Cooperation, but no reciprocity: Individual strategies in the repeated Prisoner's Dilemma

Breitmoser, Yves

28 May 2013

Cooperation, but no reciprocity: Individual strategies in the repeated Prisoner's Dilemma

Yves Breitmoser*

EUV Frankfurt (Oder)

May 28, 2013

Abstract

In the repeated Prisoner's Dilemma, predictions are notoriously difficult. Recently, however, Blonski, Ockenfels, and Spagnolo (2011, BOS) showed that experimental subjects predictably cooperate when the discount factor exceeds a particular threshold. I show that this threshold implies existence of an equilibrium robust to two standard refinement assumptions (utility perturbations and imperfect monitoring). The equilibrium is "Semi-Grim": Cooperate after mutual cooperation, defect after mutual defection, randomize otherwise. Testing six resulting predictions on existing data, comprising 37.000 observations, I then find that subjects indeed play Semi-Grim strategies, and switch to cooperation in round 1, when the former turn into equilibria—at the BOS-threshold.

JEL–Codes: C72, C73, C92

Keywords: Repeated Prisoner's Dilemma, experiment, equilibrium selection, cooperative behavior, reciprocity, belief-free equilibria, robustness

---

*I thank Friedel Bolle, Andreas Diekmann, Jörg Oechssler, Georg Weizsäcker, participants of the Social Science Workshop in Jena, and seminar participants from HSU Hamburg and EUV Frankfurt (Oder) for helpful comments. Financial funding of the DFG (BR 4648/1–1) is gratefully acknowledged. Address: Postfach 1786, 15207 Frankfurt(Oder), Germany, email: breitmoser@euv-frankfurt-o.de, Telephone/Fax: +3355534 2291/2390

1 Introduction

Most economic interactions are long-term relationships, and the welfare generated in long-term relationships depends primarily on whether agents manage to sustain mutual cooperation. Since this applies to all forms of relationships, including personal, industrial, and multi-national ones, the emergence of cooperation has been studied extensively. The theoretical results are well known for underlining the diversity of potential equilibrium outcomes. This diversity is theoretically robust to impatience, renegotiation, and various equilibrium refinement concepts.[1] In line with these predictions, experimental analyses showed that the existence of cooperative equilibria is necessary but not sufficient for cooperative behavior to emerge (e.g. Dal Bo, 2005).

Recently, Blonski, Ockenfels, and Spagnolo (2011, BOS) and Dal Bo and Fréchette (2011) showed that an empirically *predictable* condition for cooperation exists nonetheless. In experiments on the repeated Prisoner's Dilemma, they showed that cooperation sets in once the discount factor δ exceeds a certain threshold above the theoretically necessary condition for δ. It is currently unclear whether this threshold has a strategic interpretation, i.e. whether it is related to existence of a particular class of equilibria. The purpose of the present paper is to show that this threshold is the existence condition of a particular class of Markov perfect equilibria and that subjects indeed start playing the respective strategies when they turn into equilibria. These equilibria differ from those discussed in the literature, they are mixed and lack direct reciprocity, but they are played robustly across treatments in recent experiments.

BOS derived the threshold from conditions relating δ to the base game payoffs, based on the novel idea that the emergence of cooperation depends also on the sucker's payoff in the base game. The lower the sucker's payoff, the higher is the risk of unilateral cooperation (as the opponent may defect), and hence the lower one's inclination to actually cooperate. Since the sucker's payoff is irrelevant with respect to the existence of pure-strategy equilibria such as Grim, this result broadens the perspective and its empirical robustness shows that subjects' strategies are more predictable than Folk

[1]The Folk theorem of Fudenberg and Maskin (1986) shows that all individually rational payoff profiles may be sustained along the equilibrium path if the discount factor is sufficiently close to 1. The range of equilibrium outcomes shrinks only "slightly" if discounting is not just infinitesimal (Abreu et al., 1990; Stahl et al., 1991), if players may renegotiate (Evans and Maskin, 1989; Van Damme, 1989), if monitoring is imperfect (Ely and Välimäki, 2002), or if robustness to utility perturbations is required ("purifiability", Doraszelski and Escobar, 2010).

theorems and their refinements suggest. The BOS axioms do not lend themselves to a strategic interpretation, however, and the recent econometric analyses of Dal Bo and Fréchette (2011, DF) and Fudenberg, Rand and Dreber (2012, FRD) did not identify an empirically robust strategy.

In my theoretical analysis, I show that a Markov perfect equilibrium (MPE) that is robust to two standard refinement assumptions exists if and only if the discount factor δ exceeds the BOS-threshold. These refinement assumptions are utility perturbations and imperfect monitoring. On their own, both of them have been investigated extensively in the existing literature. Robustness to utility perturbations (i.e. "purifiability", Doraszelski and Escobar, 2010) isolates equilibria that continue to exist if the opponent has payoff parameters that differ stochastically from those in the payoff matrix. A special case are logit equilibria, and following McKelvey and Palfrey (1995), logit equilibria have been applied successfully in many analyses of experimental games.[2] I am considering the limit of logit equilibria as noise disappears, i.e. limiting-logit equilibria. Robustness to imperfect monitoring, in turn, isolates equilibria that are robust to the possibility that players forget the current state, possibly due to imperfect attention in the experiment or due to imperfect recall in general. A Markov perfect equilibrium is robust to imperfect monitoring if it is completely mixed, i.e. "belief-free" (Ely et al., 2005). While the intersection of equilibria that are both generically purifiable and belief-free is empty (Bhaskar et al., 2008), I show that an equilibrium that is limiting-logit and belief-free exists iff δ exceeds the BOS-threshold. This equilibrium is in *Semi-Grim* strategies: Cooperate with high probability if both players had cooperated in the previous round, with low probability after mutual defection, and with intermediate but equal probabilities if exactly one player had cooperated in the previous round—regardless of who had cooperated, i.e. without direct reciprocity.

