-2) reported in the appendix. The memory-1 states will be abbreviated as $\emptyset, cc, cd, dc, dd$ in the following.

treatments and experiments: We refer to them as defectors, strong cooperators and cautious cooperators. The defectors play a slightly perturbed version of always defect and the cooperators both play behavior strategies predicting nearly pure behavior after cc and dd (joint cooperation and defection, respectively), and randomization after cd and dc . In round 1, the strong cooperators cooperate with high and the cautious cooperators with intermediate probability. The intuitive interpretation is that subjects expect cooperation when both cooperated in the previous round, or defection when both defected, and hence either cooperate or defect with very high probability. Otherwise, they are “unsure” and randomize.

Alternatively, in the bottom-up approach toward model selection, we start with an overly restricted model and keep removing restrictions until these restrictions stop being statistically significant. In general, the bottom-up approach may well select a different model than the top-down approach, but in our case, it converges to the very three-type mixture described above. Our starting point, the overly restricted model, is a model containing a simple behavior strategy previously hypothesized based on a small sample analysis in Breitmoser (2015), where subjects cooperate with the probabilities 0.9 after joint cooperation in the previous round, 0.3 after unilateral cooperation/defection, and 0.1 after joint defection. We obtain a prediction for round 1 by the implication of Markov perfect equilibrium that behavior in round 1 may equate with behavior in any other state. This yields a model with two types of “cooperating” subjects, one with first-round cooperation probability of 0.9 and the other with 0.3, and a “defecting” type that plays always defect – defined as usual. Note that all three probabilities (0.9, 0.3, 0.1) above are neither optimized nor optimal, thus providing an *a priori* reasonable though presumably restrictive model as starting point.

We observe that even this very simple mixture of three constant behavioral types fits weakly better than the above upper bound for pure strategies for experienced and significantly better for inexperienced subjects. This shows that behavior strategies fit better than pure strategies even without exploiting any free parameters or other sources of econometric magic, which we find quite intuitive: Average cooperation probabilities reasonably close to 0.9, 0.3, and 0.1 are observed in all experiments on the repeated PD with perfect monitoring (see Table 2 in the appendix for the exact numbers). Our analysis simply considers the possibility that subjects actually play the behavior strategy that we observe so robustly on average—instead of ruling out this possibility *ex-ante* and assuming that there must be hidden types of pure strategies with highly variable type shares that implicitly reproduce this average behavior again and again. Second, we observe that relaxing the three probabilities (0.9, 0.3, 0.1) to be treatment specific improves the goodness-of-fit significantly for experienced subjects. These results answer our second question: Subjects are heterogeneous –

the bottom-up and the top-down approach congruently lead us to a three-type model with defectors, strong cooperators, and cautious cooperators. Let us emphasize that this mixture of types is not necessarily optimal for any single treatment, but it *fits best on average* in our analysis, where we average across 32 treatments from 12 experiments to get robust results.

Regarding the third question, the shares of the three subject types vary significantly across treatments, while the cooperation probabilities in the five memory-1 states of the strategies vary relatively little across treatments. The latter are largely uncorrelated with treatment parameters or other known predictors of cooperation, but the distribution of types is highly predictable based on the discount factor δ and the Blonski et al. (2011) (BOS) threshold of cooperation δ^* . As δ approaches δ^* , the share of defectors decreases relative to cooperators, and as δ is raised further, the strong cooperators start to outnumber the cautious cooperators, see Figure 3. That is, allowing for behavior strategies in the estimation implies that the distribution of subject types stops being erratic and becomes predictable. This finding extends existing results on the determinants of round-1 cooperation (see e.g. Dal Bó and Fréchette, 2018; Embrey et al., 2018) towards determinants of defective, cautiously cooperative and strongly cooperative behavior, and it partially answers our third question. It suggests that subjects are aware of δ and other parameters when picking their strategy, while the actual strategies seem largely uncorrelated with δ (see Figure 2 below). In other words, subjects seem to be choosing one of three strategies depending on the environment, but hardly adapt the strategy as such to the environment. This impression is reminiscent of the automata discussed by Rubinstein (1986) and Schmidt (1993), amongst others, and reinforced by our final result that the cooperation probabilities in the five memory-1 states are not rationalizable as responses to expected payoffs under rational expectations. Overall, we conclude that while strategy shares are predictable and strategies as such are largely invariant (hence, also predictable), the strategies are not immediately rationalizable.

2 Background information

Definitions The *prisoner's dilemma* (PD) involves two players choosing whether to cooperate (c) or defect (d). In the normalized PD, each player's payoff is 1 if both cooperate and 0 if both defect. If exactly one player cooperates, the cooperating player's payoff is $-l$ ($l > 0$) and the defecting player's payoff is $1 + g$ ($g > 0$). The infinitely repeated PD is an infinite repetition of this constituent game, mostly assuming future payoffs are discounted exponentially (using factor $\delta < 1$). Laboratory experiments implement the indefinitely repeated PD, which is terminated with probability $1 - \delta$ after each round and where payoffs accrue for

all rounds at once after termination of the game. The two games are strategically equivalent under expected utility. We will refer to these games jointly as repeated PD (or, *supergame*). Given $g, l > 0$, cooperation is dominated in the one-shot game but may be sustained along the path of play in subgame-perfect equilibria of the repeated PD (depending on δ).

A *strategy* σ in the repeated PD maps all finite histories to probabilities of cooperation in the next round. The strategy has *memory-1* if it prescribes the same cooperation probability for any two histories not differing in the actions chosen in their respective last rounds. It has *memory-2* if the same holds for the last two rounds. We denote memory-1 strategies as $\sigma = (\sigma_0, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$ corresponding to the five memory-1 histories $\{\emptyset, cc, cd, dc, dd\}$, called *states* in the following. For example, σ_{cd} denotes the probability of cooperation when a player’s most recent action is c and her co-player’s most recent action is d , σ_0 denotes the action in the first round. A strategy is a *pure strategy* if it prescribes degenerate cooperation probabilities after all histories ($\sigma \in \{0, 1\}^5$), and it is a *behavior strategy* otherwise (Selten, 1975). It is a *mixed strategy*, when a player randomizes over the set of pure strategies prior to the start of each supergame, but sticks to the drawn pure strategy throughout the supergame. In contrast, when playing a behavior strategy, she randomizes during the supergame.⁴

Table 1: Overview of the most commonly analyzed strategies (see Table 13 in the appendix for a more comprehensive list)

Strategy	Abbreviation	Description	$(\sigma_0, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$
Always Defect	AD	Always defects	(0,0,0,0,0)
Always Cooperate	AC	Always cooperates	(1,1,1,0,0)
Grim	G	Only cooperate in R1 and after cc	(1,1,0,0,0)
Tit-for-Tat	TFT	Start with c, then copy opponent	(1,1,0,1,0)
Suspicious TFT	STFT, D-TFT	Start with d, then copy opponent	(0,1,0,1,0)
Win-Stay-Lose-Shift	WSLS	Cooperate in R1, cc and dd	(1,1,0,0,1)
Semi-Grim	SG	Behavior strategy satisfying ...	$\sigma_{cc} > \sigma_{cd} \approx \sigma_{dc} > \sigma_{dd}$

Note: The conventional definition of AC is (1, 1, 1, 1, 1), which is behaviorally equivalent to (1, 1, 1, 0, 0). The definition used above implies that any memory-1 behavior strategy that might be observed on average can be rebuilt using some combination of AD, AC, Grim, TFT and WSLS.

Related behavioral literature We will keep the literature review short and focused due to the availability of an excellent recent survey by Dal Bó and Fréchette (2018). The modern experimental research on the repeated PD started with Dal Bó (2005), who criticized earlier experiments for implementing experimental designs that let subjects play against comput-

⁴Including the case when she would switch between pure strategies within a supergame. Note that Kuhn’s Theorem, as generalized to infinite extensive-form games in Aumann (1964), implies that every mixed strategy has a behavior strategy representation, but not vice versa, implying that the distinction is relevant.

erized opponents. The first wave of experiments following Dal Bó (2005) includes Dreber et al. (2008), Duffy and Ochs (2009), Blonski et al. (2011) and Kagel and Schley (2013), and focuses on analyzing first-round and total cooperation rates. A second wave comprising Dal Bó and Fréchette (2011, 2019), Bruttel and Kamecke (2012), Camera et al. (2012), Fudenberg et al. (2012), Sherstyuk et al. (2013), Breitmoser (2015), and Fréchette and Yuksel (2017) analyzes the strategies actually chosen by players. The general theme in the reported results is that round-1 cooperation rates depend on the strategic environment. More specifically, the results indicate that subgame perfection of Grim is necessary but not sufficient for cooperation to emerge (first reported in Dal Bó, 2005), and that subsequent cooperation of subjects depends on their own and their opponent’s actions, primarily on those in the previous round. The central importance of initial cooperation is also demonstrated in Fudenberg and Karreskog (2023, forthcoming). Many of the second-wave analyses classify individual subjects’ strategies into varying sets of pre-selected strategies. Even allowing for noise, these analyses clearly show that subjects do not homogeneously follow a given pure strategy across all supergames. The studies differ in their assumptions of what subjects might be doing instead—whether they are playing pure, mixed, or behavior strategies—and consequently in their conclusions about behavior.

Many analyses assume that decisions are made only prior to the first supergame of a session, with subjects then sticking to a *pure strategy* (with trembles) for the rest of the session. Given this restriction to pure strategies, these analyses typically conclude that the majority of subjects play either AD, TFT, or Grim, with each being attributed weights around 20–30%. For example, Result 6 of Dal Bó and Fréchette (2018, DF18) states that these three strategies account for “most of the data”, specifically they “account for 70 percent of strategies in most treatments”, but importantly, this result is obtained after a-priori restricting attention to (a subset of) pure strategies without further validating this restriction. We refer to this statement as the *pure-strategy conjecture* and will seek to evaluate the identifying assumption.

A second, less common approach is based on the assumption that subjects randomly switch pure strategies between supergames, which resemble *mixed strategies* in the game-theoretical sense. For example, DF18 report that 84 percent of choices in supergames lasting more than one round are accounted for by five pure strategies (now also including AC and suspicious TFT) when they allow for strategy switching between supergames (DF18, Footnote 38).⁵ The difficulty now is to explain this strategy switching; otherwise, the impression

⁵Specifically, DF18’s observation states that subjects’ behavior is described “exactly” even if one “does not allow for any mistakes.”

of a perfect fit, not requiring a complicated analysis allowing for noise, is intriguing, but it is only true in a post-hoc sense. Ex-ante, the strategy chosen by a given subject is not predictable, and the game-theoretical concept closest to such random choice over varying pure strategies over time is that of a mixed strategy. The probabilities of choosing different pure strategies over time may be path dependent, given the path-dependency they may be degenerate, and they may be heterogeneous between subjects. Below, we shall explicitly allow for these possibilities by considering Markov-switching models to capture strategy switching that contain pure, mixed, and path-dependent mixtures as special cases.⁶ This will be one of the major novelties of our analysis and will enable us to determine an upper bound for the goodness-of-fit of pure, mixed and Markov-switching strategies.

A third and growing group of studies attempts to validate the restrictions to pure strategies by allowing subjects to randomize in each round of each supergame, as in the game-theoretical concept of *behavior strategies*. Relaxing the restriction to pure strategies, Breitmoser (2015) observed that cooperating subjects play a semi-grim behavior strategy ($\sigma_{cc} > \sigma_{cd} \approx \sigma_{dc} > \sigma_{dd}$) approximating $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.9, 0.3, 0.3, 0.1)$, without specifying σ_0 (*behavior-strategy conjecture*, c.f. Breitmoser, 2015, p. 2889). The intuition attributed to this observation is that subjects expect cooperation after *cc* and then cooperate with high probability, that they expect defection after *dd* and then defect with high probability, and that they are unsure after the mixed histories *cd*, *dc* and then randomize. This intuition directly entails a prediction for behavior in round 1: Subjects expecting cooperation will cooperate with high probability, subjects that are unsure will randomize, while subjects expecting defection would play always defect as usually assumed. The game-theoretic foundation for this prediction is that, if $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.9, 0.3, 0.3, 0.1)$ is part of an equilibrium strategy, which is approximately the case if symmetric belief-free equilibria exist (Breitmoser, 2015), then the three possible completions of this strategy to a symmetric equilibrium strategy are $\sigma_0 = 0.9$, $\sigma_0 = 0.3$, and $\sigma_0 = 0.1$ (round 1 can be equated with any of the subsequent states). Additionally, semi-grim behavior strategies are found to better capture behavior than certain mixtures of pure memory-1 strategies (Breitmoser, 2015).⁷ Recently, Fudenberg and Karreskog (2023, forthcoming) report evidence highlighting the predictive power of semi-grim strategies in repeated PDs, though the analysis focuses mainly on round-1 behavior.

A recent study by Romero and Rosokha (2023) elicits subjects' strategies by letting them construct memory-1 behavior strategies to be played by an automaton. Complementary to our results, they find that even in cold-play, that is, if subjects ex-ante must com-

⁶Note, Markov-switching is not covered by the definition of a mixed strategy.

⁷Few other studies investigate behavior strategies, e.g. Fudenberg et al. (2012), who include the strategy "generous TFT" with randomization after opponent's defection, and more recently Dvorak and Fehrler (2018).

mit to a rule of play for a sequence of 20 super-games and know that their opponent does too, many of them choose behavior strategies that involve randomization.⁸ This is striking evidence for the empirical relevance of behavior strategies. Interestingly, their study also shows that the average cooperation rates change from an approximate semi-grim pattern $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.86, 0.35, 0.45, 0.13)$ in hot-play, as in the experiments in our data, towards a so-called “mixed tit-for-tat” pattern $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.78, 0.17, 0.46, 0.14)$ in the cold-play tournament where the selected strategies have to play against 20 other supergame strategies without the possibility of interim modifications. This indicates that the necessity to commit to a strategy for 20 supergames (against varying opponents) indeed affects behavior. Such commitment is known to affect the set of equilibrium strategies (from subgame perfection towards Nash) favoring “unbeatable” strategies such as TFT (Axelrod, 1980a,b; Duersch et al., 2014); and it affects individual incentives substantially by offering “hedging” over bundles of 20 supergames.⁹ In light of this, behavior seems to be surprisingly robust, given that the effect is largely confined to a reduction of the cooperation probability in state cd by 18 percentage points.

The behavioral assumption that decisions are made in each round, instead of say once at the start of a session (as in the pure-strategy conjecture), seems intuitive.¹⁰ However, there are several concerns about Breitmoser’s results that might explain why the behavior-strategy conjecture faces skepticism: the data set might be fortunately selected in Breitmoser (2015), behavior might be more complex than memory-1 admits, strategies may be behavior strate-

⁸They elicit perturbed versions of TFT, which have slightly higher cooperation rates in state dc than cd and find that after allowing for re-specification a larger share of subjects specifies rules similar to pure strategies. Note the important strategic differences between this elicitation experiment and hot-play in standard PD games that are reflected in the average cooperation rates across the memory-1 states, discussed below.

⁹To see the impact of hedging on individual incentives, consider a lottery that pays \$10 with probability .5 and \$0 with probability .5, while the decision maker is slightly risk averse and would not be willing to buy the lottery ticket for \$4 (say $\alpha = 0.75$ in a CRRA utility $u(x) = x^\alpha$). This would be the round-by-round attitude towards risk. Now consider the case that she is offered the option to buy 20 such lottery tickets in a bundle, at the price \$80, while she will be informed only of the total payoff after all tickets had been checked by a computer. Would she be willing to buy this bundle of lotteries? Yes, and indeed she would be willing to pay \$99.35 (her certainty equivalent) for it – now she is “almost” risk neutral by comparison (thanks to bundling the risks and the implied hedging). Thus, her risk attitude would be much closer to risk neutrality when she commits to one strategy for 20 rounds (per supergame) and much more indeed if she commits to one strategy for 20 such supergames, as in Romero and Rosokha (2023). The structure in the experiment of Romero and Rosokha offers exactly this bundling of lotteries and thus is a strategically different situation than hot play – giving rise to the use of riskier strategies than we would use in hot play.

¹⁰Here follow the interpretation of behavior in repeated matching pennies games (e.g. Goeree et al., 2003), whereby the description of subjects randomizing say 50-50 each round is simply the best-possible description for the outside observer. Subjects themselves typically do perceive their decisions to be deliberate each round. Similarly, we consider a behavior strategy implying randomization each round to be the best-possible description available to observers of seemingly random but subjectively deliberate decisions that subjects make each round.

gies other than semi-grim, subjects might switch strategies as the session progresses, and round-1 behavior was not included in the estimation of strategies. In the next two sections, we address all of these concerns and report arguably conclusive answers to the following questions:

Question 1. *Which qualitative strategies do subjects play? Are they playing pure or non-degenerate behavioral strategies? Are they choosing different strategies across supergames?*

Question 2. *Which particular strategies do subjects play? Is there heterogeneity across subjects?*

The case for memory-2 strategies had been made by Fudenberg et al. (2012), who analyze the repeated PD with imperfect monitoring and show that if we assume subjects play pure strategies, then there must be subjects with memory-2, based on evidence for 2TFT and "lenient" Grim2 strategies. Similar ideas are expressed in Aoyagi and Fréchette (2009) and Bruttel and Kamecke (2012). We will provide robustness checks with memory-2, but in addition, we will relax the restriction to pure strategies, which seems critical since behavior strategies also generate decision patterns resembling memory-2 or -3.

The data We analyze the exact same set of experiments reviewed in Dal Bó and Fréchette (2018). This set comprises most of the modern experiments on the repeated Prisoner's Dilemma with perfect monitoring, i.e. those published since Dal Bó (2005), and consists in total of data from 12 studies, 32 treatments¹¹, more than 1900 subjects, and almost 145,000 decisions. The set of studies equates with the studies listed in Table 2. A brief review and an overview table is in Appendix B, but for a detailed discussion, see DF18. Due to its enormous size, the wide range of experiments covered (from different experimenters at various universities and in various countries), and its comprehensive character with respect to the recent list of experiments on the repeated PD, this data set appears to be optimal for our purposes. In addition, by sticking exactly to the list of experiments reviewed by Dal Bó and Fréchette (2018), we can rule out the notion that data selection biases the results in favor of any of the hypotheses we intend to test.

¹¹Most studies have several treatments with varying payoff parameters and continuation probabilities. We estimate our results at treatment level.

3 A model-free overview of behavior

In order to provide a foundation for the subsequent analysis and discussion, let us first provide an overview of behavior in the repeated PD without imposing restrictions reflecting any of the above stated three conjectures. To this end, we simply report average cooperation rates in both the first and second halves of sessions of all experiments and discuss how these average strategies align with expected payoffs across states.

Average behavior Table 2 reports the average cooperation rates across experiments in each of the four memory-1 states after round 1 and tests for significance of differences. For brevity, we aggregate across all treatments per experiment here but provide results by treatment upon request. Initially, we skip round-1 behavior as it varies substantially across treatments, as discussed below, but the cooperation rates in the remaining states are fairly similar across treatments and indeed across experiments, as Table 2 shows. In state cc , cooperation rates are above 0.9, in state dd they are mostly at or below 0.1 (with the sole exception of Aoyagi and Frechette, 2009), and after the mixed histories cd and dc , cooperation rates fluctuate somewhat in the range $[0.2, 0.5]$. The differences between inexperienced and experienced subjects, which we distinguish by first and second halves of sessions, are very minor overall. The aggregate cooperation probabilities shift by at most five percentage points. This observation notwithstanding, it is customary to distinguish experienced and inexperienced behavior, which we maintain also for this paper. In robustness tests, we examine alternative definitions of experienced behavior (to be behavior after 30 or 60 decisions) and in line with the very minor experience effects visible in Table 2, the results are highly robust.

Re-analyzing four experiments, Breitmoser (2015) made the observation that average memory-1 strategies have a “semi-grim” pattern $\sigma_{cc} > \sigma_{cd} \approx \sigma_{dc} > \sigma_{dd}$, with the approximation $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.9, 0.3, 0.3, 0.1)$. Based on the vastly extended data set analyzed here, we can scrutinize whether this somewhat surprising observation was related to (involuntary) data selection. To begin with, Table 2 shows that $(0.9, 0.3, 0.3, 0.1)$ is clearly no more than an approximation, but in some steps of our analysis, we shall use it nonetheless in order to avoid post-hoc specification adaptations of semi-grim.