Re-analyzing the data from four recent experiments on the repeated Prisoner's Dilemma, I then show that subjects indeed play Semi-Grim strategies. In addition to the data of BOS and FD, which in aggregate contain 16 treatments where δ is comparably near the BOS-threshold, I also consider Duffy and Ochs (2009) and FRD, which contain one treatment each where δ is substantially above the BOS-threshold (the remaining treatments of their experiments are non-standard repeated games with either

[2] This includes the centipede game (Fey et al., 1996), the traveler's dilemma (Capra et al., 1999), auctions (Goeree et al., 2002b), public goods games (Goeree et al., 2002a), monotone contribution games (Choi et al., 2008), and beauty contests (Breitmoser, 2012).

fixed rematching or exogenous noise). This provides two robustness checks. I use the raw data from the second halves of these experiments, when learning has largely stabilized. The analyzed data set thus contains about 37.000 decisions of 748 subjects in four different experiments and 18 different treatments.

The main results are summarized briefly. Behavior is well-described by memory-1 Markov strategies. In round 1, behavior adapts strongly to the treatment parameters and cooperation starts to occur systematically (in 50% of the games) almost exactly as δ crosses the BOS-threshold. The continuation strategies are fairly invariant to the treatment parameters and have the outlined Semi-Grim structure across treatments. The predicted absence of direct reciprocity is confirmed. I then analyze individual strategies by estimating the latent weights of standard equilibrium strategies and Semi-Grim MPEs, and find that subjects switch strategies as δ increases. Starting with Always-Defect, they move via Grim toward Semi-Grim MPEs. This is compatible with the observation that the average strategy is Semi-Grim in most treatments. As δ crosses the BOS-threshold, about 50% of the population plays Semi-Grim MPEs, and further above the threshold, they play Semi-Grim MPEs almost exclusively.

Thus, subjects switch to cooperation in round 1 and to Semi-Grim in the continuation simultaneously, when the limiting-logit, belief-free MPE appears. This explains the emergence of cooperation at the BOS-threshold. It also shows that strategies in repeated games are robust—behaviorally robust across treatments and strategically robust to utility perturbations and imperfect monitoring—and indeed predictable, at least to the degree that mixed strategies are predictable.[3] Both strategies and behavior in repeated games thus appear to be far more predictable (and systematic) than previously assumed. It seems reasonable to expect that the results extend to general repeated games, as the discussed refinement concepts generalize. Potential applications thus range from interpersonal relationships over repeated public goods problems to industrial competition.

Section 2 reviews the literature and provides a first look at the data. Sections 3 and 4 derive testable predictions from the refinements concepts considered here. Sections 5 and 6 test the predictions with respect to aggregate and individual behavior, respectively. Section 7 concludes. The supplementary material contains proofs, robustness checks, and parameter estimates.

[3] The observation that the strategies are mixed also explains that the sucker's payoff is of strategic relevance, since it implies that one may end up cooperating unilaterally.

Figure 1: Prisoner's dilemma games ($p_{dc} > p_{cc} > p_{dd} > p_{cd}$ and $b > a > 1$)

(a) Prisoner's Dilemma

	c	d
c	p_{cc}, p_{cc}	p_{cd}, p_{dc}
d	p_{dc}, p_{cd}	p_{dd}, p_{dd}

(b) "Standardized" PD

	c	d
c	a, a	$0, b$
d	$b, 0$	$1, 1$

2 Related literature and a first look at the data

Early experimental research, such as Roth and Murnighan (1978) and Murnighan and Roth (1983), showed that cooperation increases when equilibria sustaining cooperation exist. The existence of cooperative equilibria did not seem to yield substantial or even robust cooperation, however. Dal Bo (2005) discusses this literature in detail and conducted a novel series of experiments showing that behavior is largely in line with theoretical predictions. Cooperative equilibria exist if the discount factor satisfies $\delta \geq (p_{dc} - p_{cc})/(p_{dc} - p_{dd})$, using the notation in Table 1a, and the existence of cooperative equilibria is merely necessary, not sufficient for the emergence of cooperation. The resulting question, whether a predictive threshold for δ existed at all, has recently been answered in the affirmative by Blonski et al. (2011, BOS) and Dal Bo and Fréchette (2011, DF). They showed that cooperation occurs predictably if δ exceeds a threshold significantly stronger than existence of cooperative equilibria, namely if

$$\delta \geq \frac{p_{dc} + p_{dd} - p_{cd} - p_{cc}}{p_{dc} - p_{cd}} =: \delta^{\text{BOS}}. \tag{1}$$

While BOS derived this threshold axiomatically, their analysis gives little insight into related equilibrium strategies, since their axioms relate to the base game payoffs, not to existence of equilibria such as Grim.[4] Indeed, it is unclear whether δ^{BOS} has a strategic interpretation.

In order to obtain a first, preliminary picture of the strategies used in experiments, I estimated the average memory-1 strategies in the four experiments mentioned in the

[4]In short, the axioms are: Invariance with respect to linear payoff transformations; Monotonicity of cooperation in δ; Cooperation disappears as $p_{cd} \to -\infty$; Cooperation occurs iff a condition $\mu \geq 0$ where μ is additively separable with respect to $p_{cc} - p_{dd}, p_{dc} - p_{cc}, p_{dd} - p_{cd}$; the two differences $p_{dc} - p_{cc}$ and $p_{dd} - p_{cd}$ have the same weight in this condition.