We test for differences in the cooperation rates using bootstrapped p -values, resampling at the subject level, and distinguishing two levels of significance: the conventional level 0.05 and the tighter level $0.002 \approx 0.05/24$. The latter implements the Bonferroni correction for tests across 12 experiments and the two session halves. Naturally, we shall focus on this corrected level of significance, but for clarity we also report the conventional level that does

Table 2: Few subjects play pure strategies and assuming pure strategies yields a striking bias even in large mixture models

Experiment	Actual cooperation rates				Number of subjects not randomizing 50-50						
	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$	(c,c)	(c,d)	(d,c)	(d,d)			
First halves per session											
<i>Aoyagi and Frechette (2009)</i>	0.917	»	0.45	≈	0.408	≈	0.336	32/38	1/23	3/20	7/21
<i>Blonski et al. (2011)</i>	0.89	»	0.279	≈	0.193	»	0.034	13/17	1/5	3/3	124/135
<i>Bruttel and Kamecke (2012)</i>	0.91	»	0.286	≈	0.228	»	0.08	12/18	6/23	8/21	32/36
<i>Dal Bó (2005)</i>	0.922	»	0.212	<	0.342	»	0.089	13/13	0/3	2/2	42/54
<i>Dal Bó and Fréchette (2011)</i>	0.951	»	0.334	≈	0.331	»	0.063	94/106	28/117	51/128	218/253
<i>Dal Bó and Fréchette (2019)</i>	0.94	»	0.297	≈	0.335	»	0.057	216/243	37/137	62/147	404/474
<i>Dreber et al. (2008)</i>	0.904	»	0.217	≈	0.213	»	0.036	15/25	3/19	12/18	45/48
<i>Duffy and Ochs (2009)</i>	0.904	»	0.301	≈	0.33	»	0.111	43/57	4/25	10/24	61/82
<i>Fréchette and Yuksel (2017)</i>	0.943	»	0.141	≈	0.266	≈	0.091	21/28	0/0	2/2	5/8
<i>Fudenberg et al. (2012)</i>	0.982	»	0.4	≈	0.427	»	0.066	38/43	1/6	5/11	20/25
<i>Kagel and Schley (2013)</i>	0.935	»	0.263	≈	0.295	»	0.051	71/81	20/71	32/60	98/111
<i>Sherstyuk et al. (2013)</i>	0.945	»	0.328	≈	0.371	»	0.117	37/44	10/36	12/34	41/52
Pooled	0.938	»	0.304	≈	0.322	»	0.065	605/713	111/465	202/470	1097/1299
Second halves per session											
<i>Aoyagi and Frechette (2009)</i>	0.958	»	0.398	≈	0.517	≈	0.375	33/37	0/12	1/12	5/9
<i>Blonski et al. (2011)</i>	0.923	»	0.287	≈	0.231	»	0.02	26/32	10/25	11/16	172/178
<i>Bruttel and Kamecke (2012)</i>	0.947	»	0.221	≈	0.297	»	0.041	13/15	8/17	9/12	31/35
<i>Dal Bó (2005)</i>	0.92	»	0.242	<	0.388	»	0.064	18/27	0/3	0/1	50/65
<i>Dal Bó and Fréchette (2011)</i>	0.979	»	0.376	≈	0.362	»	0.041	132/137	34/89	62/100	196/215
<i>Dal Bó and Fréchette (2019)</i>	0.976	»	0.315	<	0.402	»	0.035	340/365	52/162	77/146	448/497
<i>Dreber et al. (2008)</i>	0.917	»	0.128	«	0.39	»	0.009	14/18	6/11	6/12	41/43
<i>Duffy and Ochs (2009)</i>	0.977	»	0.367	≈	0.391	»	0.082	80/87	5/35	16/43	60/68
<i>Fréchette and Yuksel (2017)</i>	0.97	»	0.233	≈	0.398	»	0.069	33/37	1/6	2/10	20/25
<i>Fudenberg et al. (2012)</i>	0.971	»	0.487	≈	0.412	»	0.083	41/44	2/8	4/10	14/17
<i>Kagel and Schley (2013)</i>	0.966	»	0.262	≈	0.332	»	0.025	87/90	16/56	30/46	91/97
<i>Sherstyuk et al. (2013)</i>	0.973	»	0.482	≈	0.437	»	0.078	44/48	7/24	17/23	23/29
Pooled	0.971	»	0.327	<	0.376	»	0.039	861/937	141/448	235/431	1151/1278

Note: The “actual cooperation rates” are the relative frequencies estimated directly from the data. The relation signs encode bootstrapped p -values (resampling at the subject level with 10,000 repetitions) where $<$, $>$ indicate rejection of the Null of equality at $p < .05$ and \ll , \gg indicating $p < .002$. Following Wright (1992), we accommodate for the multiplicity of comparisons within data sets by adjusting p -values using the Holm-Bonferroni method (Holm, 1979). As a result, if a data set is considered in isolation, the .05-level indicated by “ $>$, $<$ ” is appropriate. If all 24 treatments are considered simultaneously, the corresponding Bonferroni correction requires to further reduce the threshold to $.002 \approx .05/24$, which corresponds with “ \gg , \ll ”. Note that all econometric details here exactly replicate Breitmöser (2015), i.e. the statistical tests are not adapted post-hoc. The “number of subjects not randomizing 50-50” indicates the number of subjects with cooperation rates in the various states differing significantly from 50-50 (in subject-level two-sided binomial tests), conditioning on subjects having moved at least five times in the respective state. The required level of significance is set at $p = 0.0625$ such that five observations are sufficient to trigger statistical significance if the subject plays a pure strategy.

not correct for multiple testing.¹²

Out of all the 24 observations, considering first and second halves separately, only one observation, based on one session half in one experiment (Dreber et al., 2008), indicates a significant violation of the key restriction $\sigma_{cd} \approx \sigma_{dc}$, while the other two restrictions $\sigma_{cc} > \sigma_{cd,dc}$ and $\sigma_{cd,dc} > \sigma_{dd}$ are never violated significantly. In 45/48 cases they are even confirmed significantly at the tight 0.002 level surviving the Bonferroni correction. Pooling all observations from all experiments, $\sigma_{cd} \approx \sigma_{dc}$ is not rejected in the first halves of sessions but at the 0.05 level it is rejected in the second halves of sessions. The difference of σ_{cd} and σ_{dc} remains small, however, and is not significant at the 0.025 level surviving the Bonferroni correction considering that we run two simultaneous tests for the pooled data (one for the first halves of sessions and one for the second halves). Given this range of observations on a vastly extended data set, we conclude that Breitmoser’s observation passed the out-of-sample test on non-selected data, i.e. that average behavior indeed exhibits the semi-grim pattern.

Of course, this is but a first indicator for semi-grim patterns, not evidence for the individual strategies played. That is, if there is subject heterogeneity, mean cooperation rates provide unbiased estimates of the true cooperation rates but are not necessarily unbiased estimates of the mean strategies (e.g. due to selection effects after round 1). The behavior-strategy conjecture postulates that this semi-grim pattern does not only characterize the behavior on average but also the strategies of *individual* subjects. We will rigorously test this conjecture below. For now we add the thought that otherwise, the observation that this pattern recurs across all treatments and experiments would appear to be a surprising coincidence—considering that pure strategies are estimated to be played in strikingly varying weights across treatments (Dal Bó and Fréchette, 2018).

The results of a first simple test of this hypothesis are reported in the last four columns of Table 2. These columns list the number of subjects (per experiment) that deviate significantly from randomizing 50-50 in the four memory-1 states. We focus on subjects with at least five observations per state, which is sufficient to trigger significance in two-sided Fisher tests if subjects play a pure strategy. The results are fairly clear: In state *cd*, i.e. after unilateral defection of the opponent, all standard pure strategies (except AC, which is rarely observed though) agree on the (pure) prediction that one should defect. This state is unique with respect to the unanimity of the prediction. For this state, however, we find the lowest number of subjects significantly deviating from randomizing 50-50—only around a quarter of the subjects do so, putting a rather tight bound on the number of subjects potentially playing

¹²In Table 2, <, > indicate significance at the conventional level and <<, >> indicate significance surviving the Bonferroni correction (see the table notes for details).

pure strategies.

To further illustrate this bound, assume that subjects do use pure strategies: On one hand, given that the semi-grim pattern results on average, there have to be subjects that systematically cooperate after unilateral defection of opponents (state cd). These subjects are rarely found in analyses, as indicated most clearly by the aforementioned Result 6 of Dal Bó and Fréchette (2018), stating that “always defect” (AD), Grim, and tit-for-tat (TFT) are the “three strategies [that] account for most of the data”. This directly contradicts the observation that $\sigma_{cd} \approx \sigma_{dc} > \sigma_{dd}$, unless in addition to the strategies accounting for most of the data a substantial number of subjects systematically cooperate in state cd . However, the strategies predicting at least occasional cooperation after cd , such as always-cooperate and tit-for-2-tats, were found to fit behavior of only very few subjects in Dal Bó and Fréchette (2018). This contradiction foreshadows what we will find below: even allowing for drastic data mining, pure strategies cannot be pushed to fit behavior as well as a simple behavior strategy does.

Relation to monetary incentives Complementing the model-free description of behavior, let us look at what subjects should be doing under rational expectations. While relating the decisions “cooperate” and “defect” to expected payoffs in each state is a standard behavioral piece of information in analyses of static games, it is novel in analyses of repeated games. The underlying question, whether the actions chosen are at least qualitatively plausible, is of obvious relevance in any attempt to understand behavior.

For this initial model-free exposition, we will estimate the expected payoffs of cooperate and defect, in each state, from the perspective of an agent who assumes continuation play follows the average relative frequencies of cooperation observed above. These relative frequencies are denoted as the behavior strategy $\sigma = (\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$. Given σ , the expected payoff in state $\omega \in \{\emptyset, cc, cd, dc, dd\}$ is denoted as π_ω , with

$$\pi_\omega = \sigma_\omega \pi_\omega(c) + (1 - \sigma_\omega) \pi_\omega(d), \quad (1)$$

where $\pi_\omega(c)$ and $\pi_\omega(d)$ denote the expected payoffs of playing c and d in state ω ,

$$\pi_\omega(c) = \sigma_{\omega'} (\delta \pi_{cc} + (1 - \delta) \times 1) + (1 - \sigma_{\omega'}) (\delta \pi_{cd} + (1 - \delta) \times (-l)), \quad (2)$$

$$\pi_\omega(d) = \sigma_{\omega'} (\delta \pi_{dc} + (1 - \delta) \times (1 + g)) + (1 - \sigma_{\omega'}) (\delta \pi_{dd} + (1 - \delta) \times 0), \quad (3)$$

with continuation probability δ and ω' being the state ω from the opponent’s point of view,

such that σ_ω is the probability of cooperation by the opponent. Using the treatment-specific average behavior strategies σ from above, we can solve the linear equation system, Eqs. 1–3 for all ω , and obtain the expected payoffs $\pi_\omega(c)$ and $\pi_\omega(d)$ for each state in each treatment.

The monetary incentive to cooperate is $\pi_\omega(c) - \pi_\omega(d)$, for each ω . Figure 1 provides an overview of the results: We plot the relative frequencies of cooperation across treatments against the respective monetary incentives to cooperate for each state, separately for first and second halves of sessions. The states cd and dc are pooled for simplicity. Figure 1 additionally shows the best-fitting logistic curves, estimated without intercept such that neutral incentives $\pi_\omega(c) - \pi_\omega(d) = 0$ yield a predicted cooperation probability of 0.50. The pseudo- R^2 of the logistic curves indicate how much of the null deviance is explained by allowing for logistic errors in utility maximization.

The observations can be summarized as follows: For each state, we have observations from treatments with net incentives ranging from around -0.5 to $+1$, i.e. from cases where $\pi_\omega(c) - \pi_\omega(d)$ is highly negative to cases where it is highly positive. Essentially, the former obtains in treatments where Grim is not a subgame-perfect equilibrium strategy and the latter obtains in treatments where the discount factor δ is substantially above the threshold for Grim to be a subgame-perfect equilibrium strategy. Despite this range of induced monetary incentives, relative probabilities of cooperation and monetary incentives are highly correlated only in round 1 (state \emptyset). They are statistically close to independence in all states after round 1. For example, in second halves of sessions, when subjects have gained experience, the Pseudo- R^2 of the logit model is about 0.8 in round 1 and below 0.2 in all states afterwards. Obviously, this model-free analysis has the drawback of neglecting subject heterogeneity, which we will address below, but it seems that behavior in states cc and dd may be difficult to align with monetary incentives. For this reason, we raise the following question, which will be addressed in Section 5.

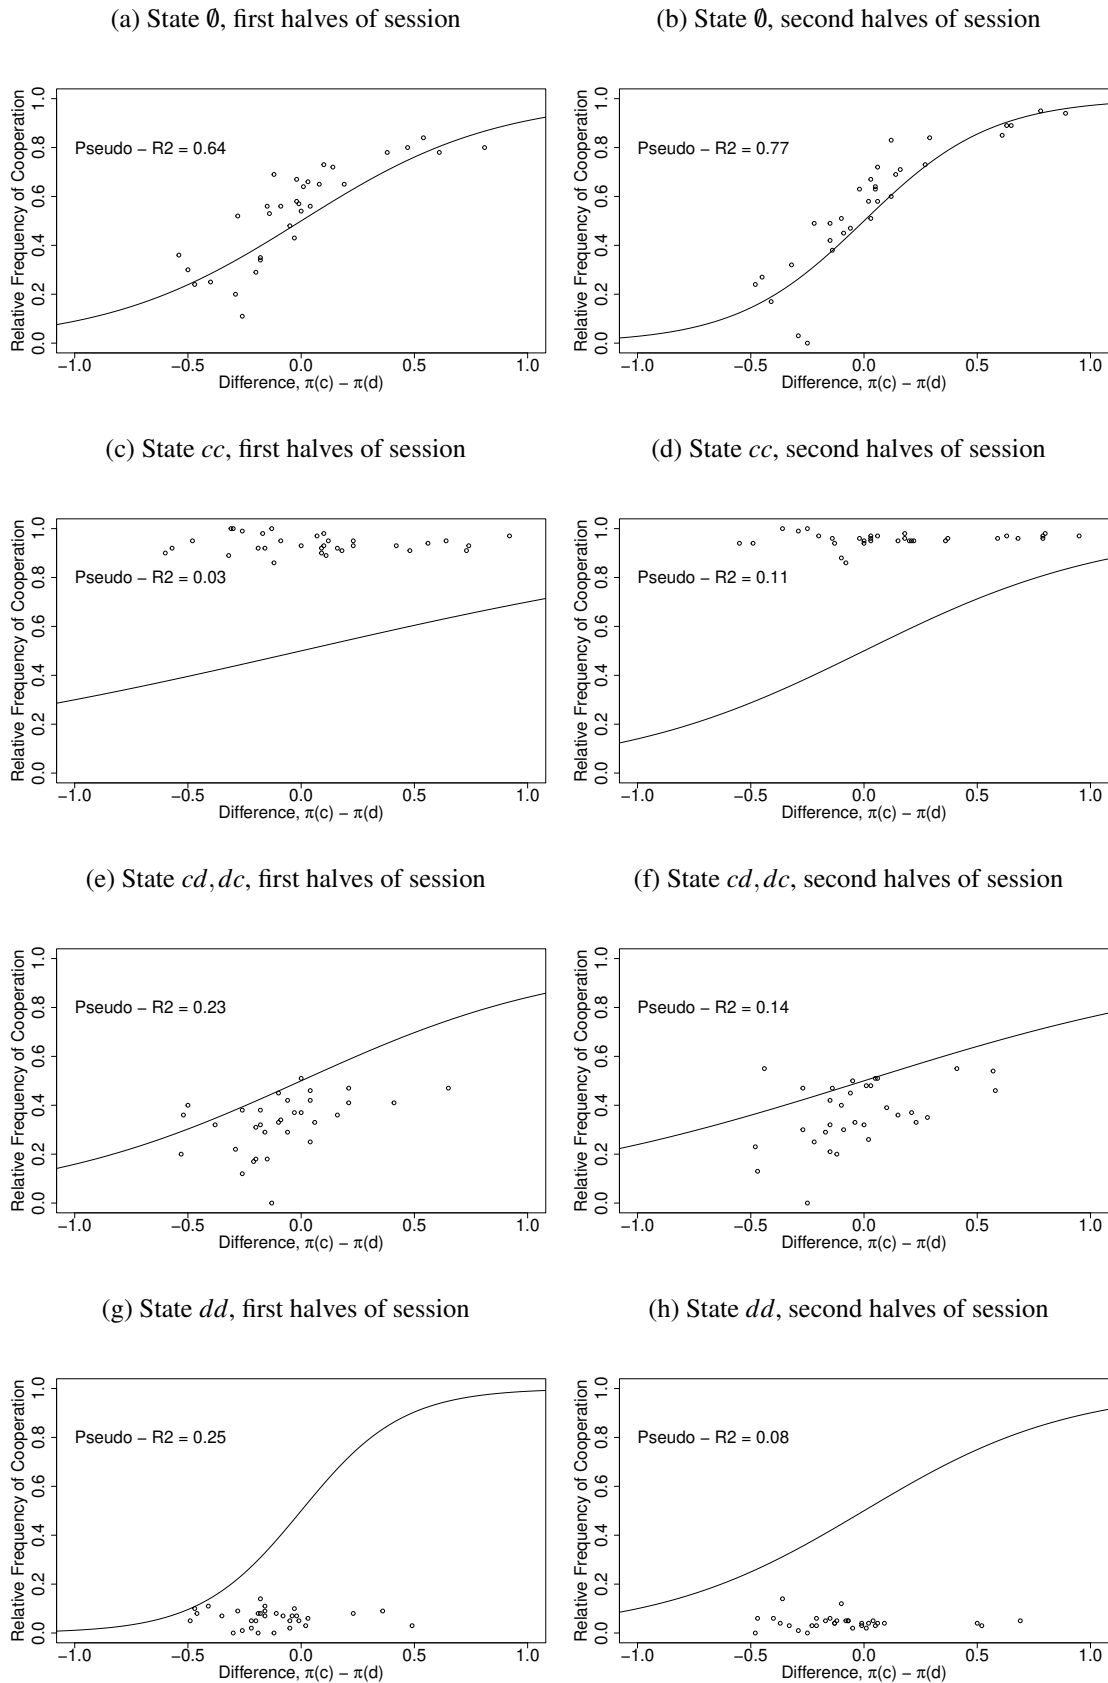
Question 3. *Can the strategies played be predicted from parameters of the game?*

4 Approach towards strategy estimation

In this section, we lay out our econometric approach and provide simulation results on the econometric identification of strategy types.

Econometric approach and identification The general understanding of behavior in repeated games experiments is that there is a finite number of subject types and that each type is

Figure 1: Relation of monetary incentives and cooperation rates across states (naive beliefs)



Note: The expected payoff cooperating in that state $\hat{\pi}(c)$, the expected payoff of defecting in that state $\hat{\pi}(d)$, the “predicted” probability of cooperation based on the logistic regression of cooperation rates on monetary incentive $\hat{\pi}(c) - \hat{\pi}(d)$, and the absolute deviation of that prediction.

represented by a specific strategy. Since we observe actions instead of strategies (in hot-play experiments), types are not directly observable. Formally, let K denote the finite set of types, where type $k \in K$ plays strategy σ_k and has population share ρ_k . Our data set is denoted O where $o_{s,t}$ denotes the action of subject $s \in S$ in round t of their experimental session. The probability of observing choice $o_{s,t}$, conditional on subject s being of type k is denoted as $\Pr(o_{s,t}|\sigma_k)$. Hence, the probability that s generates the observations $o_s = \{o_{s,t}\}_t$ conditional on being of type k is

$$\Pr(o_s|\sigma_k) = \prod_t \Pr(o_{s,t}|\sigma_k),$$

and unconditionally, using the prior type shares $\rho = \{\rho_k\}_k$, the probability that subject s generates the observations o_s is

$$\Pr(o_s|\sigma, \rho) = \sum_{k \in K} \rho_k \Pr(o_s|\sigma_k).$$

As McLachlan and Peel (2004) discuss in detail, by the far the most common (and easiest) consistent approach towards estimation of type shares ρ and strategies σ is maximum likelihood, where we maximize

$$\ln \mathcal{L}(\sigma, \rho|O) = \sum_{s \in S} \ln \Pr(o_s|\sigma, \rho), \quad (4)$$

usually using the EM algorithm. Maximum likelihood is standard practice also in the estimation of strategies in the repeated prisoner's dilemma, and to our knowledge, there is no alternative approach towards consistent estimation given the above finite mixture model of the subject pool. Note, that this method is well established in many fields of the empirical economics literature, including literature on lifetime decision making and public policy for a long time and more recently experimental economics.¹³

A subject using a *pure strategy* acts equivalently whenever a given state is reached and she uses the same pure strategy across all supergames in the considered data set (say, first halves of sessions). This is directly represented by the above model. A subject using a *behavior strategy* is also assumed to use the same strategy across all supergames in the considered

¹³The approach of using mixture models in order to uncover decision rules in experimental data has been established by Stahl and Wilson (1994) and El-Gamal and Grether (1995) and subsequently used in many analyses of level- k reasoning and stochastic choice, see e.g. Houser and Winter (2004) and Houser et al. (2004), to unravel individual decision rules. A special case of finite mixture modeling is the Strategy Frequency Estimation Method (SFEM) employed by Dal Bó and Fréchet (2011), Fudenberg et al. (2012), Rand et al. (2015), Dal Bó and Fréchet (2019), Fréchet and Yuksel (2017).

data set and thus is captured by the above model. A subject using a *mixed strategy* uses a pure strategy within supergames but randomizes over pure strategies prior to supergames. This can be captured straightforwardly in finite mixture models by assuming that subjects are represented by different agents in different supergames, where the likelihoods are aggregated over agents instead of subjects, which we will call random-switching model for clarity.

Can we reliably identify strategy types? We estimate the strategy proportions and parameters (such as tremble probabilities) by maximum likelihood and evaluate model differences by the robust Schennach-Wilhelm likelihood-ratio tests, which guarantees reliable identification for large samples. Our sample is comparably large, but is it large enough and more generally, what is the probability of misidentification in standard experimental samples? In order to answer these questions, we simulate populations that play pure (or, mixed) strategies exactly in the proportions estimated in DF18—for all 17 treatments analyzed in DF18. Is our maximum-likelihood approach sufficiently robust to correctly identify the type of strategy (pure, mixed, behavior) if the sample size is small (50 subjects)?

Appendix A.2 provides a detailed analysis answering this question, but the key results for our purpose are provided in Table 3. If subjects play pure strategies as assumed in DF18, in the proportions estimated in DF18, then our maximum very robustly identifies these strategies as pure strategies even in small samples of just 50 subjects and 20 observations per subject past round 1 (93% probability of correct identification, averaged across all 17 treatments). Mistaken identification is in almost all cases in favor of mixed strategies (as opposed to behavior strategies), and our “data mining” approach further benefits pure and mixed strategies (as discussed below). Table 3 also shows that mixed strategies are identified with very high reliability across conditions. Combined, these results very clearly demonstrate that populations with subjects playing pure strategies as estimated in DF18, or similar populations playing mixed strategies, will not be misidentified as populations playing behavior strategies.