Introduction. In addition to the 16 treatments of BOS and FD, where δ is below or near the BOS-threshold, I consider Duffy and Ochs (2009) and FRD, which contain one treatment each where δ is substantially above the BOS-threshold, to verify robustness. I focus on the second halves of these experiments, when learning has largely stabilized. Table 1 provides a detailed overview of these data sets and the estimated memory-1 strategies. The set of memory-1 histories is $\{cc, cd, dc, dd\}$ and the corresponding memory–1 Markov strategy is denoted as $\sigma = (\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$. For example, σ_{cd} denotes the probability that a player cooperates when his most recent action is c (cooperate) and his opponent's most recent action is d (defect).

Across treatments and experiments in Table 1, the average strategy has a surprisingly constant structure, approximately $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.9, 0.3, 0.3, 0.1)$. I refer to strategies of the form $\sigma_{cc} \approx 1$, $\sigma_{cd} = \sigma_{dc} \in (0, 1)$, and $\sigma_{dd} \approx 0$ as Semi-Grim: similar to Grim, they do not return to cooperation after mutual defection, but they are milder than Grim, as they return to cooperation after mixed histories. Note that the robustness of Semi-Grim strategies on average does not guarantee that any individual subject plays Semi-Grim, although it seems likely. This will be analyzed (and confirmed) in Section 6. The observation that $\sigma_{cd} = \sigma_{dc}$ is rejected in just one treatment, across all four experiments, is perhaps most surprising. It shows that subjects do not act directly reciprocal—a previously cooperating player does not cooperate less with a previously defecting one than the other way around.

The recurrence of Semi-Grim strategies does not seem to have been noticed in the literature. The only references related to this observation (that I found) are in Rapoport and Mowshowitz (1966), who observed $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (0.81, 0.43, 0.37, 0.22)$, which is discussed briefly also in Erev and Roth (2001), and in Bruttel and Kamecke (2012), who elicit strategies via three different methods (hot play, strategy method, and a Moore procedure following Engle-Warnick and Slonim, 2004, 2006) and find, but not discuss, strategies of the form $\sigma_{cc} > \sigma_{cd} \approx \sigma_{dc} > \sigma_{dd}$ in their logistic regressions (see their Table 4).

The conjecture that subjects play mixed Semi-Grim strategies would explain several results in the literature, in particular of DF and FRD. They both found that Tit-for-Tat (TFT, mimic the opponent's action of the previous round) had not been used systematically. This observation is compatible with mixed Semi-Grim, as $\sigma_{cd} < \sigma_{dc}$ would otherwise tend to result. Similarly, $\sigma_{dd} \approx 0$ explains why Win-Stay-Lose-Shift

Table 1: Overview of the experiments: Treatment parameters, numbers of observations, and average memory-1 strategies

Treatment	Game parameters b	a	δ	Sizes of the data sets #Subj	$\frac{\text{\#Games}}{\text{Subj}}$	$\frac{\text{\#Obs}}{\text{Subj}}$	#Obs	Average memory-1 strategy $\hat{\sigma}_{cc}$		$\hat{\sigma}_{dc}$		$\hat{\sigma}_{cd}$		$\hat{\sigma}_{dd}$
Blonski et al. (2011)														
6	1.429	1.286	0.5	20	8	17	340					0.133	$\approx$	0.006
3	2.5	1.5	0.5	20	9	21	420			0.333	$\approx$	0.117	$\approx$	0.014
2	1.25	1.125	0.75	40	8–9	41–42	1660							0.003
7	1.429	1.286	0.75	40	8–9	22–25	940	0.982	$\approx$	0.411	$\approx$	0.313	$>$	0.006
4	2.5	1.5	0.75	60	8–9	27–39	2020	0.912	$\gg$	0.128	$\approx$	0.226	$\gg$	0.032
8	1.429	1.286	0.875	40	5–6	33–34	1340	0.968	$>$	0.088	$\approx$	0.234	$\gg$	0.027
5	2.5	1.5	0.875	40	5–6	41–46	1740	0.971	$\gg$	0.28	$\approx$	0.211	$\gg$	0.039
9	2.4	1.8	0.75	40	8–9	28–38	1320	0.887	$\gg$	0.117	$\approx$	0.244	$>$	0.037
1	3	2	0.75	40	8–9	30–50	1600	0.908	$>$	0.286	$\approx$	0.285	$\gg$	0.022
10	4.667	3	0.75	40	8–9	34–40	1480	0.854	$\gg$	0.254	$\approx$	0.195	$>$	0.064
Dal Bo and Fréchette (2011)														
1	2.923	1.538	0.5	44	30–36	55–80	2988	1	$\approx$	0.464	$\approx$	0.406	$>$	0.03
3	2.923	2.154	0.5	50	36	62–81	3614	0.951	$\gg$	0.364	$\approx$	0.403	$\gg$	0.122
2	2.923	1.538	0.75	44	14–17	70–103	3606	0.946	$\gg$	0.38	$\approx$	0.358	$\gg$	0.03
5	2.923	2.769	0.5	46	34–39	67–78	3398	1	$\gg$	0.21	$\approx$	0.371	$\gg$	0.037
4	2.923	2.154	0.75	38	12–24	49–75	2524	0.959	$\gg$	0.555	$>$	0.328	$\gg$	0.103
6	2.923	2.769	0.75	44	15–18	59–83	3140	0.976	$\gg$	0.347	$\approx$	0.28	$\approx$	0.074
Duffy and Ochs (2009), "random rematching" treatment														
	3	2	0.9	56	4–7	28–103	3276	0.924	$\gg$	0.37	$\approx$	0.347	$\gg$	0.123
Fudenberg et al. (2012), "no-noise" treatment														
6	5	4	0.875	48	3–5	24–42	1800	0.935	$\gg$	0.465	$\approx$	0.465	$\gg$	0.078