Importantly, this does not stand in contrast to the estimates of Romero and Rosokha, who simulate populations where no agent plays semi-grim behavior strategies and then estimate population shares of a set of strategies including semi-grim. They find that a positive share of subjects is estimated to play semi-grim subjects although actually no subject did so. The explanation for the misidentification is rather simple: many of the subjects play strategies that are not in the set of candidate strategies considered by Romero and Rosokha, so they must be misclassified, and some of them end up being misclassified as playing semi-grim. This is not a concern in our simulations, which builds on Result 6 of DF18 that “Three

Table 3: Identification of strategies with very few observations in simulated populations with type shares of the 17 different treatments estimated by DF18

	Type proportions following DF18, Tab. 10 (σ_{AD} , σ_{Grim} , σ_{TFT} , σ_{AC} , σ_{STFT})	In 100 simulation runs, identified as		
		Pure strategies	Mixed strategies	AD + Semi-Grim
<i>Simulated agents play: Pure strategies ("No switching")</i>				
Dreber et al. (2008), T1	(0.64, 0.07, 0.15, 0, 0.14)	100	0	0
Dreber et al. (2008), T2	(0.3, 0.21, 0.4, 0, 0.09)	100	0	0
Dal Bo and Frechette (2011), T1	(0.91, 0, 0.07, 0, 0.02)	99	1	0
Dal Bo and Frechette (2011), T2	(0.76, 0, 0.06, 0, 0.08)	94	6	0
Dal Bo and Frechette (2011), T3	(0.49, 0, 0.24, 0.01, 0.04)	100	0	0
Dal Bo and Frechette (2011), T4	(0.66, 0, 0.23, 0, 0)	100	0	0
Dal Bo and Frechette (2011), T5	(0.11, 0.04, 0.21, 0, 0.08)	100	0	0
Dal Bo and Frechette (2011), T6	(0, 0.02, 0.55, 0.02, 0)	22	78	0
Fudenberg et al. (2012)	(0.06, 0.12, 0.15, 0.24, 0)	83	4	13
Rand et al. (2015)	(0.18, 0.43, 0.27, 0, 0.05)	98	0	2
Frechette and Yuksel (2014)	(0.14, 0.32, 0.39, 0, 0.02)	98	0	2
Dal Bo and Frechette (2015), T1	(0.53, 0.06, 0.05, 0, 0.14)	100	0	0
Dal Bo and Frechette (2015), T2	(0.25, 0.36, 0.19, 0.03, 0.03)	96	0	4
Dal Bo and Frechette (2015), T3	(0.47, 0.1, 0.1, 0.02, 0.12)	100	0	0
Dal Bo and Frechette (2015), T4	(0.12, 0.35, 0.3, 0.08, 0)	94	0	6
Dal Bo and Frechette (2015), T5	(0.14, 0.17, 0.39, 0.02, 0.07)	100	0	0
Dal Bo and Frechette (2015), T6	(0.22, 0.06, 0.25, 0, 0)	100	0	0
Probability of identification overall		0.932	0.052	0.016
<i>Simulated agents play: Mixed strategies ("Random switching")</i>				
Dreber et al. (2008), T1	(0.64, 0.07, 0.15, 0, 0.14)	0	100	0
Dreber et al. (2008), T2	(0.3, 0.21, 0.4, 0, 0.09)	0	100	0
Dal Bo and Frechette (2011), T1	(0.91, 0, 0.07, 0, 0.02)	0	100	0
Dal Bo and Frechette (2011), T2	(0.76, 0, 0.06, 0, 0.08)	0	100	0
Dal Bo and Frechette (2011), T3	(0.49, 0, 0.24, 0.01, 0.04)	0	100	0
Dal Bo and Frechette (2011), T4	(0.66, 0, 0.23, 0, 0)	0	100	0
Dal Bo and Frechette (2011), T5	(0.11, 0.04, 0.21, 0, 0.08)	0	100	0
Dal Bo and Frechette (2011), T6	(0, 0.02, 0.55, 0.02, 0)	2	98	0
Fudenberg et al. (2012)	(0.06, 0.12, 0.15, 0.24, 0)	0	100	0
Rand et al. (2015)	(0.18, 0.43, 0.27, 0, 0.05)	0	100	0
Frechette and Yuksel (2014)	(0.14, 0.32, 0.39, 0, 0.02)	0	100	0
Dal Bo and Frechette (2015), T1	(0.53, 0.06, 0.05, 0, 0.14)	0	100	0
Dal Bo and Frechette (2015), T2	(0.25, 0.36, 0.19, 0.03, 0.03)	0	100	0
Dal Bo and Frechette (2015), T3	(0.47, 0.1, 0.1, 0.02, 0.12)	0	100	0
Dal Bo and Frechette (2015), T4	(0.12, 0.35, 0.3, 0.08, 0)	0	100	0
Dal Bo and Frechette (2015), T5	(0.14, 0.17, 0.39, 0.02, 0.07)	0	100	0
Dal Bo and Frechette (2015), T6	(0.22, 0.06, 0.25, 0, 0)	0	100	0
Probability of identification overall		0.001	0.999	0

Note: Analysis based on simulated data sets comprising 50 subjects and 20 observations (past round 1) per subject, reporting the frequency of identification of the different strategy classes using our econometric methodology (where identification of either pure, mixed or behavior strategies is based on the ICL-BICs of the corresponding no-switching, random-switching, or semi-grim structures).

strategies account for most of the data: AD, grim, and TFT” (page 18), to which we add two more strategies (AC and S-TFT) and extensive data mining to be on the safe side, implying that very few pure-strategy subjects should not be accounted for in our analysis. Based on Result 6 of DF18 and our simulation results, misclassification of pure-strategy subjects as playing semi-grim strategies therefore appears to be practically impossible in our analysis.

Markov-switching As indicated, our analysis extends prior work by allowing also for Markov-switching models towards strategy selection, which generalize the finite-mixture and random-switching models used in previous analyses of repeated game strategies. They allow us to capture a potentially heterogeneous group of agents (in our case, subjects potentially playing different strategies), where each agent is characterized by a “state of mind” (the strategy to be played), and agents may change their states of mind over the course of time, but both states and transitions are latent and thus not directly observable. Let us refer to Ansari et al. (2012), Breitmoser et al. (2014) and Shachat et al. (2015) for earlier applications in behavioral analyses. The identifying assumption is that transitions between supergames follow a Markov process, that is transitions can only depend on the strategy in the previous supergame. This generalizes the finite mixture model, with degenerate transition probabilities, and the random switching model, where the strategy choices are independent of the strategy choice in the previous supergame.¹⁴ Note that Markov switching does not represent a pure or mixed strategy in the standard sense, but we include it in the estimation in order to help variations of pure strategies achieve the best-possible fit in our estimation of the upper bound for their goodness-of-fit.

As above, the estimation of Markov-switching models proceeds by maximum likelihood. Formally, let G denote the set of supergames and let $\kappa \in K^G$ denote a transition paths across types. For example, by $\kappa = (\kappa_1, \kappa_2, \dots)$ the subject in question is of type $\kappa_1 \in K$ in supergame 1, of type κ_2 in supergame 2, and so on. Further, let $g_s(t)$ denote the number g of the supergame that subject s plays when she is in round t of her experimental session. Thus, $\kappa_{g_s(t)}$ denotes the type according to transition path κ in the supergame where subject s is in round t of her experimental session. Hence, the probability that s generates the observations $o_s = \{o_{s,t}\}_t$ conditional on following the transition path $\kappa \in K^G$ is

$$\Pr(o_s | \sigma, \kappa) = \prod_t \Pr(o_{s,t} | \sigma_{\kappa_{g_s(t)}}),$$

¹⁴The latter, with the additional assumption that choice probabilities for the strategies are constant over time, ensures identification of random-switching models.

and unconditionally, using $\rho(\kappa)$ to denote the prior probability of transition path $\kappa \in K^G$, the probability that subject s generates the observations o_s is

$$\Pr(o_s|\sigma, \rho) = \sum_{\kappa \in K^G} \rho(\kappa) \Pr(o_s|\sigma_\kappa).$$

As indicated, we assume that type transitions follow a Markov process, the specification of which defines the path probabilities $\rho(\kappa)$. Consistent estimation of the model parameters is achieved by maximum likelihood, and to this end, we use the standard Baum-Welch algorithm (Bilmes et al., 1998).

Mechanically, the more complex a model (the more parameters and the more subject types), the larger a model’s capacity to fit the data—and implicitly, the larger its fallacy to overfit the data. This is conventionally captured by evaluating model adequacy based on information criteria such as the Bayes information criterion (BIC), which penalize for the degrees of freedom in a theoretically adequate manner. Mixture and switching models additionally contain freedom in defining the number of subject types, which provides an additional source for overfitting aside from the number of parameters used. Following Biernacki et al. (2000), we address these concerns by using the information-classification likelihood Bayes-information criterion (ICL-BIC), a criterion that penalizes both model complexity and the failure of the mixture model to provide a classification in well-separated strategy clusters.

We address the observation that modeling mixtures of pure, mixed, and behavior strategies induces sophisticated nesting structures between models, and the concern that indeed all models may be misspecified by evaluating model differences using the robust Schennach-Wilhelm likelihood ratio tests (Schennach and Wilhelm, 2017).

Finally, we allow for stochastic choice in the form of trembles (after all histories of play) following Harless and Camerer (1994), i.e. in each round the minimal probability of any action is equal to $\gamma \geq 0$ where γ is a free (noise) parameter in the estimation. In our context, this approach is econometrically equivalent to a logistic-error approach, which is traditionally used in random utility modeling as opposed to strategy estimation.¹⁵ In strategy estimation, the perturbation approach has the advantage that it does not perturb choice

¹⁵To see the equivalence in strategy estimation, consider a type playing a pure strategy. In states where the strategy prescribes cooperation with probability 1, the logistic-error approach prescribes cooperation with probability $1/(1 + \exp\{\lambda \cdot (0.5 - 1)\})$, which implies the tremble probability $\gamma = 1 - 1/(1 + \exp\{\lambda \cdot (0.5 - 1)\})$. In states where cooperation is prescribed with probability 0, the logistic cooperation probability is $1/(1 + \exp\{\lambda \cdot (0.5 - 0)\})$, which yields the same value for the tremble probability $\gamma = 1/(1 + \exp\{\lambda \cdot (0.5 - 0)\})$. Hence, there is a bijection between logistic precision λ and tremble probability γ , implying the econometric equivalence.

probabilities of subjects that are already randomizing, which enables direct estimation of strategies. Importantly, it is customary to assume that trembles, i.e. unintended deviations from pure strategies, are symmetric across states. Behavior strategies in general, and semi-grim strategies in particular, are not symmetric in this sense: they allow for deviations from purity to be more pronounced in some states than in others. This is the statistical foundation for a distinction of pure strategies with trembles and behavior strategies. If the symmetry condition is violated significantly, then behavior strategies fit better than pure strategies, and if the violations of the symmetry are mild, then the more parsimonious pure strategies fit better (by our information criterion). Arguably, if the tremble probabilities exceed the .05 level, they induce a “significant” deviation from pure strategies, and the resulting strategies might be labeled behavior strategies even if trembles are symmetric across states. To give the pure-strategy conjecture the best-possible chance, however, we will not restrict the magnitude of tremble probabilities in any way.

5 Analysis

In this section, we identify the set of candidate strategies and data mine for the best possible (post-hoc) mixtures of (generalized) pure strategies for *each treatment*. We will not penalize the model for data mining best mixtures but treat the resulting mixtures treatment-by-treatment as if they had been hypothesized ex-ante. As we discuss below, this provides us with an upper bound for the goodness-of-fit of pure and mixed strategies, which we will compare to a simple model that contains only defectors playing AD and two types of (potential) cooperators playing unrestricted memory-1 behavior strategies. Due to the one-sided data mining, involving optimizing the post-hoc mixture of pure and mixed strategies, augmented by Markov-switching strategies, measured against a simple mixture of AD and two behavior strategies, this analysis is heavily lopsided in favor of modeling behavior using pure and mixed strategies. In this sense, we give the pure- and mixed-strategy conjectures the best possible chance.

We then estimate the number of subject types and the strategies played in both a top-down and a bottom-up approach towards model selection. The top-down approach starts with the general model and iteratively eliminates insignificant components, while the bottom-up approach starts with a basic model and iteratively adds model components identified as significant. Both approaches will converge to the same model distinguishing defectors playing AD from cautious and strong cooperators playing semi-grim strategies. Section C in the appendix demonstrates robustness to longer memory lengths by showing that model adequacy

does not improve by equipping subjects with memory-2, neither for (generalizations of) pure strategies nor for semi-grim. That is, while increasing memory length slightly improves the goodness-of-fit, this increase does not make up for the increased complexity of strategies as evaluated using the Bayesian information criterion.

Candidate strategies Our objective is to obtain an upper bound for the goodness-of-fit of (generalized) pure and mixed strategies. To this end, we first need to define the candidate set of pure strategies that are being considered, as it is obviously impossible to consider all pure strategies explicitly. Nor is it necessary, however, as many strategies are simply not capturing subjects' behavior even by classifications focusing on pure strategies only. To begin with, previous studies have found that in the repeated PD with perfect monitoring, strategies with memory-2 or more seem to be of little relevance statistically (Breitmoser, 2015; Dal Bó and Fréchette, 2019). In the appendix, we report a robustness test of our results with respect to memory-2 and similarly find memory-2 to be statistically insignificant. For these reasons, and following the existing literature, we shall focus on memory-1 strategies in this paper. This leaves us with $2^5 = 32$ pure strategies, but many of them are behaviorally equivalent. For example, always defect $(\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0, 0, 0, 0, 0)$ is behaviorally equivalent to suspicious Grim $(0, 1, 0, 0, 0)$ and suspicious AC $(0, 1, 1, 0, 0)$. When we eliminate these duplicates (see Appendix A.1 for details) and the strategically dubious strategies predicting cooperation in some state but defection in cc , this leaves us with the 10 pure strategies shown in Table 4. Not all of them seem to be equally plausible candidate strategies, however. In order to assess their potential relevance in capturing behavior observed in the experiments we analyze, we next check for how many subjects these strategies best explain their play. That is, we classify subjects into the best fitting strategy at the individual level.

Our classification is based on maximum likelihood¹⁶, which is congruent with the other estimates reported here, and essentially asks which strategy reports the observed choices of a given subject with the highest probability. We report the results for two possible levels of tremble probabilities $\gamma = 0.01$ and $\gamma = 0.03$, which cover the relevant range based on noise levels estimated in Breitmoser (2015), in Table 4. But before we discuss the results, a word of caution is in order. Despite our maximum-likelihood approach towards the classifications, the reported type shares are not the maximum-likelihood estimates of the type shares in the population. The latter estimates can only be obtained by estimating all type shares jointly, i.e. by estimating the aforementioned finite-mixture model for the entire population. This

¹⁶We are grateful to a referee for the suggestion to classify subjects in order to justify our choice of candidate strategies. The suggestion was a classification based on maximum likelihood, which we adopt here, but appendix 9 reports robustness checks showing that Euclidean distance yields similar classifications.

joint estimation implicitly accounts for the relative weights of evidence, as for example subjects that are just between Grim and AD, with a slight advantage for Grim, will count in the relatively correct way for Grim and AD weights in the population-level estimation reported below, while they count fully as Grim and zero as AD in the subject-level estimation reported in Table 4. With this being the main reason for the population-level (“finite-mixture”) estimation, the population-level estimation is also the only consistent approach to simultaneously estimate strategy parameters such as tremble probabilities and cooperation probabilities in behavior strategies (although the results are fairly robust in this respect).

Table 4: Share of strategies by maximum likelihood classification with and without allowing for behavior strategies

	Strategy definition	Best fitting strategy		Robustness
	$(\sigma_0, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$	Shares	Pure only	Shares
AD	(0, 0, 0, 0, 0)	0.21	0.32	0.23
S-TFT	(0, 1, 0, 1, 0)	0.02	0.13	0.03
S-WSLS	(0, 1, 0, 0, 1)	0	0	0
S-AC1	(0, 1, 1, 1, 0)	0.01	0.04	0.01
S-AC2	(0, 1, 1, 0, 1)	0	0	0
S-AC3	(0, 1, 1, 1, 1)	0	0	0
Grim	(1, 1, 0, 0, 0)	0.05	0.22	0.07
AC	(1, 1, 1, 0, 0)	0.02	0.07	0.03
TFT	(1, 1, 0, 1, 0)	0.08	0.21	0.1
WSLS	(1, 1, 0, 0, 1)	0	0.01	0
Cautious SG	(.3, .9, .3, .3, .1)	0.35	–	0.31
Strong SG	(.9, .9, .3, .3, .1)	0.26	–	0.21
<i>Sum of shares:</i>				
All		1	1	
AD + C/S SG		0.82	–	0.75
AD + Grim + TFT		0.34	0.74	0.37

Note: Strategy shares by maximum likelihood classification allowing for 1% tremble probabilities (“Best fitting strategy”) and 3% tremble probabilities (“Robustness”). The set includes all possible memory-1 strategies (see Appendix A.1 for details) and two behavior strategies cautious and strong cooperators derived from Breitmoser (2015). If two strategies fit equally well, we prioritize known strategies over unknowns (S-ACx, S-WSLS). And if then still two strategies fit equally well we assign the subject half to each of the strategies. For average actual play of subjects assigned to these strategies, see Tables 7 and 8 in the Appendix.

As shown in Table 4, the results are robust to the choice of tremble probabilities considered here. Hence, let us focus on the best-fitting strategies by the $\gamma = 0.01$ classification. In the classification restricted to pure strategies, 74% of subjects are assigned to play either AD, TFT or Grim, with 22% to Grim and 21% to TFT, replicating Dal Bó and Fréchette

(2018, Result 6).¹⁷ Once we allow for the two types of semi-grim strategies discussed above, the weights of Grim and TFT drop to 5% and 8%, respectively, while the two semi-grim strategies represent the best fit for 61% of the subjects (35% cautious and 26% strong semi-grim)—without any adjustments to the coarsely rounded cooperation probabilities (.9, .3, .1) mentioned in Breitmoser (2015). These observations again foreshadow the results of the population-level analysis allowing for general behavior strategies: any restriction to pure strategies is behaviorally inadequate and will be statistically rejected.

At present, however, we are defining the the set of candidate strategies, which we consider to be played in pure, mixed, and/or generalized form, as the set of strategies attracting at least 0.02 weight in the above classification, i.e. AD, Grim, TFT, STFT and AC, which replicates the observation of DF18 (Footnote 38) that these five strategies are the most relevant pure strategies. The candidate WSLS, which we hypothesized to be a promising candidate as a number of studies established its evolutionary robustness (Nowak and Sigmund, 1993; Imhof et al., 2007) receives only a negligible share below 1%.¹⁸

In order to extend the scope of pure and mixed strategies, we extend this set of strategies by adding generalized versions. These generalized versions have a free parameter per strategy prescribing first-round cooperation rates σ_0 , thus allowing subjects' first-round cooperation rates to be different from 0 in AD and STFT and different from 1 in the other strategies, not just in the form trembles (which are the same for all strategies), but individually for each strategy type. Note that round-1 randomization has trivial mixed-strategy representations, since mixed strategies also randomize in round 1. This way, we allow for some types to play pure and other types to play (certain) mixed strategies randomizing over cautious and cooperative play in round 1. The definition of the continuation behavior remains unchanged, such that $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) \in \{0, 1\}^4$ aside from trembles. We refer to these strategies as *generalized pure strategies*. In addition, we will of course consider the pure strategies in their original form (with trembles), thereby covering the possibility that in at least some treatments the generalizations do not improve the goodness-of-fit, allowing us to post-hoc save parameters.

Pure, mixed or behavior strategies? In order to answer this question, we now estimate an upper bound of the goodness-of-fit of pure strategies. As indicated, the models that we consider for estimating this upper bound go far beyond what one would usually consider to

¹⁷The actual average play of subjects assigned to these strategies is far from being pure, though (see Table 7 in the Appendix).

¹⁸Results with WSLS are available in previous working paper versions and upon request.

be pure strategies, but we relax some restrictions to put pure and mixed strategies at an advantage. In combination with the way we pick the best specifications at the treatment-level, preceded by several optimization steps, deliberately puts them in a head start position compared to what is generally considered best practice in model selection (hence, upper bound). This is true particularly in comparison to the simple mixture of AD and two unrestricted behavior strategies without post-hoc model selection.

Given the set of candidate strategies defined above, our approach towards data mining pure-strategy mixtures across treatments is as follows. First, we evaluate independently for each treatment which mixture of pure or generalized pure strategies best captures behavior. That is, we determine for each treatment, which combination of *pure* strategies fits best, which combination of *generalized pure strategies* fits best, and which of the best combinations fits best. Following the pure-strategy conjecture, we assume the best combination always contains at least TFT, AD, and Grim. We add AC, STFT, or pure noise players (randomizing 50-50 in all states) when this improves the goodness-of-fit by ICL-BIC (see above). Thus, we choose the best out of 9 as promising conjectured memory-1 mixtures, for each of the 32 treatments and each of the two half-sessions independently.¹⁹ In total, we therefore evaluate 9^{32} models per level of experience and afterwards pick the best-fitting model by ICL-BIC. Second, we do all of this separately for the three “switching models” designed to capture the three possibilities of strategy switching between supergames: "No Switching" (pure strategy), "Random Switching" (mixed strategy), and "Markov Switching" (strategy switching between supergames follows a Markov process).

The results for each of the three switching models are reported in columns 2-5 of Table 5. The leftmost column contains the results for the baseline model comprising AC, AD, TFT, Grim, and STFT without data mining, which can serve as a reference for how much of the goodness-of-fit is due to data mining. The three columns “No Switching”, “Random Switching” and “Markov Switching” contain the ICL-BICs of the best fitting combination of pure and generalized pure strategies—optimally picked by treatment, but for the sake of readability, we report ICL-BICs aggregated by experiment.²⁰

The random switching model in column 3 of Table 5 capturing mixed strategies generally fits worst, by the enormous amount of more than 2000 points on the log-likelihood

¹⁹For each of the two classes of strategies (pure and generalized pure), we consider mixtures containing AD, TFT and Grim and in addition either (i) no other strategy, (ii) AC, (iii) STFT, and (iv) AD + STFT. This makes 8 combinations in total. In addition, in the case of pure strategies, we allow for a mixture containing noise players (randomizing 50-50 in all states) as type besides AD, TFT and Grim.

²⁰Treatment-wise ICL-BICs are provided upon request. Each entry in the aggregated table represents the sum of ICL-BICs of the best out of 9 models for each respective treatment.

scale. This shows that subjects are reasonably consistent in their strategy choice. The no-switching model capturing pure strategies (column 2) fits worse than the Markov-switching model (column 4) in the first halves of sessions, but weakly better in the second halves of sessions. If these models captured behavior well, this could suggest that subjects initially experiment with different pure strategies, though not randomly, as in mixed strategies, but systematically, as in a stochastic Markov process, to then converge to individual choices for strategies as the session proceeds.

Overall, the aggregate effect achieved by data mining for the best-fitting combination of (generalized) pure strategies and switching model is highly significant in relation to the baseline model. Modeling the behavior of inexperienced subjects (first halves of sessions), our generalizations and data mining combined yield a gain of 1000 points on the log-likelihood scale, comparing the baseline model to the best-fitting Markov switching models, and modeling experienced subjects (second halves), generalization and data mining combined yield a gain of about 700 points compared to the baseline model. Since these scores do not account for the degrees of freedom inherent in the model selection during data mining, they do not imply that the baseline model has to be rejected, but they clearly show that our approach yields an enormous improvement in fit over the memory-1 strategies identified as the candidate strategies and in the literature. Further, since we attempted to include all specifications that may be considered compatible with either the pure- or the mixed-strategy conjecture, and picked the best one for each treatment, we can consider this data-mined specification to be a generous upper bound of the adequacy of these memory-1 models to describe behavior.

Second, this upper bound, reported in column 5 (“Best Switching”) of Table 5, allows us to test the pure- and mixed-strategy conjectures, extended by Markov-switching, against the behavior-strategy conjecture. While the behavior-strategy conjecture suggests that the behavior of cooperating subjects is well-described using semi-grim strategies after round 1, we will initially use no such prior insights and estimate strategy mixtures involving unrestricted memory-1 strategies $(\sigma_0, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$. Specifically, we estimate mixtures dubbed “2×P5+AD” involving an AD type and two such unrestricted behavior-strategy types (with five free parameters each, hence “P5”). This mixture generalizes the three-type mixture AD+TFT+Grim hypothesized by DF18, simply by lifting all restrictions, without imposing ex-ante plausible restrictions, and without any model selection to improve the information criterion. For this reason, we refer to it as a lower bound of the goodness-of-fit attainable when considering behavior strategies. We will explore possible improvements incorporating semi-grim structure below.

Specifically, we compare the simple three-type model with an invariant and sub-optimal

Table 5: **Best mixtures of pure or generalized strategies in relation to behavior strategies.** Strategy mixtures are estimated treatment-by-treatment. The resulting ICL-BICs are pooled for experiments and overall (less is better, relation signs point to better models)

	Best mixture of pure or generalized strategies												
	Baseline Model	No Switching	Random Switching	Markov Switching	Best Switching	Unrestr Beh 2×P5+AD	Best Mixture Best Switching By Treatment						
Specification													
# Models evaluated	1	9 ³²	9 ³²	9 ³²	3 × 9 ³²	1	27 ³² ≈ 10 ⁴⁶						
# Pars estimated (by treatment)	5	48	48	180	276	13	276						
# Parameters accounted for	5	3–10	3–10	12-35	3–30	13	3–30						
First halves per session													
<i>Aoyagi and Frechette (2009)</i>	888.46	≈	843.08	≈	834.4	≈	845.5	≈	845.5	≫	744.76	<	834.4
<i>Blonski et al. (2011)</i>	1126.33	≫	1069.54	≈	1103.29	≪	1220.75	≈	1220.75	≈	1209.41	≫	1060.2
<i>Bruttel and Kamecke (2012)</i>	827.86	≈	821.99	≈	835.56	>	785.49	≈	785.49	≈	759.46	≈	785.49
<i>Dal Bó (2005)</i>	639.5	≈	623.19	≪	674.07	≈	641.52	≈	641.52	≈	609.66	≈	619.34
<i>Dal Bó and Fréchette (2011)</i>	7143.23	≫	6874.99	≪	7459.16	≫	6378.54	≈	6378.54	≈	6273.56	≈	6378.54
<i>Dal Bó and Fréchette (2019)</i>	8590.44	≫	8367.55	≪	9152.85	≫	8181.42	≈	8181.42	≫	7775.32	≪	8161.75
<i>Dreber et al. (2008)</i>	840.44	>	789.22	<	863.52	≫	744.21	≈	744.21	≈	767.3	≈	744.21
<i>Duffy and Ochs (2009)</i>	1401.73	≈	1396.68	<	1467.36	≫	1372.98	≈	1372.98	≈	1345.12	≈	1372.98
<i>Fréchette and Yuksel (2017)</i>	314.71	≈	300.87	<	339.64	>	297.74	≈	297.74	≈	285.33	≈	297.74
<i>Fudenberg et al. (2012)</i>	437.5	≈	437.5	≈	432.38	≈	435.85	≈	435.85	≫	372.32	≪	432.38
<i>Kagel and Schley (2013)</i>	2660.58	≈	2660.58	≪	2992.72	≫	2439.06	≈	2439.06	≈	2398.74	≈	2439.06
<i>Sherstyuk et al. (2013)</i>	1318.22	≈	1299.14	≪	1450.71	≫	1274.09	≈	1274.09	>	1186.92	<	1274.09
Pooled	26371.37	≫	25710.49	≪	27804.74	≫	25307.72	≈	25307.72	≫	24202.07	≪	24919.14
Second halves per session													
<i>Aoyagi and Frechette (2009)</i>	530.1	≈	492.28	≈	484.04	≈	482.66	≈	492.28	≫	408.59	<	482.66
<i>Blonski et al. (2011)</i>	1473.77	≫	1373.41	<	1479.64	≪	1592.81	≈	1373.41	<	1458.47	≫	1369.08
<i>Bruttel and Kamecke (2012)</i>	522.14	≈	493.79	≪	611.49	≫	498.18	≈	493.79	≈	471.73	≈	493.79
<i>Dal Bó (2005)</i>	715.1	≈	712.63	≪	782	>	740.37	≈	712.63	>	687.86	<	712.63
<i>Dal Bó and Fréchette (2011)</i>	5355.14	≫	5045.08	≪	6396	≫	4860.82	≈	5045.08	≫	4493.1	≪	4820.25
<i>Dal Bó and Fréchette (2019)</i>	8273.95	>	8107.4	≪	9413.23	≫	7837.73	≈	8107.4	≫	7152.05	≪	7785.24
<i>Dreber et al. (2008)</i>	603.7	≈	580.69	<	661.75	≫	558.39	≈	580.69	>	519.56	≈	558.39
<i>Duffy and Ochs (2009)</i>	1914.65	≈	1910.09	≈	1992.7	≫	1883.18	≈	1910.09	≫	1598.03	≪	1883.18
<i>Fréchette and Yuksel (2017)</i>	458.49	≈	433.18	<	474.9	≈	429.95	≈	433.18	≫	382.11	≪	429.95
<i>Fudenberg et al. (2012)</i>	505.1	≈	505.1	≈	540.33	≈	523.98	≈	505.1	≫	421.81	<	505.1
<i>Kagel and Schley (2013)</i>	1681.03	≈	1681.03	≪	2347.19	≫	1621.42	≈	1681.03	≫	1527.49	<	1621.42
<i>Sherstyuk et al. (2013)</i>	901.02	≈	890.27	≪	1137.49	≫	896.87	≈	890.27	>	815.26	<	890.27
Pooled	23116.57	≫	22462.25	≪	26533.34	≫	22747.03	≈	22462.25	≫	20410.22	≪	22069.97

Note: Relation signs encode p -values of Schennach-Wilhelm likelihood-ratio tests where $<$, $>$ indicate rejection of the Null of equality at $p < .05$ and \ll , \gg indicating $p < .002$, which implements the Bonferroni correction of 24 simultaneous tests per hypothesis. “No Switching” assumes that subjects chooses a strategy prior to the first supergame and plays this strategy constantly for the entire half session. “Random Switching” assumes that subjects randomly chooses a strategy prior to each supergame (by i.i.d. draws), and “Markov Switching” allows that these switches follow a Markov process.

number of 13 free parameters per treatment (column 6: “Unrestr Beh, $2 \times P5 + AD$ ”), to the “Best Switching” model (column 5) that was post-hoc picked from 3×9^{32} models, after estimating 276 parameters for each of the 32 treatments, but without accounting for the degrees of freedom used in the model selection process (solely accounting for the 3–10 parameters of the best-fitting model that is finally used—in line with the data mining ideal).

Despite this abuse of statistical power, the behavior strategy mixture fits significantly better than the mined mixture of generalized pure or mixed strategies: it improves on the data-mined model by more than 1100 points in the first-halves of sessions and even by more than 2000 points in the second halves of sessions. Since AD players are contained in all models, this demonstrates that the behavior of subjects not playing AD—i.e. behavior of cooperating subjects—is much better described by behavior strategies than using any mixture of standard or generalized pure strategies. Indeed, the effect gets more pronounced as subjects gain experience. This is substantial and perhaps surprising, but in the end, it is simply a reflection of the deficiency of deterministic choice rules in capturing behavior as we observed already in the subject-level classification. A robustness check clarifying that this observation also holds true after accounting for memory-2 is reported in the appendix.

Third, we evaluate the arguably extreme model, which identifies the best-fitting combination of (generalized) pure strategies (out of 9 combinations) *and* the best-fitting switching model (out of 3) treatment by treatment without any consistency requirement. Thus, we choose the best-fitting model from 27 models for each treatment, amounting to the enormous selection of the best out of 27^{32} models across all experiments. Note that such analysis without imposing consistency requirements across treatments does not yield economically useful estimates, but if anything, this provides an even more generous upper bound of the behavioral content of pure and generalized pure strategies of memory-1. The results are reported in the right-most column (“Best Switching By Treatment”). In total, this exhaustively mined model still fits worse than the behavior-strategy mixture, by more than 700 points in first halves of sessions and by almost 1600 points in second halves of sessions.²¹ Hence, it is key to allow cooperating subjects to play strategies involving non-trivial randomization in strategy estimation, and restrictions to pure strategies are invalid. We summarize these observations as follows.

Result 1 (Question 1). *Cooperating subjects seem to use memory-1 behavior strategies. The upper bound of the goodness-of-fit that we can attain with pure or mixed strategies, even allowing for Markov-switching, is significantly below the (lower bound of) goodness-of-fit of*

²¹Section C in the appendix demonstrates that this result is robust to allowing for memory-2, where we find that memory-2 is overall insignificant.

a model allowing for cooperating subjects to play behavior strategies.

Heterogeneity of cooperators and strategy estimation Let us now examine to what extent the cooperating subjects are heterogeneous and if there is any structure in the behavior strategies they seem to be playing. We estimate the extent of heterogeneity in two ways, by a top-down approach and by a bottom-up approach. In the top-down approach, we start with a general model allowing for four different subject types (per treatment), one of which plays AD and three that play memory-1 behavior strategies without any restrictions, and successively impose restrictions until they reduce the goodness-of-fit in a statistically significant manner. In Table 6, we refer to the first model as “ $3 \times P5 + AD$ ”, where $P5$ indicates the unrestricted five-parameter behavior strategy. Note that we explicitly distinguish AD and behavior strategies, as the behavior strategy conjecture is that *cooperating* subjects play (semi-grim) behavior strategies.

In the bottom-up approach, we start at a highly restricted model with just 3 free parameters per treatment, to capture type shares and trembles (column 2: “Fixed SG, $1.5 \times SG + AD$ ”, Table 6)²², and successively lift restrictions until this stops improving the goodness-of-fit in a statistically significant manner. As starting point, we use the cautious and strong semi-grim strategies discussed above, which we identified to be played by 61% of subjects overall, which amounted to about 75% of cooperating subjects in the individual-level analysis.

Table 6 provides detailed information on a range of models connecting these two extremes, distinguishing either up to three cooperating types playing general behavior strategies or up to three types playing semi-grim (SG) strategies, which we define as strategies satisfying $\sigma_{cc} = 1 - \sigma_{dd} > \sigma_{cd} = \sigma_{dc}$. These intermediate models represent our prior hypothesis and allows us to implement the two approaches toward model selection. Finally, Table 6 includes a model with just one type of cooperating subjects (“ $P5 + AD$ ” in the right-most column), as a robustness check if this might suffice, but unsurprisingly, this model fits significantly worse.

Before delving into the details, one point is worth noting: Table 6 reports on a large range of models where cooperating subjects always are assumed to play unrestricted or semi-grim behavior strategies. All of these models improve on the best of the 10^{46} models assuming subjects play pure or generalized pure strategies (“Best Mixture, Best Switching” in the left-most column of Table 6). That is, our earlier result on the inadequacy of pure and generalized pure strategies is confirmed very robustly: whatever specification we use, allow-

²²Slightly abusing notation, 1.5 semi-grim types indicates that the two cooperating types have different cooperation probabilities in round 1 of each supergame but equivalent continuation strategies.

ing for two types of cooperating subjects to play flexible semi-grim or unrestricted behavior strategies fits behavior much better. Importantly, this would not be observed if the pure-strategy conjecture was empirically valid: Besides AD, our behavior-strategy estimation can accommodate cooperating subjects to play any mix of up to three cooperative strategies like TFT, Grim, and say WLSL, STFT or AC depending on treatment (in $3 \times P5 + AD$). If they actually did so, then the (generalized) pure strategy mixture would fit at least as well as the behavior strategy mixture without using as many free parameters, which improves the ICL-BIC score. And if there were more than three different types of cooperating subjects in some treatments, say TFT, Grim, STFT, and AC which the pure strategy mixture can capture very parsimoniously treatment-by-treatment, then the (generalized) pure strategies would fit substantially better than the behavior strategy models. This is not the case, however.

Now, in the top-down approach using “ $3 \times P5 + AD$ ” as starting point, we can analyze which form of heterogeneity is most suitable for describing behavior. Starting with four subject types seems to be sufficient ex-ante, and will turn out to be sufficient ex-post. In Table 6, the two right-most columns report on the adequacy of nested models that distinguish only two types or one type of cooperating subjects (besides the AD type). It turns out that distinguishing just two types of cooperating subjects (“ $2 \times P5 + AD$ ”) weakly improves on distinguishing three types, while models with just one cooperating type (“ $P5 + AD$ ”) fit significantly worse. The latter further corroborates that cooperating subjects are not homogeneous. To the left of column “ $3 \times P5 + AD$ ”, Table 6 details information on models assuming the cooperating subjects play semi-grim strategies rather than unrestricted memory-1 strategies. To be exhaustive, we consider models distinguishing three semi-grim types (“ $3 \times SG + AD$ ”), two semi-grim types (“ $2 \times SG + AD$ ”) and 1.5 semi-grim types (“ $1.5 \times SG + AD$ ”), besides the model with fixed semi-grim strategies (“Fixed SG, $1.5 \times SG + AD$ ”) defined above.

At this point, the discussion can be kept rather short as the results are fairly clear: All models distinguishing at least two types of cooperating subjects and flexible semi-grim behavior strategies fit about equally well. The differences between these models are at best weakly significant, while all of them fit significantly better than the model assuming cooperating subjects are homogeneous (“ $P5 + AD$ ”). Compared to the model specification with the fixed semi-grim strategies used above, the differences are insignificant in first halves of sessions but become significant in second halves of sessions. Initially, that is, cooperating subjects seem to be well-described by $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.9, 0.3, 0.3, 0.1)$, while their behavior becomes more nuanced and treatment-dependent as they gain experience. Further, the best-fitting mixtures generally involve semi-grim strategies, indicating that the semi-grim restrictions ($\sigma_{cc} = 1 - \sigma_{dd}$ and $\sigma_{cd} = \sigma_{dc}$) are statistically insignificant even in these large

Table 6: Examining heterogeneity of cooperating subjects and semi-grim structure of their strategies

	Best Mixture		Fixed SG		Treatment-specific SG specification										
	Best Switching		1.5×SG+ AD		1.5×SG+AD	2×SG+AD	3×SG+AD	3×P5+AD	2×P5+AD	P5+AD					
Specification															
# Models evaluated	27 ³² ≈ 10 ⁴⁶		1		1	1	1	1	1	1	1	1	1		
# Pars estimated (by treatment)	276		3		7	9	13	19	17	17	17	11	11		
# Parameters accounted for	3-30		3		7	9	13	19	17	17	17	11	11		
First halves per session															
<i>Aoyagi and Fréchette (2009)</i>	834.4	>	793.63	≈	792.28	≈	777.81	≈	782.63	>	742.29	≈	744.76	≈	744.06
<i>Blonski et al. (2011)</i>	1060.2	≈	1043.19	≪	1104.1	<	1138.64	≪	1236.36	≪	1333.92	≫	1209.41	≫	1104.45
<i>Bruttel and Kamecke (2012)</i>	785.49	≈	763.66	≈	771.14	≈	762.83	≈	748.06	≈	751.87	≈	759.46	≈	803.58
<i>Dal Bó (2005)</i>	619.34	≈	600.65	<	618.12	≈	601.72	<	627.73	≈	640.66	>	609.66	≈	620.39
<i>Dal Bó and Fréchette (2011)</i>	6378.54	≈	6458.01	≈	6352.64	≈	6304.98	≈	6198.13	≈	6217.26	≈	6273.56	≪	6553.24
<i>Dal Bó and Fréchette (2019)</i>	8161.75	≫	7912.58	>	7829.75	≈	7810.64	≈	7830.56	≈	7829.08	≈	7775.32	≪	7969.32
<i>Dreber et al. (2008)</i>	744.21	≈	774.76	≈	764.44	≈	763.51	≈	766.74	≈	766.69	≈	767.3	≈	783.46
<i>Duffy and Ochs (2009)</i>	1372.98	≈	1325.28	≈	1361.13	≈	1320.67	≈	1297.82	≈	1291.38	<	1345.12	≈	1361.84
<i>Fréchette and Yuksel (2017)</i>	297.74	≈	284.66	≈	289.54	≈	284.1	≈	289.67	≈	294.37	≈	285.33	≈	291.69
<i>Fudenberg et al. (2012)</i>	432.38	≈	421.46	>	377.96	≈	370.01	≈	381.2	≈	381.32	≈	372.32	≈	377.33
<i>Kagel and Schley (2013)</i>	2439.06	≈	2473.59	≈	2450.24	≈	2421.34	≈	2384.98	≈	2354.05	≈	2398.74	≪	2551.75
<i>Sherstyuk et al. (2013)</i>	1274.09	≈	1243.95	≈	1234.52	≈	1200.28	≈	1184.82	≈	1177.26	≈	1186.92	≪	1286.15
Pooled	24919.14	≫	24204.84	≈	24201.19	≈	24084.8	≈	24202.88	<	24473.15	>	24202.07	≪	24702.58
Second halves per session															
<i>Aoyagi and Fréchette (2009)</i>	482.66	≈	460.38	≫	421.21	≈	422.29	≈	423.58	>	404.94	≈	408.59	≈	409.05
<i>Blonski et al. (2011)</i>	1369.08	≈	1350.39	≈	1373.9	≈	1393.26	<	1452.56	≪	1561.35	≫	1458.47	≫	1382.85
<i>Bruttel and Kamecke (2012)</i>	493.79	≈	487.8	≈	480.47	≈	478.23	≈	470.25	≈	443.83	≈	471.73	<	528.54
<i>Dal Bó (2005)</i>	712.63	≈	688.66	≈	677.22	≈	679	<	698.19	≈	707.19	≈	687.86	≈	696.4
<i>Dal Bó and Fréchette (2011)</i>	4820.25	≈	4966.19	≫	4565.87	≈	4545.09	≈	4428.79	≈	4431.94	≈	4493.1	≪	5045.23
<i>Dal Bó and Fréchette (2019)</i>	7785.24	≈	7820.35	≫	7306.18	≈	7310.31	>	7171.6	≈	7089.53	≈	7152.05	≪	7683.74
<i>Dreber et al. (2008)</i>	558.39	≈	545.25	≈	544.66	≈	541.83	≈	539.8	≈	520.49	≈	519.56	<	563.51
<i>Duffy and Ochs (2009)</i>	1883.18	>	1764.77	>	1656.55	≈	1602.92	>	1518.65	≈	1509.7	<	1598.03	≪	1715.88
<i>Fréchette and Yuksel (2017)</i>	429.95	≈	436.46	≈	422.52	≈	382.45	≈	377.48	≈	384.1	≈	382.11	<	409.93
<i>Fudenberg et al. (2012)</i>	505.1	≈	493.46	≫	433.74	≈	416.17	≈	406.54	≈	410.66	≈	421.81	≈	448.37
<i>Kagel and Schley (2013)</i>	1621.42	≈	1713.66	≫	1572.95	≈	1541.38	>	1488.44	≈	1477.83	≈	1527.49	≪	1748.01
<i>Sherstyuk et al. (2013)</i>	890.27	≈	901.89	>	834.73	≈	823.06	≈	801.05	≈	801.52	≈	815.26	≪	935.01
Pooled	22069.97	≈	21738.7	≫	20545.33	≈	20464.29	>	20251.1	≈	20436.11	≈	20410.22	≪	21821.84

Note: This table verifies a number of possible mixtures involving semi-grim types as a robustness check for the sufficiency of focusing on the mixtures examined above. E.g. “3× SG” refers to a model containing three different versions of memory-1 semi-grim with allowing for heterogeneity of randomization parameters across subjects.

data sets and thus behaviorally adequate.

These results provide strong evidence for heterogeneity and the behavior-strategy conjecture, and as indicated, let us next implement the top-down and bottom-up approaches towards model selection. By the top-down approach, we start with the most general model ($3 \times P5 + AD$) and successively reduce its complexity until such reductions dampen its adequacy significantly. The simplest model that we reach this way without a significantly negative impact on adequacy is $1.5 \times SG + AD$ —with fixed semi-grim strategies in first halves of sessions and with flexible ones in second halves of sessions. This confirms the above result that cautious and strong semi-grim account for the vast majority of cooperating subjects by individual classification, and importantly, if the populations would be better captured by replacing one of the semi-grim types by any other strategy (say STFT), post-hoc selected at the treatment level, which would contradict the individual classification, then the restriction to semi-grim would have been rejected significantly. In turn, by the bottom-up approach, we start with the simplest model (Fixed SG, $1.5 \times SG + AD$) and successively increase its complexity as long as these increments significantly improve model adequacy. Starting with this, adequacy improves significantly only in second halves of sessions, then by allowing for flexible semi-grim strategies, but beyond that, further increments again are not significant in a manner surviving the Bonferroni correction (indicated by \gg or \ll in Table 6).