Note: The "game parameters" are standardized as in Figure 1b. The strategy components $\hat{\sigma}_{cc}, \hat{\sigma}_{cd}, \hat{\sigma}_{dc}, \hat{\sigma}_{dd}$ are estimated as follows: For each state, the individual cooperation probability in this state is estimated for each subjects, and then averaged over all subjects with at least 3 observations per subject; the result is printed here if there are at least three such subjects. The relation signs indicate the p-values of Wilcoxon matched-pair tests (using the cooperation probabilities of individual subjects as independent observations): "$\gg, \ll$" indicate the p-value is $p < .001$, "$>, <$" indicate $p < .01$, and "$\approx$" indicates p-values greater than 0.01. As σ_{cd}, σ_{dc} mostly do not differ significantly, the relative frequencies of $\sigma_{cd,dc}$ have been pooled in tests against either σ_{cc} or σ_{dd}.

(WSLS, cooperate iff both players acted equally in the previous round) had generally been insignificant in their analyses.[5] Finally, FRD estimated weights that are fluctuating (across treatments) between strategies such as Tit-for-2-Tats (TF2T: Play C unless partner played D in both of the last 2 rounds), 2-Tits-for-Tat (2TFT: Play C unless partner played D in either of the last 2 rounds), Lenient Grim2 (Play C until 2 consecutive rounds occur in which either player played D, then play D forever), and corresponding memory-3 strategies. This is compatible with Semi-Grim, which randomizes after mixed histories fairly uniformly (see FRD in Table 1), and then exactly these paths of play will materialize depending on how the chips fall.

3 Remaining definitions

Consider the infinitely repeated Prisoner's dilemma (PD, Figure 1a) where players discount future payoffs using $\delta \in (0,1)$. In each round, the players may either cooperate (c) or defect (d), and in the theoretical analysis, I assume that their actions depend only on the choices made in the most previous round. The memory-1 strategies are denoted as $\sigma = (\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$, as introduced above. Following most of the literature, I focus on strategy profiles (σ, σ) that are symmetric between players, and consequently I do not use a player index in denoting strategies. Thus, if a player cooperates with probability $\sigma_{s',s''}$ in state (s', s''), then his opponent considers the state to be (s'', s') and cooperates with probability $\sigma_{s'',s'}$.

Given strategy profile (σ, σ) and a state state $(s', s'') \in S \times S$, define the expected payoff of choosing c in this state as $\pi_{s',s''}(c)$, the expected payoff of choosing d as $\pi_{s',s''}(d)$, and the expected payoff of playing according to $\sigma_{s',s''}$ as $\pi_{s',s''}$. These payoffs are determined by solving a linear equation system where for all $(s', s'') \in S \times S$

$$\pi_{s',s''} = \sigma_{s',s''} \cdot \pi_{s',s''}(c) + (1 - \sigma_{s',s''}) \cdot \pi_{s',s''}(d) \tag{2}$$

$$\pi_{s',s''}(c) = \sigma_{s'',s'} \cdot \big(\delta \pi_{cc} + (1-\delta)\, p_{cc}\big) + (1 - \sigma_{s'',s'}) \cdot \big(\delta \pi_{cd} + (1-\delta)\, p_{cd}\big) \tag{3}$$

$$\pi_{s',s''}(d) = \sigma_{s'',s'} \cdot \big(\delta \pi_{dc} + (1-\delta)\, p_{dc}\big) + (1 - \sigma_{s'',s'}) \cdot \big(\delta \pi_{dd} + (1-\delta)\, p_{dd}\big). \tag{4}$$

[5] Both TFT and WSLS are theoretically promising strategies. Axelrod (1980a,b) showed that TFT is effective against opponents with unknown strategies, and Nowak et al. (1993) and Imhof et al. (2007) showed that WSLS is even more effective than TFT when most subjects play cooperative strategies (as WSLS can restore cooperation after unilateral defection).

Solving the equation system for $(\pi_{s',s''})$ yields the expected payoffs as functions of σ. Relatedly, define the *cooperation incentive* $\tilde{\pi}_{s',s''}$ in state $(s',s'') \in S \times S$ to be the difference of expected payoffs from one-time cooperation and one-time defection in this state, with continuation play evolving according to σ. That is, $\tilde{\pi}_{s',s''} := \pi_{s',s''}(c) - \pi_{s',s''}(d)$. The player is strictly best off cooperating in state (s',s'') if $\tilde{\pi}_{s',s''} > 0$, he is best off defecting if $\tilde{\pi}_{s',s''} < 0$, and he may randomize if $\tilde{\pi}_{s',s''} = 0$. A strategy profile (σ,σ) is called Markov perfect equilibrium (MPE), or "equilibrium" for short, if

$$\sigma_{s',s''} > 0 \Rightarrow \tilde{\pi}_{s',s''} \geq 0 \qquad \text{and} \qquad \sigma_{s',s''} < 1 \Rightarrow \tilde{\pi}_{s',s''} \leq 0 \tag{5}$$

for all $(s',s'') \in S \times S$. In the following, I assume payoffs standardized as in Figure 1b.