That is, both the top-down and the bottom-up approach converge to the same conclusion that we need to distinguish two types of cooperating subjects, whose behaviors differ only in round 1 of each supergame. On average, the less cooperative type cooperates with probabilities in $[0.2, 0.5]$ in round 1, similar to the cooperation probabilities after mixed histories cd/dc , and the more cooperative type cooperates with probabilities above 0.9 in most treatments, similar to cooperation probabilities after cc (see also Figure 2).

Result 2 (Question 2). *The analysis identifies two types of cooperating subjects playing the same semi-grim continuation strategy but different cooperation probabilities in round 1 (**cautious cooperators** and **strong cooperators**) and a subject type playing a strategy close to always defect (**defectors**). A model with this subject composition, and any other model allowing for two types of cooperating subjects playing behavior strategies, fits significantly better than all 10^{46} models assuming pure or generalized pure strategies.*

Hence, allowing for non-trivial randomization particularly in the states cd and dc is crucial for understanding subject behavior. As the fixed and flexible model fit equally well in the first halves of sessions, the actual strategies seem to be largely independent of treatment parameters for inexperienced subjects. This is not the case for experienced subjects, however.

We investigate this in detail next.

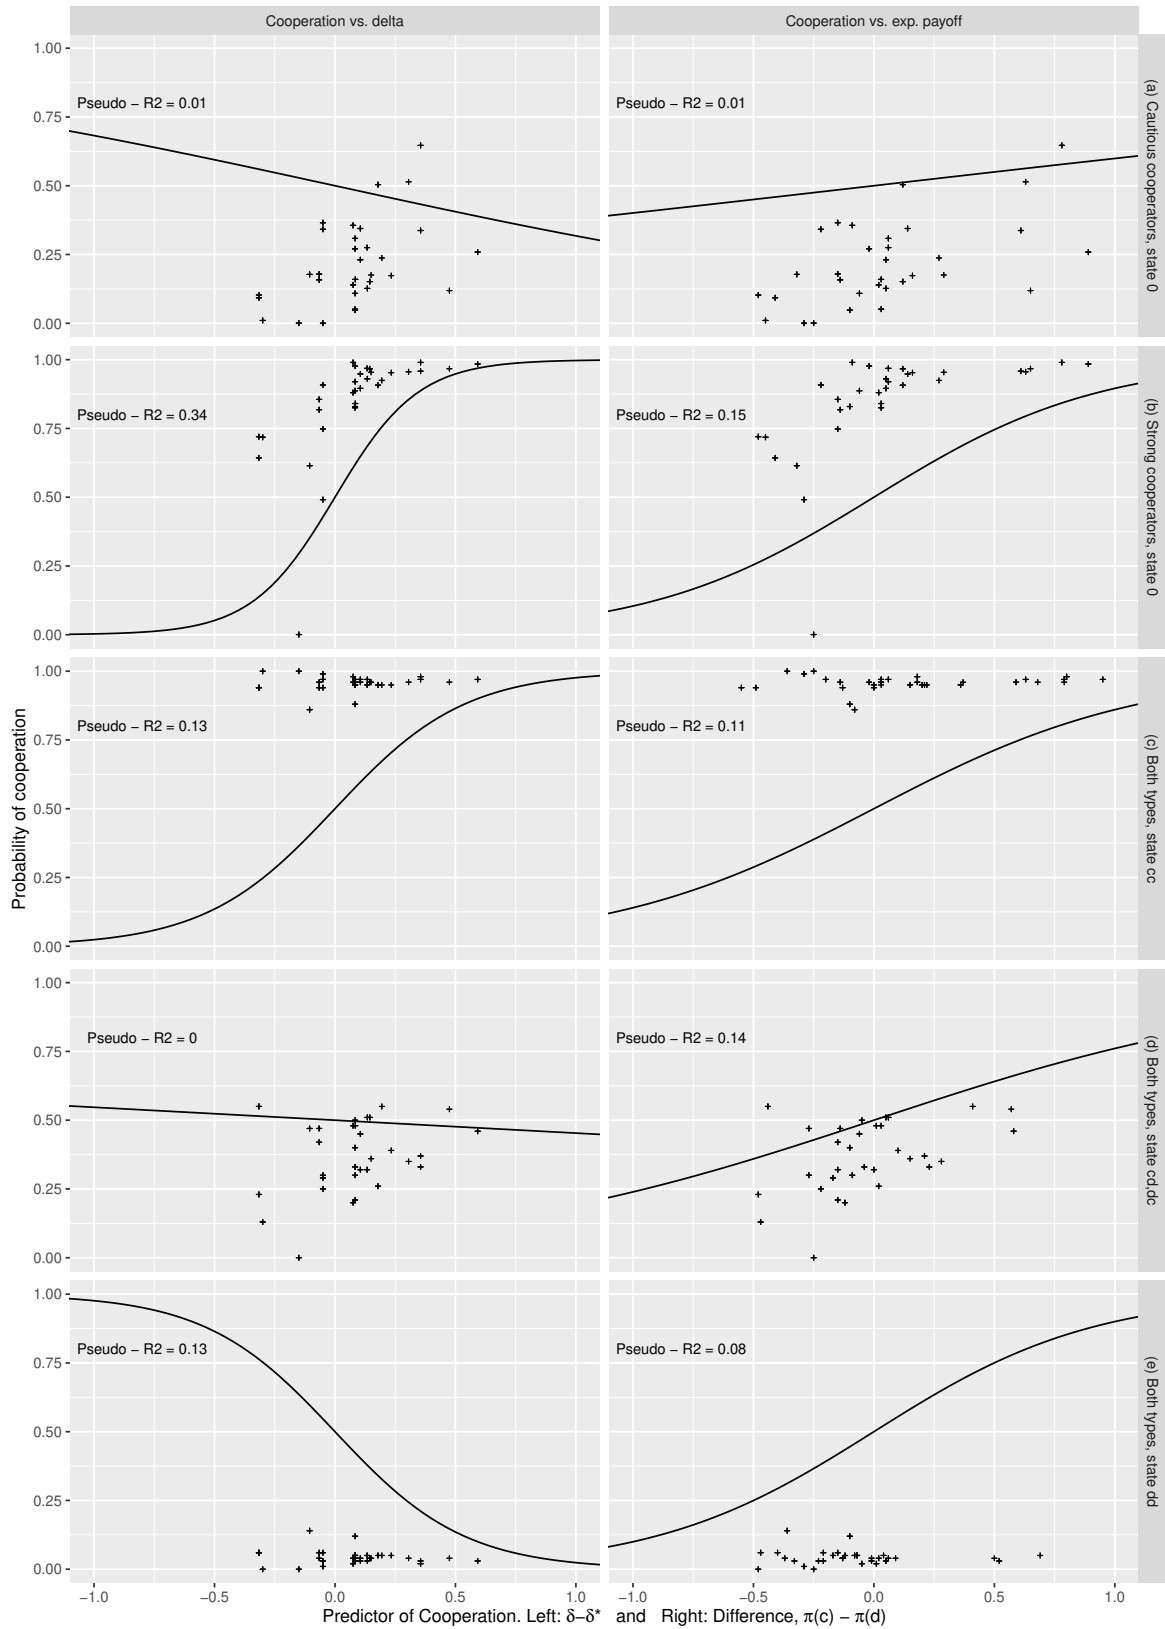
6 How do strategies relate to supergame parameters?

Having estimated the number of subject types, and along the way their strategies, we can revisit Question 3 and ask to what extent the subjects' strategies are functions of treatment parameters and to what extent they may be predictable. In light of the above results, we distinguish defecting and cooperating subjects. The defecting subjects play slightly perturbed strategies close to AD, which are essentially invariant to treatment parameters and rationalizable to the extent that AD is rationalizable (AD is a best response to itself in all supergames considered here). For this reason, we shall focus on the strategies played by cooperating subjects. By Result 2, there are two types of cooperating subjects, both identified as playing semi-grim supergame strategies with significant differences found in the probability of cooperation in round 1. The strategies are significantly treatment-dependent when subjects are experienced, i.e. in second halves of sessions, on which we shall focus in the following.

Overview Recall that Figure 1 plots the average cooperation rates across states against the expected payoffs from cooperation, which suggested that subjects act highly rationally in round 1 but ignore expected payoffs afterwards. We suspected confounds due to looking at raw cooperation rates, most notably possible selection effects, and our maximum-likelihood estimates of the strategies of (cooperating) subject types allow us to resolve these concerns. Figure 2 now plots the cooperation probabilities according to the estimated strategies of cooperating subjects against two predictors of cooperation (expected payoffs and $\delta - \delta^*$). In the left column of plots, we see how the cooperation probabilities across states relate to the difference of discount factor δ and BOS threshold²³ $\delta^* = (g + \ell)/(1 + g + \ell)$. In the right column of plots, we see how the probability of cooperation relates to the monetary incentive to cooperate, $\pi_\omega(c) - \pi_\omega(d)$ as defined above, Eqs. 1–3, for each state ω . For these plots, we assume that subjects hold “false consensus” beliefs that their opponent plays the same strategy that they play. That is, strong cooperators believe they face strong cooperators and weak cooperators believe they face weak cooperators. In comparison to Figure 1, the results do not change substantially: Behavior is still close to being independent of the predictors of cooperation in most states (bottom three panels), arguably with the exception of strong cooperators in round 1 (panel b).


²³The threshold is axiomatically derived in Blonski et al. (2011) and characterized by risk dominance in Blonski and Spagnolo (2015).

Figure 2: Relation of δ (left) and monetary incentives (right) to cooperation rates (second halves of sessions)



Note: Cooperation probabilities of estimated strategies by states and by two predictors of cooperation: Difference between δ and the BOS threshold δ^* (left panel) and the monetary incentive $\hat{\pi}(c) - \hat{\pi}(d)$ (right panel), with $\hat{\pi}(c)$ the expected payoff of cooperating in that state and $\hat{\pi}(d)$ the expected payoff of defecting in that state.

Figure 3: Relation of $\delta - \delta^*$ to shares of cooperators (second halves of sessions)



Note: This figure shows how the ratios of the three strategies – defectors, cautious cooperators, and strong cooperators change with the distance of δ to the BOS cooperation threshold δ^* across treatments. The solid line represents the best fitting logistic curve estimated with intercept. Panel (a) displays the total share of both cooperators, panel (b) the relative share of cautious cooperators among cooperators, panel (c) the share of cautious cooperators overall, panel (d) the share of strong cooperators overall.

Recall that we know from Dal Bó (2005) and subsequent work that average cooperation rates change as payoff parameters change, and above we have seen that most of these changes can be reduced to changes in round 1 (see also Dal Bó and Fréchette, 2018). Yet, as just seen, the types of strategies are largely independent of the payoff parameters. To illuminate this further, we next test the complementary statistic and examine how the shares of the three subject types change as parameters change. Figure 3 plots the shares of cooperators as a function of the discount factor δ in relation to the BOS-threshold δ^* . We see two relatively strong effects: As δ approaches δ^* , the overall share of cooperators increases, i.e. defectors become cooperators, and at the same time, the relative share of cautious cooperators declines, i.e. cautious cooperators turn into strong cooperators.

Result 3 (Question 3). *The shares of subjects playing either of the three strategies change highly predictably. As δ increases defectors are replaced by cooperators and as it passes the BOS-threshold δ^* the strong cooperators start to outweigh the cautious cooperators ($\bar{R}^2 \geq$*

0.2 in all cases). The strategies themselves are largely invariant to treatment parameters and monetary incentives. The only exception satisfying $\tilde{R}^2 \geq 0.2$ are strong cooperators in round 1, whose strategies correlate with treatment parameters ($\delta - \delta^*$) but not with monetary incentives, and only in round 1.

That is, the behavioral changes regarding average and round 1 cooperation observed in the literature (e.g. Dal Bó and Fréchette, 2018, Result 4) are mainly transitions from defection to cautious cooperation and from cautious cooperation to strong cooperation. These transitions are neatly predictable, being logistic functions of $\delta - \delta^*$, which is a substantial result in relation to previous work that found no reliable association between the individual strategies used and the payoff parameters (Dal Bó and Fréchette, 2018). This result directly follows from the unrestricted estimation of strategies, which thus not only fits better but also renders type shares predictable. In turn, the actual strategies associated with these seemingly archetypical behavioral types are largely invariant of payoff parameters, which also is a novel result raising questions about their potential rationalizability for future work.

7 Concluding discussion

Revisiting our initial research questions, we summarize our main results as follows.

Regarding the first two questions, re-analyzing 12 experiments, we find robust evidence for the use of behavior strategies and subject heterogeneity. The most parsimonious population model consistently capturing behavior in 12 experiments contains three different types of subjects per treatment: defectors, playing a strategy close to AD, and cautious and strong cooperators who play behavior strategies with a semi-grim structure that differ only in their first-round cooperation probability. This confirms the subject-level classification, which assigns about 80% of the subjects to either AD or stylized semi-grim strategies.

Regarding our third question, we find that the estimated strategies are largely independent of treatment parameters but the shares of subjects picking either of the three strategies depend strikingly on the continuation probability δ in relation to the BOS-threshold δ^* (Blonski et al., 2011). To be more precise, strategies are statistically treatment-invariant in first halves of sessions (when subjects are inexperienced), and strategy shares are predictable (see also Figure 4 in the appendix). In second halves of sessions, the strategy choice is still predictable, but the treatment invariance is rejected highly significantly—suggesting that subjects have started to adapt their behavior to the environment, but not in a way that is immediately clear or even predictable based on predictors of cooperation.

Moreover, following rounds where at least one player cooperated, the cooperators cooperate systematically even in supergames where Grim is not a subgame-perfect equilibrium. In such supergames, cooperation is not rationalizable if we assume that subjects seek to maximize pecuniary payoffs, i.e. it is never a best response, regardless of the subjective beliefs about the opponent's strategy.

By using the advantages of a meta-study in aggregating plenty of observations, we identify how future research can analyze behavior in the repeated PD: We need to distinguish three subject types and estimate the behavior strategies used by the cooperating subjects. Based on this, we can aim to identify the strategies and predict behavior in repeated games more generally, which at present seems to be an ambitious long-term goal.

The above observation that subjects cooperate even in supergames where it is not rationalizable based on pecuniary payoffs may suggest that subjects have interdependent preferences, but the recurring semi-grim strategies are also instances of belief-free equilibria (Ely et al., 2005) for appropriate discount factors. Furthermore, the predictive BOS-threshold δ^* is the threshold for the existence of belief-free equilibria satisfying $\sigma_{cd} = \sigma_{dc}$. However, the cooperation probabilities exhibit semi-grim patterns even if belief-free equilibria sustaining cooperation do not exist, which is not an obvious implication of the theory as it is currently developed. Having said this, cooperating subjects do not generally play subgame-perfect equilibria of the games they face, hence also not belief-free equilibria of the games they face. This may be indicative of analogical reasoning in the sense of Samuelson (2001), who argued that subjects initially play the version of the game they know from prior experience and over time adapt to the new circumstances, which may explain the divergence from treatment-invariant behavior over time.

This might also explain how semi-grim strategies are being played although belief-free equilibria are not particularly robust, e.g. to utility perturbations (Doraszelski and Escobar, 2010), strategy perturbations (Heller, 2017) or evolutionary mutation (Kandori, 2011), and appear “unrealistically complex” (Compte and Postlewaite, 2015): subjects are not particularly adaptive to the game that they face in the first place. If the recurrence of strategies similar to semi-grim is not a coincidence, belief-free equilibrium may thus help us to predict behavior in repeated games more generally in the long run.

References

- Ansari, A., Montoya, R., and Netzer, O. (2012). Dynamic learning in behavioral games: A hidden markov mixture of experts approach. *Quantitative Marketing and Economics*, 10(4):475–503.
- Aoyagi, M. and Frechette, G. (2009). Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic theory*, 144(3):1135–1165.
- Aumann, R. (1964). Mixed and behavior strategies in infinite extensive games. In Dresher, M., Shapley, L. S., and Tucker, A. W., editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, page 627–650. Princeton University Press.
- Axelrod, R. (1980a). Effective choice in the prisoner’s dilemma. *Journal of conflict resolution*, 24(1):3–25.
- Axelrod, R. (1980b). More effective choice in the prisoner’s dilemma. *Journal of conflict resolution*, 24(3):379–403.
- Bajari, P. and Hortacsu, A. (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy*, 113(4):703–741.
- Biernacki, C., Celeux, G., and Govaert, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE transactions on pattern analysis and machine intelligence*, 22(7):719–725.
- Bilmes, J. A. et al. (1998). A gentle tutorial of the em algorithm and its application to parameter estimation for gaussian mixture and hidden markov models. *International Computer Science Institute*, 4(510):126.
- Blonski, M., Ockenfels, P., and Spagnolo, G. (2011). Equilibrium selection in the repeated prisoner’s dilemma: Axiomatic approach and experimental evidence. *American Economic Journal: Microeconomics*, 3(3):164–192.
- Blonski, M. and Spagnolo, G. (2015). Prisoners’ other dilemma. *International Journal of Game Theory*, 44:61–81.
- Breitmoser, Y. (2015). Cooperation, but no reciprocity: Individual strategies in the repeated prisoner’s dilemma. *American Economic Review*, 105(9):2882–2910.
- Breitmoser, Y. (2019). Knowing me, imagining you: Projection and overbidding in auctions. *Games and Economic Behavior*, 113:423–447.

- Breitmoser, Y., Tan, J. H., and Zizzo, D. J. (2014). On the beliefs off the path: Equilibrium refinement due to quantal response and level-k. *Games and Economic Behavior*, 86:102–125.
- Brunner, C., Camerer, C. F., and Goeree, J. K. (2011). Stationary concepts for experimental 2 x 2 games: Comment. *American Economic Review*, 101(2):1029–40.
- Bruttel, L. and Kamecke, U. (2012). Infinity in the lab. how do people play repeated games? *Theory and Decision*, 72(2):205–219.
- Camera, G., Casari, M., and Bigoni, M. (2012). Cooperative strategies in anonymous economies: An experiment. *Games and Economic Behavior*, 75(2):570–586.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869.
- Compte, O. and Postlewaite, A. (2015). Plausible cooperation. *Games and Economic Behavior*, 91:45–59.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.
- Dal Bó, P. (2005). Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economic Review*, 95(5):1591–1604.
- Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429.
- Dal Bó, P. and Fréchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1):60–114.
- Dal Bó, P. and Fréchette, G. R. (2019). Strategy choice in the infinitely repeated prisoner’s dilemma. *American Economic Review*, 109(11):3929–52.
- Doraszelski, U. and Escobar, J. F. (2010). A theory of regular markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoretical Economics*, 5(3):369–402.
- Dreber, A., Rand, D. G., Fudenberg, D., and Nowak, M. A. (2008). Winners don’t punish. *Nature*, 452(7185):348–351.

- Duersch, P., Oechssler, J., and Schipper, B. C. (2014). When is tit-for-tat unbeatable? *International Journal of Game Theory*, 43(1):25–36.
- Duffy, J. and Ochs, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, 66(2):785–812.
- Dvorak, F. and Fehrler, S. (2018). Negotiating cooperation under uncertainty: Communication in noisy, indefinitely repeated interactions. *IZA Discussion Paper No. 11897*.
- El-Gamal, M. A. and Grether, D. M. (1995). Are people bayesian? uncovering behavioral strategies. *Journal of the American statistical Association*, 90(432):1137–1145.
- Ely, J. C., Hörner, J., and Olszewski, W. (2005). Belief-free equilibria in repeated games. *Econometrica*, 73(2):377–415.
- Embrey, M., Fréchette, G. R., and Yuksel, S. (2018). Cooperation in the finitely repeated prisoner’s dilemma. *The Quarterly Journal of Economics*, page qjx033.
- Engle-Warnick, J. and Slonim, R. L. (2006). Inferring repeated-game strategies from actions: evidence from trust game experiments. *Economic theory*, 28(3):603–632.
- Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868.
- Fréchette, G. R. and Yuksel, S. (2017). Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics*, 20(2):279–308.
- Fudenberg, D. and Karreskog, G. (2023). Predicting cooperation with learning models. *American Economic Journal: Microeconomics*.
- Fudenberg, D. and Maskin, E. S. (1986). Folk theorem for repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554.
- Fudenberg, D., Rand, D. G., and Dreber, A. (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102(2):720–749.
- Goeree, J. K., Holt, C. A., and Pfafrey, T. R. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113.

- Harless, D. W. and Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, 62(6):1251–1289.
- Heller, Y. (2017). Instability of belief-free equilibria. *Journal of Economic Theory*, 168:261–286.
- Heller, Y. and Tubul, I. (2023). Strategies in the repeated prisoner’s dilemma: A cluster analysis. Technical report.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.
- Houser, D., Keane, M., and McCabe, K. (2004). Behavior in a dynamic decision problem: An analysis of experimental evidence using a bayesian type classification algorithm. *Econometrica*, 72(3):781–822.
- Houser, D. and Winter, J. (2004). How do behavioral assumptions affect structural inference? evidence from a laboratory experiment. *Journal of Business & Economic Statistics*, 22(1):64–79.
- Imhof, L. A., Fudenberg, D., and Nowak, M. A. (2007). Tit-for-tat or win-stay, lose-shift? *Journal of theoretical biology*, 247(3):574–580.
- Kagel, J. H. and Schley, D. R. (2013). How economic rewards affect cooperation reconsidered. *Economics Letters*, 121(1):124–127.
- Kandori, M. (2011). Weakly belief-free equilibria in repeated games with private monitoring. *Econometrica*, 79(3):877–892.
- McLachlan, G. J. and Peel, D. (2004). *Finite Mixture Models*. John Wiley & Sons.
- Nowak, M. and Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner’s dilemma game. *Nature*, 364(6432):56.
- Rand, D. G., Fudenberg, D., and Dreber, A. (2015). It’s the thought that counts: The role of intentions in noisy repeated games. *Journal of Economic Behavior & Organization*, 116:481–499.
- Romero, J. and Rosokha, Y. (2023). Mixed strategies in the indefinitely repeated prisoner’s dilemma. available on personal webpage of Julian Romero 22/02/2023.

- Rubinstein, A. (1986). Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory*, 39(1):83–96.
- Samuelson, L. (2001). Analogies, adaptation, and anomalies. *Journal of Economic Theory*, 97(2):320–366.
- Schennach, S. M. and Wilhelm, D. (2017). A simple parametric model selection test. *Journal of the American Statistical Association*, pages 1–12.
- Schmidt, K. (1993). Reputation and equilibrium characterization in repeated games with conflicting interests. *Econometrica*, 61(2):325–51.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International journal of game theory*, 4(1):25–55.
- Shachat, J., Swarthout, J. T., and Wei, L. (2015). A hidden markov model for the detection of pure and mixed strategy play in games. *Econometric Theory*, 31(4):729–752.
- Sherstyuk, K., Tarui, N., and Saijo, T. (2013). Payment schemes in infinite-horizon experimental games. *Experimental Economics*, 16(1):125–153.
- Stahl, D. O. and Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327.
- Weizsäcker, G. (2003). Ignoring the rationality of others: evidence from experimental normal-form games. *Games and Economic Behavior*, 44(1):145–171.
- Wright, S. P. (1992). Adjusted p-values for simultaneous inference. *Biometrics*, 48:1005–1013.

Appendix

A Econometric approach and validity

A.1 Possible memory-1 strategies

Given that previous literature has focused on memory-1 strategies and has claimed that memory-1 is enough to describe players behavior in games with perfect monitoring (Dal Bó and Fréchette, 2019) we use memory-1 as a starting point. In terms of pure strategies there are 16 possible memory-1 strategies where subjects cooperate in round 1, and another 16 where they do not cooperate in round 1. Reducing the number by requiring that if cooperation in any state, then also in σ_{cc} (why else cooperate?), yields the following memory-1 strategies:

With initial defection: $(\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0, 0, 0, 0, 0), (0, 1, 0, 0, 0), (0, 1, 1, 0, 0), (0, 1, 0, 1, 0), (0, 1, 0, 0, 1), (0, 1, 1, 1, 0), (0, 1, 1, 0, 1), (0, 1, 1, 1, 1)$

With initial cooperation: $(\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (1, 1, 0, 0, 0), (1, 1, 1, 0, 0), (1, 1, 0, 1, 0), (1, 1, 0, 0, 1), (1, 1, 1, 1, 0), (1, 1, 1, 0, 1), (1, 1, 1, 1, 1)$

Between them, the following strategies are behaviorally indistinguishable: $(0, 0, 0, 0, 0) \equiv (0, 1, 0, 0, 0) \equiv (0, 1, 1, 0, 0)$

Also indistinguishable: $(1, 1, 1, 0, 0) \equiv (1, 1, 1, 1, 0) \equiv (1, 1, 1, 0, 1) \equiv (1, 1, 1, 1, 1)$

Hence, there are 10 remaining memory-1 strategies: $(\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0, 0, 0, 0, 0), (0, 1, 0, 1, 0), (0, 1, 0, 0, 1), (0, 1, 1, 1, 0), (0, 1, 1, 0, 1), (0, 1, 1, 1, 1), (1, 1, 0, 0, 0), (1, 1, 1, 0, 0), (1, 1, 0, 1, 0), (1, 1, 0, 0, 1)$

Table 4 in the main body shows the classification of subjects behavior into these strategies with maximum likelihood. For the maximum-likelihood classification we have allowed for a noise level of 1% (robustness with 3%) for the pure strategies to avoid potential likelihoods of minus infinity. Based on this we calculated for each individual the log-likelihoods for playing a given strategy with

$$ll(\sigma|c, d) = \sum_{\omega} (c_{\omega} \cdot \log \sigma_{\omega} + d_{\omega} \cdot \log(1 - \sigma_{\omega}))$$

where ω are the memory-1 states, $\sigma \in \{\sigma^{AD}, \sigma^{TFT}, \dots, \sigma^{str}, \sigma^{cau}\}$ the strategies, and c_{ω} and d_{ω} the number of cooperations and defections in a given state, respectively. We classify subjects by the maximum likelihood. We allow for an exhaustive set of pure memory-1 strategies including besides well known strategies also for completely unknown candidates.

The Tables 7 and 8 show the actual strategies played by the subjects classified in the respective strategy, with and without allowing for behavior strategies. Without allowing for behavior strategies (Table 7) we see that the actual play is far from representing the strategy that they are assigned to.

Table 7: Actual strategies played by players' best fitting strategy allowing only for pure strategies (maximum likelihood classification)

	Share	"N"	Strategy definition					Average strategy played				
			\emptyset	(cc)	(cd)	(dc)	(dd)	\emptyset	(cc)	(cd)	(dc)	(dd)
"AD"	0.32	551.5	0	0	0	0	0	0.07	0.54	0.19	0.08	0.02
"S-TFT"	0.13	225	0	1	0	1	0	0.24	0.96	0.2	0.64	0.07
"S-WSLS"	0	3	0	1	0	0	1	0.24	0.47	0.19	0.32	0.58
"S-AC1"	0.04	75	0	1	1	1	0	0.22	0.96	0.66	0.67	0.09
"S-AC2"	0	4	0	1	1	0	1	0.57	0.76	0.63	0.38	0.68
"S-AC3"	0	4	0	1	1	1	1	0.26	0.67	0.53	0.72	0.68
"Grim"	0.22	373.5	1	1	0	0	0	0.82	0.96	0.21	0.2	0.07
"AC"	0.07	126.5	1	1	1	0	0	0.8	0.97	0.72	0.39	0.1
"TFT"	0.21	360	1	1	0	1	0	0.85	0.98	0.23	0.78	0.06
"WSLS"	0.01	11.5	1	1	0	0	1	0.83	0.87	0.4	0.46	0.78

Note: Strategy shares by maximum likelihood classification allowing for 1% deviation from pure strategies. The set includes all possible memory-1 strategies (see above). If two strategies fit equally well, known strategies are prioritized over unknowns (S-ACx,S-WSLS), and otherwise the subject is assigned half to one and half to the other strategies.

Table 8: Actual strategies played by players' best fitting strategy allowing for semi-grim and pure strategies (maximum likelihood classification)

	Share	"N"	Strategy definition					Average strategy played				
			\emptyset	(cc)	(cd)	(dc)	(dd)	\emptyset	(cc)	(cd)	(dc)	(dd)
"AD"	0.21	366	0	0	0	0	0	0.02	0.24	0.13	0.02	0.01
"S-TFT"	0.02	34	0	1	0	1	0	0.05	1	0.15	0.84	0.02
"S-WSLS"	0	0	0	1	0	0	1	–	–	–	–	–
"S-AC1"	0.01	14	0	1	1	1	0	0.05	1	0.7	0.9	0.04
"S-AC2"	0	0	0	1	1	0	1	–	–	–	–	–
"S-AC3"	0	0	0	1	1	1	1	–	–	–	–	–
"Grim"	0.05	91	1	1	0	0	0	0.98	0.99	0.07	0.13	0.02
"AC"	0.02	32	1	1	1	0	0	0.98	1	0.88	0.27	0.06
"TFT"	0.08	136	1	1	0	1	0	0.97	1	0.09	0.91	0.03
"WSLS"	0	1	1	1	0	0	1	1	1	0.5	–	1
"cau SG"	0.35	606	0.3	0.9	0.3	0.3	0.1	0.29	0.91	0.33	0.41	0.08
"str SG"	0.26	454	0.9	0.9	0.3	0.3	0.1	0.84	0.96	0.35	0.45	0.09

Note: Strategy shares by maximum likelihood classification allowing for 1% deviation from pure strategies. The set includes all possible memory-1 strategies (see above) and two behavior strategies cautious and strong cooperators derived from Breitmoser (2015). If two strategies fit equally well, known strategies are prioritized over unknowns (S-ACx,S-WSLS), and otherwise the subject is assigned half to one and half to the other strategies.

Table 9 shows the robustness classification via euclidean distances, i.e. classifying the subject into the strategy with the shortest distance. There is little difference in classifications.

Table 9: Classifying subjects into memory-1 strategies using euclidean distances

	Strategies					pure	
	\emptyset	(cc)	(cd)	(dc)	(dd)	" best fit"	"perfect fit"
"AD"	0	0	0	0	0	0.27	0.097
"S-TFT"	0	1	0	1	0	0.134	0.006
"S-WLS"	0	1	0	0	1	0.003	0
"S-AC1"	0	1	1	1	0	0.059	0.005
"S-AC2"	0	1	1	0	1	0.003	0
"S-AC3"	0	1	1	1	1	0.005	0
"Grim"	1	1	0	0	0	0.228	0.015
"AC"	1	1	1	0	0	0.073	0.008
"TFT"	1	1	0	1	0	0.21	0.023
"WLS"	1	1	0	0	1	0.014	0.002
"sum"						1	0.155

A.2 Validity

We test the validity of our econometric approach to distinguish pure, mixed, and behavior strategies by running the algorithm on different sets of simulated data: For each of the three strategy types, we simulate 17 populations following estimates of DF18, and each population 100 times, in order to verify if we identify the underlying strategy type based on model-fit evaluations using ICL-BIC. Before we proceed, let us emphasize that our estimation proceeds by maximum likelihood, implying that we can expect to be correct (on average) if the data set is sufficiently large. We want to find out, however, if typical experimental data sets are sufficiently large and what the probability of misidentification is. In addition, let us note that the analysis below does not engage in the kind of data mining implemented in our main analysis. This data mining yields a bias towards pure and mixed strategies, while the simulations now report the results without this bias, verifying the reliability of unbiased identification. On top of the simulation results reported next, we should therefore expect a bias in favor of pure or mixed strategies in our actual analysis.

In our analysis, we simulate populations representing type proportions implementing the actual estimates of DF18, as presented in their Table 10, for all 17 treatments analyzed in DF18. This way, we can ask whether our approach identifies pure strategies if subjects actually play the pure strategies in the proportions identified by DF18, and the same for corresponding mixtures of mixed strategies or behavior strategies. The tremble parameter is $\gamma = 0.05$. For each simulation run we determine the average ICL-BICs of the three basic

econometric models, no-switching (pure strategies), random-switching (mixed strategies), and semi-grim or $AD+1.5\times SG$ (behavior strategies). Aggregating across the 100 simulated data sets per population, we count how often which strategy type obtained the lowest ICL-BICs and got identified this way.

Tables 3 and 11 reports the results for populations with 50 subjects and 20 observations (on average) per subject past round 1, and as a robustness check, and Tables 10 and 12 report the results of simulation runs in larger data sets with 30 observations per subject past round 1. In comparison, the data sets we analyze contain on average 90 observations per subject, which we split into two halves containing about 45 observations per subject each, putting us above either range on average. In all cases, the true models are identified in the vast majority of simulation runs, with relative frequencies above 90% on average across treatments. At the treatment level, the only exceptions are treatment 6 of DF11, where subjects playing pure strategies are likely to be misidentified as playing mixed strategies, and treatment 1 of DF11, where subjects playing behavior strategies are somewhat likely to be misidentified as playing pure strategies. Critically, there is no bias in favor of behavior strategies in any condition.

Result 4. *Our approach identifies pure, mixed and behavior strategies with more than 90% reliability even in populations comprising just 50 subjects and 20 or 30 decisions per subject (past round 1). Singular treatments exhibit biases towards pure and mixed strategies, but no treatment exhibits a bias towards semi-grim strategies.*

In conclusion, let us recall that our data set overall is much larger than such small data sets, which obviously benefits the reliability of identification by the limiting properties of maximum likelihood.

Table 10: Identification of strategy types in populations with 30 observations past round 1 and type shares of the 17 different treatments estimated by DF18

	Type proportions following DF18, Tab. 10 (σ_{AD} , σ_{Grim} , σ_{TFT} , σ_{AC} , σ_{STFT})	In 100 simulation runs, identified as		
		Pure strategies	Mixed strategies	AD + Semi-Grim
<i>Simulated agents play: Pure strategies ("No switching")</i>				
Dreber et al. (2008), T1	(0.64, 0.07, 0.15, 0, 0.14)	100	0	0
Dreber et al. (2008), T2	(0.3, 0.21, 0.4, 0, 0.09)	100	0	0
Dal Bo and Frechette (2011), T1	(0.91, 0, 0.07, 0, 0.02)	99	0	1
Dal Bo and Frechette (2011), T2	(0.76, 0, 0.06, 0, 0.08)	96	4	0
Dal Bo and Frechette (2011), T3	(0.49, 0, 0.24, 0.01, 0.04)	100	0	0
Dal Bo and Frechette (2011), T4	(0.66, 0, 0.23, 0, 0)	100	0	0
Dal Bo and Frechette (2011), T5	(0.11, 0.04, 0.21, 0, 0.08)	100	0	0
Dal Bo and Frechette (2011), T6	(0, 0.02, 0.55, 0.02, 0)	38	62	0
Fudenberg et al. (2012)	(0.06, 0.12, 0.15, 0.24, 0)	94	2	4
Rand et al. (2015)	(0.18, 0.43, 0.27, 0, 0.05)	100	0	0
Frechette and Yuksel (2014)	(0.14, 0.32, 0.39, 0, 0.02)	100	0	0
Dal Bo and Frechette (2015), T1	(0.53, 0.06, 0.05, 0, 0.14)	100	0	0
Dal Bo and Frechette (2015), T2	(0.25, 0.36, 0.19, 0.03, 0.03)	100	0	0
Dal Bo and Frechette (2015), T3	(0.47, 0.1, 0.1, 0.02, 0.12)	100	0	0
Dal Bo and Frechette (2015), T4	(0.12, 0.35, 0.3, 0.08, 0)	100	0	0
Dal Bo and Frechette (2015), T5	(0.14, 0.17, 0.39, 0.02, 0.07)	100	0	0
Dal Bo and Frechette (2015), T6	(0.22, 0.06, 0.25, 0, 0)	100	0	0
Probability of identification overall		0.957	0.04	0.003
<i>Simulated agents play: Mixed strategies ("Random switching")</i>				
Dreber et al. (2008), T1	(0.64, 0.07, 0.15, 0, 0.14)	0	100	0
Dreber et al. (2008), T2	(0.3, 0.21, 0.4, 0, 0.09)	0	100	0
Dal Bo and Frechette (2011), T1	(0.91, 0, 0.07, 0, 0.02)	0	100	0
Dal Bo and Frechette (2011), T2	(0.76, 0, 0.06, 0, 0.08)	0	100	0
Dal Bo and Frechette (2011), T3	(0.49, 0, 0.24, 0.01, 0.04)	0	100	0
Dal Bo and Frechette (2011), T4	(0.66, 0, 0.23, 0, 0)	0	100	0
Dal Bo and Frechette (2011), T5	(0.11, 0.04, 0.21, 0, 0.08)	0	100	0
Dal Bo and Frechette (2011), T6	(0, 0.02, 0.55, 0.02, 0)	3	97	0
Fudenberg et al. (2012)	(0.06, 0.12, 0.15, 0.24, 0)	0	100	0
Rand et al. (2015)	(0.18, 0.43, 0.27, 0, 0.05)	0	100	0
Frechette and Yuksel (2014)	(0.14, 0.32, 0.39, 0, 0.02)	0	100	0
Dal Bo and Frechette (2015), T1	(0.53, 0.06, 0.05, 0, 0.14)	0	100	0
Dal Bo and Frechette (2015), T2	(0.25, 0.36, 0.19, 0.03, 0.03)	0	100	0
Dal Bo and Frechette (2015), T3	(0.47, 0.1, 0.1, 0.02, 0.12)	0	100	0
Dal Bo and Frechette (2015), T4	(0.12, 0.35, 0.3, 0.08, 0)	0	100	0
Dal Bo and Frechette (2015), T5	(0.14, 0.17, 0.39, 0.02, 0.07)	0	100	0
Dal Bo and Frechette (2015), T6	(0.22, 0.06, 0.25, 0, 0)	0	100	0
Probability of identification overall		0.002	0.998	0

Note: Analysis based on simulated data sets comprising 50 subjects and 60 observations (past round 1) per subject, reporting the frequency of identification of the different strategy classes using our econometric methodology (where identification of either pure, mixed or behavior strategies is based on the ICL-BICs of the corresponding no-switching, random-switching, or semi-grim structures).

Table 11: Identification of strategy in populations with 20 observations past round 1 where agents play AD + semi-grim behavior strategies in proportions representing the estimates of DF18

	Type proportions following DF18, Tab. 10 (σ_{AD} , σ_{Grim} , σ_{TFT} , σ_{AC} , σ_{STFT})	In 100 simulation runs, identified as		
		Pure strategies	Mixed strategies	AD + Semi-Grim
<i>True model: semi-grim behavior strategies</i>				
Dreber et al. (2008), T1	(0.64, 0.18, 0.18)	0	0	100
Dreber et al. (2008), T2	(0.3, 0.35, 0.35)	0	0	100
Dal Bo and Frechette (2011), T1	(0.91, 0.045, 0.045)	19	4	77
Dal Bo and Frechette (2011), T2	(0.76, 0.12, 0.12)	0	0	100
Dal Bo and Frechette (2011), T3	(0.49, 0.255, 0.255)	0	0	100
Dal Bo and Frechette (2011), T4	(0.66, 0.17, 0.17)	0	0	100
Dal Bo and Frechette (2011), T5	(0.11, 0.445, 0.445)	0	0	100
Dal Bo and Frechette (2011), T6	(0, 0.5, 0.5)	0	0	100
Fudenberg et al. (2012)	(0.06, 0.47, 0.47)	0	0	100
Rand et al. (2015)	(0.18, 0.41, 0.41)	0	0	100
Frechette and Yuksel (2014)	(0.14, 0.43, 0.43)	0	0	100
Dal Bo and Frechette (2015), T1	(0.53, 0.235, 0.235)	0	0	100
Dal Bo and Frechette (2015), T2	(0.25, 0.375, 0.375)	0	0	100
Dal Bo and Frechette (2015), T3	(0.47, 0.265, 0.265)	0	0	100
Dal Bo and Frechette (2015), T4	(0.12, 0.44, 0.44)	0	0	100
Dal Bo and Frechette (2015), T5	(0.14, 0.43, 0.43)	0	0	100
Dal Bo and Frechette (2015), T6	(0.22, 0.39, 0.39)	0	0	100
Probability of identification overall		0.011	0.002	0.986

Note: Analysis based on simulated data sets comprising 50 subjects and 20 observations (past round 1) per subject, reporting the frequency of identification of the different strategy classes using our econometric methodology (where identification of either pure, mixed or behavior strategies is based on the ICL-BICs of the corresponding no-switching, random-switching, or semi-grim structures).

Table 12: Identification of strategy in populations with 30 observations past round 1 where agents play AD + semi-grim behavior strategies in proportions representing the estimates of DF18

	Type proportions following DF18, Tab. 10 (σ_{AD} , σ_{Grim} , σ_{TFT} , σ_{AC} , σ_{STFT})	In 100 simulation runs, identified as		
		Pure strategies	Mixed strategies	AD + Semi-Grim
<i>True model: semi-grim behavior strategies</i>				
Dreber et al. (2008), T1	(0.64, 0.18, 0.18)	0	0	100
Dreber et al. (2008), T2	(0.3, 0.35, 0.35)	0	0	100
Dal Bo and Frechette (2011), T1	(0.91, 0.045, 0.045)	6	0	94
Dal Bo and Frechette (2011), T2	(0.76, 0.12, 0.12)	0	0	100
Dal Bo and Frechette (2011), T3	(0.49, 0.255, 0.255)	0	0	100
Dal Bo and Frechette (2011), T4	(0.66, 0.17, 0.17)	0	0	100
Dal Bo and Frechette (2011), T5	(0.11, 0.445, 0.445)	0	0	100
Dal Bo and Frechette (2011), T6	(0, 0.5, 0.5)	0	0	100
Fudenberg et al. (2012)	(0.06, 0.47, 0.47)	0	0	100
Rand et al. (2015)	(0.18, 0.41, 0.41)	0	0	100
Frechette and Yuksel (2014)	(0.14, 0.43, 0.43)	0	0	100
Dal Bo and Frechette (2015), T1	(0.53, 0.235, 0.235)	0	0	100
Dal Bo and Frechette (2015), T2	(0.25, 0.375, 0.375)	0	0	100
Dal Bo and Frechette (2015), T3	(0.47, 0.265, 0.265)	0	0	100
Dal Bo and Frechette (2015), T4	(0.12, 0.44, 0.44)	0	0	100
Dal Bo and Frechette (2015), T5	(0.14, 0.43, 0.43)	0	0	100
Dal Bo and Frechette (2015), T6	(0.22, 0.39, 0.39)	0	0	100
Probability of identification overall		0.004	0	0.996

Note: Analysis based on simulated data sets comprising 50 subjects and 30 observations (past round 1) per subject, reporting the frequency of identification of the different strategy classes using our econometric methodology (where identification of either pure, mixed or behavior strategies is based on the ICL-BICs of the corresponding no-switching, random-switching, or semi-grim structures).

B Information on the experiments re-analyzed

This section provides some background information on the experiments re-analyzed in this paper. Table 13 summarizes and defines the strategies considered by previous studies. Table 14 reviews the numbers of subjects and observations, average parameters, and average cooperation rates for all experiments. A table with detailed overview by treatments is available upon request.