4 Derivation of testable predictions

This section introduces six increasingly restrictive predictions derived from the assumption that subjects play robust, memory–1 Markov perfect equilibria. First, that is, I hypothesize that subjects play memory–1 Markov strategies. Existing evidence on this hypothesis is surprisingly scarce. Bruttel and Kamecke (2012) show that the memory–1 assumption is largely valid, with deviations being of limited relevance, but evidence on a broader range of data sets and treatment parameters appears desirable. In experimental analyses, the memory–1 MPE assumption is additionally understood to imply that the first-round cooperation probabilities equate with those in state (c,c). Denoting the initial state as $\emptyset$, this sub-hypothesis $\sigma_\emptyset = \sigma_{cc}$ is not formally implied by Markov perfection, however. Any state's cooperation probabilities induce an equilibrium in the first round $\emptyset$, and the validity $\sigma_\emptyset = \sigma_{cc}$ has not yet been verified. For these reasons, I am testing these two hypotheses first.

Hypothesis 1 (Markov)**.** *Subjects play memory-1 Markov strategies with* $\sigma_\emptyset = \sigma_{cc}$.

The remaining hypotheses are derived from the two refinement assumptions, robustness to utility perturbations and robustness to imperfect monitoring. I begin with robustness to utility perturbations. In conjunction with Markov perfection and the seemingly innocuous assumption $\sigma_{cc} > \sigma_{dd}$, this yields the strong hypothesis $\sigma_{cd} = \sigma_{dc}$. If exactly one player has cooperated in the previous round, both players will cooperate with the same probability in the next round—regardless of who actually cooperated. The previously cooperating player does not explicitly punish his opponent's defection in the previous round, nor does the previously defecting player explicitly reciprocate his opponent's cooperation. I therefore call these strategies non-reciprocal.

Robustness to utility perturbations is a generic property of Nash equilibria in normal-form games and of regular MPEs in stochastic games (Doraszelski and Escobar, 2010). Experimental analyses generally focus on the case that the utility perturbations have extreme-value distribution, i.e. on logit equilibria, and following McKelvey and Palfrey (1995), this has been shown to explain experimental observations in a wide range of circumstances (see Footnote 2). While the limiting equilibrium, as noise disappears, is generically independent of the noise distribution (Doraszelski and Escobar, 2010), I follow this tradition and focus on limiting-logit equilibria. Formally, a strategy profile (σ,σ) is a Markov logit equilibrium (MLE, Breitmoser et al., 2010) if there exists $\lambda \in \mathbb{R}_+$ such that for all states (s', s''),

$$\sigma_{s',s''} = \frac{\exp\{\lambda \cdot \pi_{s',s''}(c)\}}{\exp\{\lambda \cdot \pi_{s',s''}(c)\} + \exp\{\lambda \cdot \pi_{s',s''}(d)\}}. \tag{6}$$

It is a limiting-logit MPE if it is the limit of logit MPEs as λ tends to infinity.

Proposition 4.1. *Let (σ,σ) be an MLE of a repeated PD. If $\sigma_{cc} > \sigma_{dd}$ and $\sigma_{cc} > 0.5$, then $\sigma_{cd} = \sigma_{dc} < \sigma_{cc}$.*

All proofs are relegated to the appendix. Thus, $\sigma_{cd} = \sigma_{dc} < \sigma_{cc}$ results in all Markov logit equilibria (if $\sigma_{cc} > \sigma_{dd}$), and by extension in all limiting-logit MPEs. Thus, it applies in all MPEs that are robust to (extreme value) utility perturbations.[6] In conjunction, the assumptions and the implication yield the second hypothesis.

Hypothesis 2 (No reciprocity)**.** *Subjects play limiting-logit equilibria with $\sigma_{cc} > \sigma_{dd}$. This implies $\sigma_{cd} = \sigma_{dc}$ and $\sigma_{cc} > \sigma_{cd,dc}$.*

In addition, players may seek robustness to imperfect monitoring. Monitoring is imperfect if one may forget the current state with positive probability, possibly due to imperfect attention in the experiment or due to imperfect recall in general. Regular

[6]There is an alternative set of limiting-logit MPEs with $\sigma_{cc} = \sigma_{dd}$. In these equilibria, $\sigma_{cd} < \sigma_{cc} = \sigma_{dd} < \sigma_{dc}$ results. These equilibria are not of empirical relevance, as $\sigma_{cc} > \sigma_{dd}$ will be shown to be satisfied universally, i.e. in all treatments of all experiments.

MPEs, including all well-known ones such as Grim and Win-Stay-Lose-Shift, are not robust to imperfect monitoring (Kandori, 2002), but an alternative class of completely mixed MPEs—where player randomize in all states—are robust (Ely and Välimäki, 2002). In completely mixed MPEs, players are always indifferent, and their best responses are independent of their beliefs about the opponent's history. This yields the robustness to imperfect monitoring of these *belief-free MPEs* (Ely et al., 2005). To illustrate briefly, the cooperation incentives[7] generally satisfy

$$\tilde{\pi}_{cc} - \tilde{\pi}_{cd} = (\sigma_{dc} - \sigma_{cc}) \cdot \mu \qquad \tilde{\pi}_{cc} - \tilde{\pi}_{dc} = (\sigma_{cd} - \sigma_{cc}) \cdot \mu \tag{7}$$

$$\tilde{\pi}_{cc} - \tilde{\pi}_{dd} = (\sigma_{dd} - \sigma_{cc}) \cdot \mu \qquad \tilde{\pi}_{cd} - \tilde{\pi}_{dc} = (\sigma_{cd} - \sigma_{dc}) \cdot \mu \tag{8}$$

with some common factor $\mu \in \mathbb{R}$. Thus, $\mu = 0$ implies $\tilde{\pi}_{cc} = \tilde{\pi}_{cd} = \tilde{\pi}_{dc} = \tilde{\pi}_{dd}$. If $\tilde{\pi}_{cc} = 0$ holds in addition, the underlying strategy profile (σ, σ) is a completely mixed MPE, i.e. an MPE where players are indifferent in all states. These are the equilibria that are robust to imperfect monitoring. Solving for $\mu = 0$ and $\pi_{cc} = 0$, we obtain a two-dimensional manifold of such MPEs (Bhaskar et al., 2008),

$$\sigma_{cd} = \frac{(b-1)\,\delta\sigma_{cc} - (b-a)\,(1+\delta\sigma_{dd})}{(a-1)\,\delta}, \qquad \sigma_{dc} = \frac{a\delta\sigma_{dd} - \delta\sigma_{cc} + 1}{(a-1)\,\delta}. \tag{9}$$