Table 13: Pure strategies considered in behavioral analyses

Strategy	Abbreviation	Description	$(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})^\dagger$	References
Pure Strategies Non-responsive or Memory-1				
Always Defect	AD	Always defects independent of previous outcome	(0,0,0,0)	DF11, DF15, FRD12, FY17, STS13
Always Cooperate	AC	Always cooperates independent of previous outcome	(1,1,1,1) (1,1,0,0)	DF11, DF15, FRD12, B15
Grim	G	Only cooperates after cc was last outcome	(1,0,0,0)	AF09, DF11, DF15
Tit-for-Tat	TFT	Only plays C if opponent did last period	(1,0,1,0)	FRD12, FY17, STS13
Win-stay-Lose-Shift (aka Perfect TFT)	WSLS	Plays same strategy if it was successful, otherwise shifts	(1,0,0,1)	DF11, DF15, FRD12, FY17
False cooperator	C-to-AD	Play c in first round then AD	–	FRD12, FY17
Explorative TFT	D-TFT	Play d in first round then TFT	–	DF15, FRD12, FY17
Alternator	DC-Alt	Play d in first round then alternate c and d	–	FRD12, FY17
Trigger-with-Reversion	GwR	Like Grim but revert to cooperation after cc^\ddagger	(1,0,0,0)	STS13
Pure Strategies Memory-2/3				
Trigger 2 periods	T2	Player punishes defection for max. 2 periods, otherwise cooperates	$(1,0,\sigma_1^*,0)$	DF11, FY17
Tit-for-2(3)-Tats	TF2T	Defects after 2 defections	$(1,\sigma_2^*,1,\sigma_2^*)$	FRD12, FY17
2-Tits-for-2-Tats	2TF2T	Defects twice after 2 defections	$(1,\sigma_3,\sigma_3,\sigma_3)$	FRD12, FY17
2-Tits-for-1-Tats	2TFT	Defects twice after each defections	$(1,0,\sigma_4,0)$	FRD12, FY17
Grim2(3/4)	G2(3)	After 2(3) defections will play D forever	$(1,\sigma_5,0,0)$	FRD12, FY17, STS13
Win-stay-Lose-Shift-2	WSLS2	cooperate after (dd,dd),(cc,cc), (dd,cc) otherwise defect	–	FRD12
Explorative TF2(3)T	D-TF2(3)T	Play D in first round then TF2(3)T	–	FRD12, FY17
Explorative Grim2(3)	D-Grim2(3)	Play D in first round then Grim2(3)	–	FRD12, FY17
Behavior Strategies				
Semi-Grim**	SG	Similar to Grim but may cooperate after CD or DC.	$(1,\sigma_{SG},\sigma_{SG},0)$	B15
Generous TFT	GTFT	Like TFT but cooperate with prob α after CD or DD	$(1,\sigma_{GT},1,\sigma_{GT})$	FRD12, B15

[†] σ assigns cooperation probabilities after joint cooperation (cc), unilateral defection by opponent (cd), unilateral defection (dc), and joint defection (dd).

[‡] possible if players make mistakes.

* Vector assigning cooperation probabilities $\in \{0, 1\}$ depending on the state 2 periods ahead.

** σ_{SG} and σ_{GT} are mixing parameters $\in (0, 1)$.

References: **AF09** (Aoyagi and Fréchette, 2009), **B15** (Breitmoser, 2015), **DF11** (Dal Bó and Fréchette, 2011), **DF15** (Dal Bó and Fréchette, 2019), **FRD12** (Fudenberg et al., 2012), **FY17** (Fréchette and Yuksel, 2017), **STS13** (Sherstyuk et al., 2013)

Table 14: Overview of the data sets used in the analysis

Experiment	Logistics		Parameters			Average cooperation rates				
	#Subj	#Dec	δ	g	l	$\hat{\sigma}_0$	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$
First halves per session										
<i>Aoyagi and Frechette (2009)</i>	38	1650	0.9	0.333	0.111	0.465	0.917	\gg 0.45	\approx 0.408	\approx 0.336
<i>Blonski et al. (2011)</i>	200	3040	0.756	1.345	2.602	0.295	0.89	\gg 0.279	\approx 0.193	\gg 0.034
<i>Bruttel and Kamecke (2012)</i>	36	1920	0.8	1.167	0.833	0.481	0.91	\gg 0.286	\approx 0.228	\gg 0.08
<i>Dal Bó (2005)</i>	102	1320	0.75	0.939	1.061	0.342	0.922	\gg 0.212	$<$ 0.342	\gg 0.089
<i>Dal Bó and Fréchet (2011)</i>	266	17772	0.622	1.062	1.072	0.31	0.951	\gg 0.334	\approx 0.331	\gg 0.063
<i>Dal Bó and Fréchet (2019)</i>	672	22112	0.743	1.579	1.341	0.451	0.94	\gg 0.297	\approx 0.335	\gg 0.057
<i>Dreber et al. (2008)</i>	50	2064	0.75	1.488	1.488	0.488	0.904	\gg 0.217	\approx 0.213	\gg 0.036
<i>Duffy and Ochs (2009)</i>	102	3128	0.9	1	1	0.53	0.904	\gg 0.301	\approx 0.33	\gg 0.111
<i>Fréchet and Yuksel (2017)</i>	50	800	0.75	0.4	0.4	0.737	0.943	\gg 0.141	\approx 0.266	\approx 0.091
<i>Fudenberg et al. (2012)</i>	48	1452	0.875	0.333	0.333	0.756	0.982	\gg 0.4	\approx 0.427	\gg 0.066
<i>Kagel and Schley (2013)</i>	114	7600	0.75	1	0.5	0.573	0.935	\gg 0.263	\approx 0.295	\gg 0.051
<i>Sherstyuk et al. (2013)</i>	56	3052	0.75	1	0.25	0.56	0.945	\gg 0.328	\approx 0.371	\gg 0.117
Pooled	1734	65910	0.728	1.207	1.083	0.389	0.938	\gg 0.304	\approx 0.322	\gg 0.065
Second halves per session										
<i>Aoyagi and Frechette (2009)</i>	38	1400	0.9	0.333	0.111	0.424	0.958	\gg 0.398	\approx 0.517	\approx 0.375
<i>Blonski et al. (2011)</i>	200	5460	0.766	1.282	2.554	0.279	0.923	\gg 0.287	\approx 0.231	\gg 0.02
<i>Bruttel and Kamecke (2012)</i>	36	1632	0.8	1.167	0.833	0.447	0.947	\gg 0.221	\approx 0.297	\gg 0.041
<i>Dal Bó (2005)</i>	102	1650	0.75	0.961	1.039	0.297	0.92	\gg 0.242	$<$ 0.388	\gg 0.064
<i>Dal Bó and Fréchet (2011)</i>	266	19270	0.62	1.122	1.103	0.355	0.979	\gg 0.376	\approx 0.362	\gg 0.041
<i>Dal Bó and Fréchet (2019)</i>	672	29480	0.766	1.666	1.386	0.469	0.976	\gg 0.315	$<$ 0.402	\gg 0.035
<i>Dreber et al. (2008)</i>	50	1838	0.75	1.533	1.533	0.461	0.917	\gg 0.128	\ll 0.39	\gg 0.009
<i>Duffy and Ochs (2009)</i>	102	6018	0.9	1	1	0.684	0.977	\gg 0.367	\approx 0.391	\gg 0.082
<i>Fréchet and Yuksel (2017)</i>	50	1568	0.75	0.4	0.4	0.763	0.97	\gg 0.233	\approx 0.398	\gg 0.069
<i>Fudenberg et al. (2012)</i>	48	1800	0.875	0.333	0.333	0.829	0.971	\gg 0.487	\approx 0.412	\gg 0.083
<i>Kagel and Schley (2013)</i>	114	7172	0.75	1	0.5	0.704	0.966	\gg 0.262	\approx 0.332	\gg 0.025
<i>Sherstyuk et al. (2013)</i>	56	2604	0.75	1	0.25	0.646	0.973	\gg 0.482	\approx 0.437	\gg 0.078
Pooled	1734	79892	0.744	1.271	1.172	0.404	0.971	\gg 0.327	$<$ 0.376	\gg 0.039

Note: The “average cooperation rates” are the relative frequencies estimated directly from the data. The relation signs encode bootstrapped p -values (resampling at the subject level with 10,000 repetitions) where $<$, $>$ indicate rejection of the Null of equality at $p < .05$ and \ll , \gg indicating $p < .002$. Following Wright (1992), we accommodate for the multiplicity of comparisons within data sets by adjusting p -values using the Holm-Bonferroni method (Holm, 1979). Note that all details here exactly replicate Breitmoser (2015). As a result, if a data set is considered in isolation, the .05-level indicated by “ $>$,” “ $<$ ” is appropriate. If all 24 treatments are considered simultaneously, the corresponding Bonferroni correction requires to further reduce the threshold to $.002 \approx .05/24$, which corresponds with “ \gg ,” “ \ll ”.

Table 15: Overview of cooperation rates in the data

Experiment	Cooperators							Defectors						
	#Subj	#Dec	Average cooperation rates					#Subj	#Dec	Average cooperation rates				
			$\hat{\sigma}_0$	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$			$\hat{\sigma}_0$	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$
First halves per session														
<i>Aoyagi and Frechette (2009)</i>	35	1509	0.783	0.936	0.45	0.402	0.313	3	141	0.143	0.575	0.444	0.441	0.486
<i>Blonski et al. (2011)</i>	74	1145	0.685	0.896	0.31	0.356	0.056	126	1895	0.066	0.714	0.192	0.123	0.027
<i>Bruttel and Kamecke (2012)</i>	20	1062	0.75	0.926	0.253	0.267	0.113	16	858	0.144	0.806	0.375	0.198	0.055
<i>Dal Bó (2005)</i>	52	675	0.807	0.947	0.21	0.37	0.133	50	645	0.087	0.762	0.22	0.326	0.064
<i>Dal Bó and Fréchet (2011)</i>	108	7382	0.699	0.969	0.337	0.415	0.113	158	10390	0.108	0.807	0.328	0.28	0.045
<i>Dal Bó and Fréchet (2019)</i>	311	10133	0.819	0.954	0.326	0.499	0.084	361	11979	0.124	0.87	0.239	0.239	0.048
<i>Dreber et al. (2008)</i>	31	1272	0.711	0.909	0.189	0.245	0.05	19	792	0.129	0.846	0.326	0.181	0.022
<i>Duffy and Ochs (2009)</i>	63	1886	0.807	0.913	0.302	0.403	0.14	39	1242	0.097	0.866	0.298	0.25	0.087
<i>Fréchet and Yuksel (2017)</i>	41	652	0.886	0.941	0.133	0.394	0.136	9	148	0.056	1	0.25	0.129	0.039
<i>Fudenberg et al. (2012)</i>	39	1185	0.905	0.985	0.418	0.518	0.06	9	267	0.091	0.947	0.316	0.333	0.077
<i>Kagel and Schley (2013)</i>	76	5066	0.814	0.939	0.262	0.419	0.069	38	2534	0.089	0.872	0.268	0.168	0.033
<i>Sherstyuk et al. (2013)</i>	34	1920	0.828	0.968	0.33	0.518	0.119	22	1132	0.152	0.78	0.323	0.266	0.115
Pooled	884	33887	0.778	0.951	0.312	0.43	0.098	850	32023	0.111	0.843	0.283	0.242	0.049
Second halves per session														
<i>Aoyagi and Frechette (2009)</i>	34	1245	0.959	0.968	0.382	0.578	0.328	4	155	0.211	0.75	0.448	0.371	0.469
<i>Blonski et al. (2011)</i>	66	1761	0.75	0.926	0.322	0.398	0.036	134	3699	0.049	0.91	0.189	0.164	0.015
<i>Bruttel and Kamecke (2012)</i>	15	656	0.893	0.954	0.136	0.613	0.031	21	976	0.129	0.922	0.351	0.211	0.044
<i>Dal Bó (2005)</i>	60	974	0.838	0.927	0.24	0.434	0.063	42	676	0.042	0.852	0.25	0.348	0.065
<i>Dal Bó and Fréchet (2011)</i>	111	7984	0.892	0.982	0.358	0.579	0.055	155	11286	0.081	0.948	0.406	0.286	0.038
<i>Dal Bó and Fréchet (2019)</i>	319	14330	0.897	0.978	0.312	0.585	0.067	353	15150	0.089	0.965	0.322	0.315	0.024
<i>Dreber et al. (2008)</i>	22	830	0.847	0.929	0.1	0.479	0.027	28	1008	0.125	0.833	0.195	0.344	0.002
<i>Duffy and Ochs (2009)</i>	69	4206	0.943	0.978	0.376	0.408	0.083	33	1812	0.124	0.968	0.348	0.373	0.081
<i>Fréchet and Yuksel (2017)</i>	42	1322	0.909	0.973	0.227	0.507	0.115	8	246	0	0.8	0.333	0.194	0.014
<i>Fudenberg et al. (2012)</i>	41	1542	0.957	0.969	0.465	0.456	0.106	7	258	0.065	1	0.6	0.325	0.053
<i>Kagel and Schley (2013)</i>	82	5176	0.949	0.968	0.242	0.505	0.035	32	1996	0.067	0.937	0.426	0.194	0.015
<i>Sherstyuk et al. (2013)</i>	37	1674	0.907	0.978	0.489	0.558	0.124	19	930	0.123	0.946	0.456	0.382	0.053
Pooled	898	41700	0.898	0.974	0.318	0.525	0.063	836	38192	0.084	0.954	0.347	0.292	0.03

Note: “Cooperators” and “Defectors” are determined by their average cooperation rate in round 1. If above median, they are cooperators. The “average cooperation rates” are the relative frequencies estimated directly from the data. The relation signs encode bootstrapped p -values (resampling at the subject level with 10,000 repetitions) where $<$, $>$ indicate rejection of the Null of equality at $p < .05$ and \ll , \gg indicating $p < .002$. Following Wright (1992), we accommodate for the multiplicity of comparisons within data sets by adjusting p -values using the Holm-Bonferroni method (Holm, 1979). Note that all details here exactly replicate Breitmoser (2015). As a result, if a data set is considered in isolation, the .05-level indicated by “ $>$,” “ $<$ ” is appropriate. If all 24 treatments are considered simultaneously, the corresponding Bonferroni correction requires to further reduce the threshold to $.002 \approx .05/24$, which corresponds with “ \gg ,” “ \ll ”.

C Robustness check: Memory-2

We investigate memory length using a data mining approach similar to above. To this end, we extend the set of pure strategies with memory-2 strategies and then evaluate these best fitting specifications, treatment by treatment, against the above memory-1 model from above. We follow Fudenberg et al. (2012), who introduced lenient and resilient variants of the pure memory-1 strategies, i.e., strategies that punish only after the second deviation or that punish for two rounds instead of one. In order to allow for memory-2 in semi-grim we use a novel approach by allowing the cooperation probabilities in round t to depend on the behavior of one or both players in $t - 2$ more generally with a parametric approach. Here, we allow for three different specifications: cooperation probabilities may be a function of the opponent's choice in $t - 2$ (*TFT-Scheme*), a function of whether both players cooperated in $t - 2$ or not (*Grim-Scheme*), or a function of the entire choice profile in $t - 2$ (*General scheme*).

Table 16 summarizes the results. First, we mine for mixtures of pure strategies, based on the list of 8 strategies (TFT, Grim, AD, Grim2, TF2T, T2, 2TFT, 2PTFT as in Fudenberg et al. (2012)). Given the above results, we assume that subjects do not switch strategies within half-sessions, as this comes without loss of descriptive adequacy for experienced subjects and only little loss for inexperienced subjects (for whom, however, memory-2 will turn out to be of negligible relevance). For each treatment, we determine the most adequate combination of strategies from a list of five possible combinations of Fudenberg et al.'s strategies, thus providing a selection of the best of 5^{32} models overall. The resulting model (Column "Best Pure M1&M2" in Table 16) fits highly significantly worse than the selection of pure and generalized-pure strategies with memory-1 defined above ("M1" in Table 16). We may therefore discard the possibility that subjects play pure strategies (with noise) of either memory-1 or memory-2, in favor of the possibility that they play strategies allowing for non-trivial randomization.

Second, we evaluate whether behavior strategies possibly have memory-2. That is, we compare the a simple SG+AD memory-1 version with the three generalizations to memory-2 introduced above. The results are report in the three left-most columns of Table 16 and appear clear-cut: None of the memory-2 extensions improves on describing behavior by the simple memory-1 AD+SG model and the AD+1.5SG model fits significantly better. Indeed, the finer the memory-2 ramifications, the worse the model adequacy (after accounting for the additional degrees of freedom). These results are additionally compatible with a result of Breitmoser (2015) who verified the Markov assumption by testing whether subjects systematically deviate from memory-1 strategies after particular histories in memory-2. We summarize these observations as follows.

Result (Memory-2) *Model adequacy does not improve by equipping subjects with memory-2, neither for (generalizations of) pure strategies nor for semi-grim.*

Table 17 shows aggregate state-wise cooperation rates for different lagged histories (cooperation or defection of the opponent in $t - 2$) *TFT-Scheme*. Table 18 shows aggregate state-wise cooperation rates for different lagged histories (joint cooperation or not in $t - 2$) *Grim-Scheme*.

Table 16: **Memory-1 or Memory-2, and semi-grim, pure or generalized pure?** Strategy mixtures are estimated treatment-by-treatment. The resulting ICL-BICs are pooled within experiments and overall (less is better, relation signs point to better models)

	Memory-2 Generalizations of Semi-Grim + AD			AD+SG	AD+1.5×SG	Best Gen Pure		Best Pure					
	M2“General”	M2“TFT”	M2“Grim”			M1	M1 & M2						
Specification													
# Models evaluated	1	1	1	1	1	9 ³²	5 ³²						
# Pars estimated (by treatment)	12	6	6	5	7	48	32						
# Parameters accounted for	12	6	6	5	7	6–10	3–8						
First halves per session													
<i>Aoyagi and Frechette (2009)</i>	756.04	≈	764.13	≈	749.99	≈	781.86	≈	792.28	≈	843.08	≈	884.86
<i>Blonski et al. (2011)</i>	1244.76	≫	1121.17	≈	1120.87	≫	1069.26	≈	1104.1	≈	1069.54	≈	1105.98
<i>Bruttel and Kamecke (2012)</i>	807.47	≈	802.89	≈	804.16	≈	800.12	≈	771.14	≈	821.99	≈	839.97
<i>Dal Bó (2005)</i>	660.68	>	641.34	≈	642.26	≈	629.13	≈	618.12	≈	623.19	<	653.05
<i>Dal Bó and Fréchet (2011)</i>	6671.28	≈	6616.44	≈	6604.7	≈	6597.92	>	6352.64	≪	6874.99	≪	7391.89
<i>Dal Bó and Fréchet (2019)</i>	8068.37	≈	8028.83	≈	8031.59	≈	8017.56	≫	7829.75	≪	8367.55	≪	8893.78
<i>Dreber et al. (2008)</i>	805.74	>	785.48	≈	785.6	≈	782.38	≈	764.44	≈	789.22	<	863.47
<i>Duffy and Ochs (2009)</i>	1361.84	≈	1377.17	≈	1369.86	≈	1372.97	≈	1361.13	≈	1396.68	≈	1426.34
<i>Fréchet and Yuksel (2017)</i>	305.9	≈	299.72	≈	296.93	≈	299.6	≈	289.54	≈	300.87	≈	317.35
<i>Fudenberg et al. (2012)</i>	387.8	≈	379.84	≈	378.07	≈	381.01	≈	377.96	<	437.5	≈	463.4
<i>Kagel and Schley (2013)</i>	2542.02	≈	2556.45	≈	2552.09	≈	2561.76	≫	2450.24	≪	2660.58	≈	2730.66
<i>Sherstyuk et al. (2013)</i>	1311.64	≈	1307.45	≈	1303.94	≈	1303.8	≈	1234.52	<	1299.14	<	1398.69
Pooled	25434.21	≫	24972.71	≈	24931.86	≈	24779.74	≫	24201.19	≪	25710.49	≪	27115.39
Second halves per session													
<i>Aoyagi and Frechette (2009)</i>	415.47	≈	421.18	>	409.2	≈	423.68	≈	421.21	<	492.28	≈	540.47
<i>Blonski et al. (2011)</i>	1518.54	≫	1395.94	≈	1393.41	≫	1352.31	≈	1373.9	≈	1373.41	<	1564.48
<i>Bruttel and Kamecke (2012)</i>	536.19	≈	532.08	≈	529.47	≈	536.78	≫	480.47	≈	493.79	<	567.99
<i>Dal Bó (2005)</i>	727.25	≈	710.88	≈	708.32	≈	699.05	≈	677.22	<	712.63	≈	741.2
<i>Dal Bó and Fréchet (2011)</i>	5201.05	≈	5137.82	≈	5132.96	≈	5128.68	≫	4565.87	≪	5045.08	≪	5960.78
<i>Dal Bó and Fréchet (2019)</i>	7840.87	≈	7829.51	≈	7808.63	≈	7825.98	≫	7306.18	≪	8107.4	≪	9143.98
<i>Dreber et al. (2008)</i>	597.17	≈	580.63	≈	570.33	≈	589.83	>	544.66	≈	580.69	<	648.55
<i>Duffy and Ochs (2009)</i>	1706.1	≈	1753.41	≈	1719.86	≈	1761.61	≫	1656.55	≪	1910.09	≈	2003.41
<i>Fréchet and Yuksel (2017)</i>	422.32	≈	424.41	≈	419.44	≈	423.34	≈	422.52	≈	433.18	<	464.23
<i>Fudenberg et al. (2012)</i>	452.64	≈	450.08	≈	447.25	≈	452.6	≈	433.74	≪	505.1	≈	534.47
<i>Kagel and Schley (2013)</i>	1782.43	≈	1777.83	≈	1773.55	≈	1775.62	≫	1572.95	<	1681.03	≈	1830.26
<i>Sherstyuk et al. (2013)</i>	959.21	≈	952.56	≈	949.46	≈	951.34	≫	834.73	≈	890.27	<	1023.43
Pooled	22669.9	≫	22258.14	≈	22153.7	≈	22103.2	≫	20545.33	≪	22462.25	≪	25177.57

Note: The main body contains ICL-BICs aggregated at paper level. Relation signs and p -values are exactly as above, see Table 5. “M2” (“M1”) denotes strategies, whose actions may depend on actions in $t - 2$ and $t - 1$ ($t - 1$ only). The supplements “General”, “TFT”, “Grim” indicate whether parameters of behavior strategies may depend on: all four possible histories in $t - 2$ (M2 “General”), whether the opponent cooperated in $t - 2$ (M2 “TFT”), or whether there was joint cooperation in $t - 2$ (M2 “Grim”). Pure M2 strategies do not have such free parameters. Columns 1-3 contain one memory-2 version of semi-grim each. Column 4-5 are memory-1 models containing semi-grim and always defect. Column 6 contains the best mixture of generalized pure memory-1 strategies (identical to column 2 (“No Switching”) in Table 5) and the last column contains the best fitting combinations of a set of pure memory-1 and memory-2 strategies from the literature (TFT, Grim, AD, Grim2, TF2T, T2, 2TFT, 2PTFT) for definitions see Table 13 in the Appendix.