Since players are indifferent in all states, their own payoffs are independent of their own actions, and in this sense, each player sets the expected payoff for his opponent. This property has been independently discovered and further analyzed by Press and Dyson (2012, Eq. (8)) in their analysis of "zero-determinant" strategies. Their work shows that belief-free MPEs are not only robust to imperfect monitoring (or recall), but in some sense also to renegotiation. By playing belief-free MPE strategies, players essentially use "ultimatum strategies" (Stewart and Plotkin, 2012) by which they set their opponent's payoff and thus unilaterally enforce a claim of the payoffs.

As indicated before, I am interested in belief-free MPEs that are additionally robust to utility perturbations. In a first analytical step, I characterize the set of belief-free equilibria that satisfy the "no-reciprocity" condition $\sigma_{cd} = \sigma_{dc}$ implied by limiting-logit equilibrium. This additional condition refines the two-dimensional manifold to a one-dimensional one. Notably, such MPEs exists if $\delta \geq \delta^{\star}$ and they induce "Semi-

[7] Recall from Section 3 that the cooperation incentive in state (s', s'') is the long-term difference in expected payoffs between one-time cooperation and one-time defection, i.e. $\tilde{\pi}_{s',s''} = \pi_{s',s''}(c) - \pi_{s',s''}(d)$.

Grim” strategies, i.e. $\sigma_{cc} > \sigma_{cd} = \sigma_{dc} > \sigma_{dd}$.

Proposition 4.2. *A belief-free MPE satisfying $\sigma_{cd} = \sigma_{dc}$ exists iff $\delta \geq \delta^\star$, with*

$$\delta^\star = \frac{b-a+1}{b} \equiv \frac{p_{dc} + p_{dd} - p_{cd} - p_{cc}}{p_{dc} - p_{cd}} =: \delta^{BOS}. \tag{10}$$

These MPEs are in Semi-Grim strategies ($\sigma_{cc} > \sigma_{cd} = \sigma_{dc} > \sigma_{dd}$) and constitute a one-dimensional manifold consisting of all strategy profiles (σ, σ) satisfying

$$\sigma_{cc} = \frac{(b\delta-b+a-1)r+b-a+1}{bd} \qquad \sigma_{dc} = \frac{(b\delta-b+a-1)r+1}{bd} \qquad \sigma_{dd} = \frac{(b\delta-b+a-1)r}{bd} \tag{11}$$

for some $r \in [0,1]$.

That is, a repeated PD satisfies the five BOS axioms (Blonski et al., 2011) if and only if a belief-free MPE satisfying $\sigma_{cd} = \sigma_{dc}$ exists. This is a first characterization of the BOS-threshold in terms of equilibrium strategies. This class of MPEs may therefore explain both the Semi-Grim pattern across treatments (Section 2) and the predictiveness of the BOS-threshold observed by BOS and DF.

Next, I show that one these MPEs is indeed robust to (extreme-value) utility perturbations. This is not obvious, as Bhaskar et al. (2008) showed that belief-free memory–1 MPEs are not robust to generic utility perturbations. One of them is robust to extreme-value perturbations, however, and this standard assumption (following Luce, 1959) is technically not even particularly special. Belief-free MPEs that are robust to alternative distributions coexist under the same conditions, and Bhaskar et al. (2008) also showed that all belief-free memory–1 MPEs are robust to generic perturbations if we allow them to be limits of infinite-horizon strategies.

Proposition 4.3. *In any repeated PD, a memory-1 MPE that is belief-free and limiting-logit exists iff $\delta \geq \delta_{BOS}$. It satisfies $\sigma_{cc} > \sigma_{cd} = \sigma_{dc} > \sigma_{dd}$, and it satisfies Eq.* (11) *using r such that*

$$\left(\frac{((-b\delta+b-a+1)r+b\delta)((b\delta-b+a-1)r+1)}{(b\delta-b+a-1)r((-b\delta+b-a+1)r+b\delta-1)}\right)^{b-a} = \frac{((b\delta-b+a-1)r+b-a+1)((b\delta-b+a-1)r-b\delta+1)}{(b\delta-b+a-1)(r-1)((b\delta-b+a-1)r+1)}.$$

The BOS axioms are therefore equivalent to requiring existence of MPEs that are robust to both imperfect monitoring and (extreme-value) utility perturbations. Two hypotheses that immediately follow from this result are that subjects cooperate in round

1 and play Semi-Grim strategies in the continuation if this MPE exists.

Hypothesis 3 (Semi-Grim strategy)**.** *Subjects play Semi-Grim strategies* $\sigma_{cc} > \sigma_{cd,dc} > \sigma_{dd}$*, both on average and individually, if belief-free, limiting-logit equilibria exist.*

Hypothesis 4 (Round-1 cooperation)**.** *Cooperation in round 1 is increasingly frequent as* δ *increases, and it occurs "systematically" (in at least 50% of the games) if belief-free, limiting-logit equilibria exist.*

The final two hypotheses extend the previous ones in that they relate the observed strategies more tightly to the Semi-Grim MPEs. At first glance, it may seem natural to hypothesize that subjects exactly play the unique belief-free, limiting-logit MPE for all $\delta > \delta^\star$. This would imply that the probability cooperation in state (c,c) is decreasing in δ. For example, in case $b = a+1$, the belief-free, limiting-logit MPE is (applying Prop. 4.3, this MPE satisfies Eq. (11) for $r = 1/2$)

$$\sigma_{cc} = \tfrac{b\delta/2+1}{b\delta} \qquad \sigma_{cd} = \sigma_{dc} = \frac{1}{2} \qquad \sigma_{dd} = \tfrac{b\delta/2-1}{b\delta}. \tag{12}$$

Thus, at $\delta = \delta^\star$, $\sigma_{cc} = 1$ and $\sigma_{cc} = 0$, but as δ increases, σ_{cc} is predicted to decrease. The intuition is simple. Players have to be indifferent between cooperation and defection while the long-term gains from cooperation increase in δ. Thus, one continues to be indifferent between c and d only if the cooperation rate of the opponent is decreasing in δ, i.e. if σ_{cc} decreases. This would contradict fairly robust results in the existing literature, e.g. Roth and Murnighan (1978) and Dal Bo (2005), and therefore I hypothesize the following.

Hypothesis 5 (Comparative statics)**.** *The average strategy* $(\sigma_{cc}, \sigma_{cd,dc}, \sigma_{dd})$ *does not correlate (positively) with the belief-free, limiting-logit MPE for all* δ*.*

More generally, as δ increases, players may seek to balance three objectives: Efficiency $(\sigma_{cc} = 1)$, robustness to imperfect monitoring (RIM), and robustness to utility perturbations (RUP). These objectives are simultaneously satisfied only in the limiting case $\delta = \delta^\star$. As δ increases, the belief-free, limiting-logit MPE splits up into three distinctive Semi-Grim equilibria, one of which is efficient and RIM, another one is efficient and RUP, and the third one is the belief-free, limiting-logit MPE characterized in Prop. 4.3 (which is RIM and RUP). If $\delta > \delta^\star$, players therefore have to pick two of the three criteria, and I hypothesize that they always play one of the these three Semi-Grim equilibria.

Hypothesis 6 (Semi-Grim MPE)**.** *Individual subjects play one of the three Semi-Grim MPEs if they exist.*

To be specific, the MPE that is efficient ($\sigma_{cc} = 1$) and RIM obtains for $r = 1$ in Eq. (11), which yields

$$\sigma_{cc} = 1 \qquad \sigma_{dc} = \tfrac{b\delta - b + a}{bd} \qquad \sigma_{dd} = \tfrac{b\delta - b + a - 1}{bd}. \tag{13}$$

I refer to this equilibrium as the *efficient belief-free Semi-Grim* MPE. In turn, the MPE that is efficient and RUP will be called *efficient limiting-logit Semi-Grim* MPE. It structure is derived in the appendix (Lemma B.1) and satisfies $\sigma_{cc} = 1$, $\sigma_{dd} = 0$, and

$$\sigma_{cd} = \sigma_{dc} = \frac{\sqrt{2a\left(2d^2-4d+1\right)+b^2(d-1)^4-2ab(d-1)^2-2b(d-1)^2+a^2+1}+bd^2+(2-2b)d+b-a-1}{2bd^2+(-2b-2a+2)d}. \tag{14}$$

It exists under slightly weaker conditions than the belief-free, limiting-logit MPE, but in all four experiments reviewed above, there is just one treatment where this limiting-logit Semi-Grim MPE exists and the belief-free, limiting-logit MPE does not. I will therefore not analyze this subtle difference.

5 Analysis of aggregate behavior

In the next two sections, I am testing the hypotheses using the data sets reviewed in Table 1. As above, I focus on the second halves of the experiments, when learning has largely stabilized (similarly to DF and FRD). In conjunction with the measures described in the following, I believe this yields robust, replicable results. To reduce the probability of false positives, I require constant and comparably high levels of significance in all hypotheses tests. A test result will be called *significant* if the associated p-value is less than 0.01, and *highly significant* if $p < 10^{-4}$ (given the sizes of the data sets, these thresholds will be met frequently). Considering the noise in experimental data, it seems unrealistic to expect that any of the above ex-ante hypotheses is confirmed or rejected in every treatment of all experiments. I assume that correct hypotheses hold up in the (vast) majority of treatments across experiments, while incorrect ones are rejected in the majority of treatments. I will say that observations are *systematic* if they are significant in at least half the treatments, they are *highly systematic* if they are highly significant in at least half the treatments and significant in at least 75%, and they are *universal* if they are highly significant in at least 75% and significant in all treatments. The distinction between systematic and non-systematic observations is robust to changing the threshold. Indeed, there are no observations that are systematic without being highly systematic, and the threshold for "systematic observations" could be lowered to requiring significance in a third of the treatments without having to change any of the wording—there would still be no systematic observations that are not highly systematic.

The first hypothesis is that subjects play memory-1 Markov strategies. This one is crucial, as all subsequent hypothesis are derived from it. I test it in models regressing the probability of cooperation (either 0 or 1 in each round) on all memory-1 histories and as many memory-2 histories as the non-singularity condition permits, controlling for individual differences by subject-level random effects. The *p*-values of the hypothesis tests are bootstrapped. For transparency about the estimated cooperation probabilities (i.e., strategies), in particular in tests of the subsequent hypotheses, I use linear-probability models without intercept throughout this section. The results do not change notably if logit models are used instead.