Table 17: Strategies as a function of behavior in $t - 2$ (TFT scheme)

Experiment	Cooperation after $\emptyset, (c, c), (d, c)$ in $t - 2$				Cooperation after $(c, d), (d, d)$ in $t - 2$									
	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$						
First halves per session														
<i>Aoyagi and Fréchette (2009)</i>	0.93	>>	0.439	≈	0.388	≈	0.434	0.789	>>	0.463	≈	0.44	>	0.291
<i>Blonski et al. (2011)</i>	0.901	>>	0.27	≈	0.146	>>	0.053	0.667	≈	0.296	≈	0.321	>>	0.027
<i>Bruttel and Kamecke (2012)</i>	0.908	>>	0.312	≈	0.218	≈	0.151	0.944	>>	0.247	≈	0.247	>>	0.063
<i>Dal Bó (2005)</i>	0.93	>>	0.232	≈	0.31	>	0.126	0.833	>	0.147	≈	0.413	>>	0.071
<i>Dal Bó and Fréchette (2011)</i>	0.955	>>	0.352	≈	0.298	>>	0.086	0.885	>>	0.291	≈	0.41	>>	0.048
<i>Dal Bó and Fréchette (2019)</i>	0.944	>>	0.301	≈	0.277	>>	0.098	0.847	>>	0.288	≈	0.44	>>	0.044
<i>Dreber et al. (2008)</i>	0.902	>>	0.213	≈	0.189	>>	0.061	1	>	0.233	≈	0.302	>>	0.025
<i>Duffy and Ochs (2009)</i>	0.927	>>	0.316	≈	0.304	≈	0.232	0.691	>>	0.277	≈	0.361	>>	0.08
<i>Fréchette and Yuksel (2017)</i>	0.943	>>	0.153	≈	0.241	≈	0.1	1	≈		≈	0.4	≈	0.086
<i>Fudenberg et al. (2012)</i>	0.984	>>	0.394	≈	0.347	>>	0.05	0.895	>>	0.41	≈	0.579	>>	0.069
<i>Kagel and Schley (2013)</i>	0.94	>>	0.29	≈	0.25	>>	0.125	0.787	>>	0.196	≈	0.402	>>	0.032
<i>Sherstyuk et al. (2013)</i>	0.951	>>	0.329	≈	0.341	>	0.186	0.844	>>	0.328	≈	0.424	>>	0.09
Pooled	0.944	>>	0.312	>	0.279	>>	0.106	0.826	>>	0.287	≈	0.41	>>	0.05
Second halves per session														
<i>Aoyagi and Fréchette (2009)</i>	0.961	>>	0.408	≈	0.567	≈	0.447	0.867	>>	0.381	≈	0.451	≈	0.328
<i>Blonski et al. (2011)</i>	0.922	>>	0.224	≈	0.195	>>	0.029	0.944	>>	0.402	≈	0.324	>>	0.018
<i>Bruttel and Kamecke (2012)</i>	0.948	>>	0.239	≈	0.214	≈	0.118	0.923	>	0.167	≈	0.5	>>	0.018
<i>Dal Bó (2005)</i>	0.919	>>	0.264	≈	0.39	>>	0.113	0.938	>>	0.175	≈	0.383	>>	0.047
<i>Dal Bó and Fréchette (2011)</i>	0.979	>>	0.391	≈	0.29	>>	0.075	0.975	>>	0.334	≈	0.547	>>	0.022
<i>Dal Bó and Fréchette (2019)</i>	0.977	>>	0.304	≈	0.328	>>	0.064	0.927	>>	0.343	≈	0.532	>>	0.028
<i>Dreber et al. (2008)</i>	0.917	>>	0.111	<	0.311	>>	0.005	0.909	>	0.5	≈	0.629	>>	0.01
<i>Duffy and Ochs (2009)</i>	0.98	>>	0.408	≈	0.371	>	0.232	0.849	>>	0.316	≈	0.415	>>	0.058
<i>Fréchette and Yuksel (2017)</i>	0.973	>>	0.213	≈	0.286	≈	0.214	0.818	≈	0.286	≈	0.575	>>	0.038
<i>Fudenberg et al. (2012)</i>	0.974	>>	0.5	≈	0.41	>>	0.111	0.84	>	0.463	≈	0.417	>>	0.075
<i>Kagel and Schley (2013)</i>	0.967	>>	0.281	≈	0.263	>>	0.061	0.872	>>	0.188	≈	0.527	>>	0.018
<i>Sherstyuk et al. (2013)</i>	0.973	>>	0.503	≈	0.417	>>	0.12	0.968	>>	0.431	≈	0.5	>>	0.062
Pooled	0.973	>>	0.325	≈	0.315	>>	0.076	0.917	>>	0.332	≈	0.499	>>	0.028

Note: Relation signs, bootstrap procedure, and derived p -values are exactly as above, see Table 2, with the obvious adaptation that the Holm-Bonferroni correction now applies to all eight tests per data set.


Table 18: Strategies as a function of behavior in $t - 2$ (Grim scheme)

Experiment	Cooperation after $\emptyset, (c, c)$ in $t - 2$				Cooperation after $(c, d), (d, c), (d, d)$ in $t - 2$			
	$\hat{\delta}_{cc}$	$\hat{\delta}_{cd}$	$\hat{\delta}_{dc}$	$\hat{\delta}_{dd}$	$\hat{\delta}_{cc}$	$\hat{\delta}_{cd}$	$\hat{\delta}_{dc}$	$\hat{\delta}_{dd}$
First halves per session								
<i>Aoyagi and Fréchette (2009)</i>	0.939	>> 0.39	≈ 0.439	≈ 0.556	0.782	>> 0.485	≈ 0.39	> 0.32
<i>Blonski et al. (2011)</i>	0.903	>> 0.248	≈ 0.174	>> 0.045	0.714	> 0.318	≈ 0.216	> 0.031
<i>Bruttel and Kamecke (2012)</i>	0.919	>> 0.296	≈ 0.245	≈ 0.179	0.833	>> 0.278	≈ 0.213	>> 0.071
<i>Dal Bó (2005)</i>	0.926	>> 0.184	≈ 0.31	≈ 0.143	0.889	>> 0.254	≈ 0.39	>> 0.074
<i>Dal Bó and Fréchette (2011)</i>	0.961	>> 0.342	≈ 0.307	>> 0.081	0.849	>> 0.324	≈ 0.364	>> 0.054
<i>Dal Bó and Fréchette (2019)</i>	0.95	>> 0.265	≈ 0.301	>> 0.081	0.843	>> 0.328	≈ 0.369	>> 0.052
<i>Dreber et al. (2008)</i>	0.901	>> 0.154	≈ 0.217	>> 0.062	1	>> 0.359	≈ 0.203	>> 0.031
<i>Duffy and Ochs (2009)</i>	0.932	>> 0.218	≈ 0.301	≈ 0.208	0.748	>> 0.361	≈ 0.35	>> 0.102
<i>Fréchette and Yuksel (2017)</i>	0.942	>> 0.132	≈ 0.245	>> 0	1	≈ 0.182	≈ 0.364	≈ 0.111
<i>Fudenberg et al. (2012)</i>	0.985	>> 0.429	≈ 0.408	>> 0	0.921	>> 0.377	≈ 0.443	>> 0.068
<i>Kagel and Schley (2013)</i>	0.947	>> 0.236	≈ 0.288	>> 0.133	0.763	>> 0.298	≈ 0.305	>> 0.042
<i>Sherstyuk et al. (2013)</i>	0.953	>> 0.312	≈ 0.395	>> 0.172	0.875	>> 0.343	≈ 0.349	>> 0.107
Pooled	0.949	>> 0.278	≈ 0.3	>> 0.091	0.825	>> 0.333	≈ 0.346	>> 0.059
Second halves per session								
<i>Aoyagi and Fréchette (2009)</i>	0.965	>> 0.438	≈ 0.625	≈ 0.333	0.846	>> 0.371	< 0.443	≈ 0.378
<i>Blonski et al. (2011)</i>	0.922	>> 0.157	≈ 0.232	>> 0.027	0.941	>> 0.425	≈ 0.23	>> 0.019
<i>Bruttel and Kamecke (2012)</i>	0.946	>> 0.156	≈ 0.233	≈ 0.173	0.958	>> 0.327	≈ 0.4	>> 0.019
<i>Dal Bó (2005)</i>	0.918	>> 0.178	< 0.4	> 0.131	0.937	>> 0.32	≈ 0.373	>> 0.052
<i>Dal Bó and Fréchette (2011)</i>	0.981	>> 0.373	≈ 0.323	>> 0.077	0.95	>> 0.38	≈ 0.416	>> 0.025
<i>Dal Bó and Fréchette (2019)</i>	0.98	>> 0.264	< 0.366	>> 0.058	0.904	>> 0.369	≈ 0.44	>> 0.031
<i>Dreber et al. (2008)</i>	0.913	>> 0.029	≪ 0.314	>> 0.007	0.955	>> 0.417	≈ 0.611	>> 0.009
<i>Duffy and Ochs (2009)</i>	0.981	>> 0.362	≈ 0.433	≈ 0.226	0.889	>> 0.369	≈ 0.368	>> 0.077
<i>Fréchette and Yuksel (2017)</i>	0.976	>> 0.173	≈ 0.308	≈ 0.222	0.75	> 0.294	≈ 0.49	>> 0.06
<i>Fudenberg et al. (2012)</i>	0.976	>> 0.473	≈ 0.509	≈ 0.2	0.854	>> 0.5	≈ 0.328	>> 0.077
<i>Kagel and Schley (2013)</i>	0.969	>> 0.218	≈ 0.293	> 0.098	0.868	>> 0.332	≈ 0.394	>> 0.02
<i>Sherstyuk et al. (2013)</i>	0.974	>> 0.465	≈ 0.486	>> 0.107	0.952	>> 0.505	≈ 0.369	>> 0.072
Pooled	0.975	>> 0.282	≪ 0.351	>> 0.07	0.908	>> 0.378	≈ 0.404	>> 0.033

Note: Relation signs, bootstrap procedure, and derived p -values are exactly as above, see Table 2, with the obvious adaptation that the Holm-Bonferroni correction now applies to all eight tests per data set.

D Further robustness checks for Sections 4 and 5

Figure 4: Relation of $\delta - \delta^*$ to shares of cooperators (first halves of sessions)



Note: This figure shows how the ratios of the three strategies – defectors, cautious cooperators, and strong cooperators change with the distance of δ to the BOS cooperation threshold δ^* across treatments. The solid line represents the best fitting logistic curve estimated with intercept. Panel (a) displays the total share of both cooperators, panel (b) the relative share of cautious cooperators among cooperators, panel (c) the share of cautious cooperators overall, panel (d) the share of strong cooperators overall.

Table 19: Robustness of Table 5 towards “30 rounds of play till experience”: **Best mixtures of pure or generalized strategies in relation to behavior strategy.**

	Baseline Model		Best mixture of pure or generalized strategies					Unrestr Beh 2×P5+AD	Best Mixture				
			No Switching	Random Switching	Markov Switching	Best Switching	Best Switching		Best Switching	By Treatment			
Specification													
# Models evaluated	1		9 ³²	9 ³²	9 ³²	9 ³²	3 × 9 ³²	1		27 ³² ≈ 10 ⁴⁶			
# Pars estimated (by treatment)	5		48	48	180	276	276	13		276			
# Parameters accounted for	5		3–10	3–10	12-35	3–30	3–30	13		3–30			
First halves per session													
<i>Aoyagi and Frechette (2009)</i>	704.77	≈	679.07	≈	667.62	≈	684.64	≈	684.64	≫	603.81	<	667.62
<i>Blonski et al. (2011)</i>	2017.43	>	1942.75	<	2061.54	≈	2102.97	≈	2102.97	≫	1993.47	≈	1941.45
<i>Bruttel and Kamecke (2012)</i>	568.74	≈	549.7	≈	560.79	≈	543.15	≈	543.15	≈	538.89	≈	543.15
<i>Dal Bó (2005)</i>	1225.33	≈	1210.87	≪	1366.68	≫	1210.8	≈	1210.8	>	1161.77	≈	1192.55
<i>Dal Bó and Fréchette (2011)</i>	3884.21	>	3735.25	≪	3911.56	≫	3686.29	≈	3686.29	>	3535.53	<	3633.61
<i>Dal Bó and Fréchette (2019)</i>	8963.07	≫	8730.66	≪	9631.78	≫	8233.2	≈	8233.2	≈	8098.25	≈	8115.28
<i>Dreber et al. (2008)</i>	580.87	≈	559.49	≈	584.75	≈	534.29	≈	534.29	≈	549.76	≈	534.29
<i>Duffy and Ochs (2009)</i>	1677.02	≈	1665.73	≈	1720.54	≈	1605.79	≈	1605.79	≫	1481.56	≪	1605.79
<i>Fréchette and Yuksel (2017)</i>	603.84	≈	590.89	<	649.85	>	568.13	≈	568.13	≈	547.57	≈	568.13
<i>Fudenberg et al. (2012)</i>	455.89	≈	455.89	≈	459.83	≈	457.81	≈	457.81	≫	394.37	<	455.89
<i>Kagel and Schley (2013)</i>	1621.93	≈	1621.93	<	1754.14	≈	1542.63	≈	1542.63	≈	1489.87	≈	1542.63
<i>Sherstyuk et al. (2013)</i>	919.02	≈	914.61	<	967.22	>	922.45	≈	922.45	≫	835.9	≪	914.61
Pooled	23404.48	≫	22898.28	≪	24522.31	≈	22773.53	≈	22773.53	≈	21704.93	<	22161.03
Second halves per session													
<i>Aoyagi and Frechette (2009)</i>	719.05	≈	656.92	≈	656.65	≈	646.25	≈	656.92	≫	545.5	<	646.25
<i>Blonski et al. (2011)</i>	615.13	≈	529.97	≈	528.82	≪	701.88	≈	529.97	≪	725.95	≈	489.09
<i>Bruttel and Kamecke (2012)</i>	766.66	≈	754.7	≪	885.07	≈	734.16	≈	754.7	>	690.3	≈	734.16
<i>Dal Bó (2005)</i>	196.84	≈	148.82	≈	105.12	≪	152.48	≈	148.82	≪	198.02	≈	105.12
<i>Dal Bó and Fréchette (2011)</i>	8850.02	≈	8206.12	≪	9898.77	≈	7445.82	≈	8206.12	≈	7300.56	≈	7445.82
<i>Dal Bó and Fréchette (2019)</i>	7851.06	≈	7644.65	≪	8960.18	≈	7778.52	≈	7644.65	≈	6865.95	≪	7410.23
<i>Dreber et al. (2008)</i>	832.58	>	778.4	≪	930.98	≈	751.43	≈	778.4	>	723.17	≈	751.43
<i>Duffy and Ochs (2009)</i>	1692.08	≈	1663.82	≈	1738.03	≈	1647.57	≈	1663.82	≈	1404.25	≪	1647.57
<i>Fréchette and Yuksel (2017)</i>	158.72	>	133.03	≈	153.88	≈	151.25	≈	133.03	≈	134.58	≈	133.03
<i>Fudenberg et al. (2012)</i>	486.81	≈	486.81	≈	515.02	≈	505.03	≈	486.81	≈	402.93	<	486.81
<i>Kagel and Schley (2013)</i>	2690.35	≈	2690.35	>	2347.19	≈	2481.64	≈	2690.35	≈	2419.78	≈	2347.19
<i>Sherstyuk et al. (2013)</i>	1292.68	≈	1267.42	≈	1137.49	≈	1229.61	≈	1267.42	≈	1182.85	≈	1137.49
Pooled	26334.35	≫	25193.46	≪	28083.76	≈	24996.21	≈	25193.46	≈	23068.02	<	23874.33

Note: For a detailed description of the columns, see Table 5. This table differs from table 5 only by the definition of experience. Here we call subjects experienced for supergames starting after round 30. We call them inexperienced in all other supergames.

Table 20: Robustness of Table 6 towards “30 rounds of play till experience”: Examining heterogeneity of cooperating subjects and semi-grim structure of their strategies

	Best Mixture		Fixed SG		Treatment-specific SG specification										
	Best Switching		1.5×SG+ AD		1.5×SG+AD	2×SG+AD	3×SG+AD	3×P5+AD	2×P5+AD	P5+AD					
Specification															
# Models evaluated	27 ³² ≈ 10 ⁴⁶		1		1	1	1	1	1	1	1	1	1		
# Pars estimated (by treatment)	276		3		7	9	13	19	17	17	17	11	11		
# Parameters accounted for	3-30		3		7	9	13	19	17	17	17	11	11		
First halves per session															
<i>Aoyagi and Frechette (2009)</i>	667.62	≈	637.07	≈	643.49	≈	629.81	≈	634.94	>	605.64	≈	603.81	≈	602.46
<i>Blonski et al. (2011)</i>	1941.45	≈	1898.44	≈	1908.87	≈	1919.49	<	1987.72	≪	2085.1	≫	1993.47	≈	1968.36
<i>Bruttel and Kamecke (2012)</i>	543.15	≈	522.33	≈	524.84	≈	538.67	≈	524.91	≈	531.34	≈	538.89	≈	551.79
<i>Dal Bó (2005)</i>	1192.55	≈	1164.98	≈	1164.09	≈	1187.03	≈	1177.08	≈	1170.35	≈	1161.77	<	1219.45
<i>Dal Bó and Fréchet (2011)</i>	3633.61	≈	3521.82	≈	3544.24	≈	3536.45	≈	3559.72	<	3611.59	>	3535.53	≈	3589.33
<i>Dal Bó and Fréchet (2019)</i>	8115.28	≈	8284.59	>	8202.48	≈	8139.52	≈	8142.63	≈	8130.33	≈	8098.25	≪	8326.76
<i>Dreber et al. (2008)</i>	534.29	≈	541.42	≈	540.96	≈	542.95	≈	556.19	≈	559.2	≈	549.76	≈	538.03
<i>Duffy and Ochs (2009)</i>	1605.79	≈	1537.47	>	1488.44	<	1548.72	>	1486.31	≈	1484.36	≈	1481.56	≪	1571.38
<i>Fréchet and Yuksel (2017)</i>	568.13	≈	580.86	≈	577.75	≈	556.27	≈	554.35	≈	551.45	≈	547.57	<	575.86
<i>Fudenberg et al. (2012)</i>	455.89	≈	444.08	>	398.99	≈	391.72	≈	402.04	≈	401.02	≈	394.37	≈	399.46
<i>Kagel and Schley (2013)</i>	1542.63	≈	1516.89	≈	1523.91	≈	1492.6	≈	1483.6	≈	1473.5	≈	1489.87	<	1561.95
<i>Sherstyuk et al. (2013)</i>	914.61	≈	869.61	≈	866.61	≈	841.39	≈	833.88	≈	829.05	≈	835.9	≪	889.51
Pooled	22161.03	>	21628.99	≈	21640.02	≈	21652.87	≈	21817.56	≪	22125.94	≫	21704.93	≪	22049.69
Second halves per session															
<i>Aoyagi and Frechette (2009)</i>	646.25	≈	610.11	≫	569.62	≈	566.52	≈	564.62	≈	541.38	≈	545.5	≈	548.76
<i>Blonski et al. (2011)</i>	489.09	<	534.49	≪	614.55	≪	648.01	≪	753.61	≪	899.58	≫	725.95	≫	572.14
<i>Bruttel and Kamecke (2012)</i>	734.16	≈	706.9	≈	699.2	≈	697.16	≈	688.82	≈	659.27	≈	690.3	<	764.9
<i>Dal Bó (2005)</i>	105.12	≪	157.19	<	167.69	>	160.88	≪	224.9	≪	249.73	≫	198.02	≫	127.22
<i>Dal Bó and Fréchet (2011)</i>	7445.82	≪	7944.17	≫	7429.41	≈	7340.41	≈	7179.94	≈	7129.04	<	7300.56	≪	8095.2
<i>Dal Bó and Fréchet (2019)</i>	7410.23	≈	7515.99	≫	6964.58	≈	6990.43	≫	6833.14	≈	6826.27	≈	6865.95	≪	7367.82
<i>Dreber et al. (2008)</i>	751.43	≈	769.57	≈	759.05	≈	746.3	≈	740.91	≈	715.77	≈	723.17	<	806.32
<i>Duffy and Ochs (2009)</i>	1647.57	>	1544.71	>	1448.68	≈	1419.69	>	1348.05	≈	1343.95	≈	1404.25	≪	1499.54
<i>Fréchet and Yuksel (2017)</i>	133.03	≈	148.65	>	115.87	≈	126.62	≈	123.56	≈	130.91	≈	134.58	≈	124.53
<i>Fudenberg et al. (2012)</i>	486.81	≈	471.81	≫	415.35	≈	397.53	≈	387.56	≈	393.1	≈	402.93	≈	427.25
<i>Kagel and Schley (2013)</i>	2347.19	<	2625.27	≫	2451.52	≈	2428.78	>	2346.78	≈	2336.24	≈	2419.78	≪	2728.25
<i>Sherstyuk et al. (2013)</i>	1137.49	<	1293.71	>	1211.48	≈	1193.64	≈	1163.22	≈	1170.62	≈	1182.85	≪	1325
Pooled	23874.33	<	24431.99	≫	23102.33	≈	23044.26	>	22829.28	<	23088.88	≈	23068.02	≪	24642.26

Note: For a detailed description of the columns, see Table 6. This table differs from Table 6 only by the definition of experience. Here we call subjects experienced for supergames starting after round 30. We call them inexperienced in all other supergames.