Table 2 categorizes the resulting levels of significance across treatments; Table 9 (in the appendix) additionally provides the estimated coefficients. Treatment 6 of BOS contains virtually no cooperation, which implies that meaningful regression analysis is impossible there. In the remaining 17 treatments, each memory-1 history is highly significant in at least 10 treatments, and significant in at least 15, from which I conclude that every memory-1 history is of highly systematic relevance. In turn, no memory-2 history is significant in more than 5 of these 17 treatments, and thus no memory-2 history is of systematic relevance. The fact that one or two memory-2 histories are significant in most cases, varying between treatments, confirms the chosen approach of analyzing multiple data sets to mitigate the multiple-testing problem and obtain replicable results. I conclude that only memory-1 histories are systematically relevant and focus on memory-1 Markov strategies in the following.

Table 2: Testing the memory-1 Markov assumption: Levels of significance of memory-1 and memory-2 Markov strategies (bootstrapped p-values; the coefficients are provided in the appendix, Table 9)

	Memory-1 histories				Memory-2 histories $(t-1)\times(t-2)$										
Tr	(c,c)	(d,c)	(c,d)	(d,d)	$cc\times cc$	$cc\times cd$	$cc\times dc$	$cd\times cc$	$cd\times dc$	$cd\times dd$	$dc\times cc$	$dc\times cd$	$dc\times dd$	$dd\times cd$	$dd\times dc$
Blonski et al. (2011)															
6															
3	$++$														
2	$++$			$+$											
7	$++$	$++$	$+$						$+$						
4	$++$	$++$	$+$	$++$										$--$	
8	$++$	$++$	$++$	$++$										$-$	
5	$++$	$++$	$++$	$++$					$+$						$+$
9	$++$	$++$	$+$	$+$					$+$						
1	$++$	$++$	$++$	$++$						$-$		$++$			
10	$++$	$++$	$+$	$++$				$-$	$+$						
Dal Bo and Fréchette (2011)															
1	$++$	$++$	$+$	$++$								$+$		$-$	
3	$++$	$++$	$++$	$++$											
2	$++$	$++$	$++$	$++$							$+$	$++$			
5	$++$	$++$	$++$	$++$								$+$			
4	$++$	$++$	$++$	$++$		$-$								$--$	
6	$++$	$++$	$++$	$++$											
Duffy and Ochs (2009), "random rematching" treatment															
	$++$	$++$	$++$	$++$		$-$			$+$						$++$
Fudenberg et al. (2012), "no-noise" treatment															
6	$++$	$++$	$++$	$++$											

Note: $^{++}$ indicates significance of a positive coefficient at $p<10^{-4}$, $^{+}$ indicates significance of a positive coefficient at $p<10^{-2}$, $^{--}$ and $^{-}$ indicate the respective levels of significance for negative coefficients

The second claim in Hypothesis 1 is that subjects approach the first round as if they had cooperated before, i.e. $\sigma_{cc} = \sigma_{\emptyset}$. It can be tested similarly, in linear-probability models with subject-level random effects, now excluding the memory-2 histories due to their predicted and observed unsystematic relevance. In addition, to the four non-empty histories (cc, cd, dc, dd), this regression includes the empty history $\emptyset$. As before, I include subject-level random effects and bootstrap the p-values. The results of evaluating $H_0 : \sigma_{cc} = \sigma_{\emptyset}$ are provided in the third column of Table 3. The remaining columns indicate the results on tests of subsequent hypotheses. As for $H_0 : \sigma_{cc} = \sigma_{\emptyset}$, Table 3 shows that it is rejected highly significantly in all treatments, i.e. universally.

Result 1 (Markov)**.** *Subjects play memory-1 Markov strategies with $\sigma_{\emptyset} < \sigma_{cc}$.*

This shows that we cannot pool first-round behavior and choices after (c, c) when we estimate strategies, and I will therefore continue to distinguish the state (c, c) from round 1, i.e. from state $\emptyset$.

Hypothesis 2 claims that $\sigma_{cc} > \sigma_{dd}$ would imply the no-reciprocity condition $\sigma_{cd} = \sigma_{dc}$ as well as $\sigma_{cc} > \sigma_{cd,dc}$ (derived from the assumption that subjects play limiting-logit equilibria). This is analyzed with the set of regression models already used to test $\sigma_{cc} = \sigma_{\emptyset}$, and the results are provided in the remaining columns of Table 3. The first column shows that $\sigma_{cc} > \sigma_{dd}$ applies universally, i.e. $\sigma_{cc} = \sigma_{dd}$ is rejected highly significantly in all treatments. The theoretical implication $\sigma_{cd} = \sigma_{dc}$ is maintained in 13 of 17 treatments, again systematically. Surprisingly, the four rejections of $\sigma_{cd} = \sigma_{dc}$ are all in Dal Bo and Fréchette (2011), which suggests the existence of a cohort effect. As there is no systematic evidence in favor of rejecting $\sigma_{cd} = \sigma_{dc}$, however, I will pool the states (c, d) and (d, c) in the following and use an otherwise equivalent regression model to test $\sigma_{cc} > \sigma_{cd,dc}$. The null $\sigma_{cc} = \sigma_{cd,dc}$ is rejected universally in favor of $\sigma_{cc} > \sigma_{cd,dc}$, as the last column of Table 3 shows. Thus, both predictions of limiting-logit MPE are confirmed.

Result 2 (No reciprocity)**.** *The hypothesis that subjects play limiting-logit equilibria is sustained. $\sigma_{cc} > \sigma_{dd}$ and $\sigma_{cc} > \sigma_{cd,dc}$ are universal, while $\sigma_{cd} = \sigma_{dc}$ is not violated systematically.*

I proceed by testing the stronger hypotheses that subjects' behavior aligns with the limiting-logit, belief-free equilibrium. Hypothesis 3 claims that the average subject plays a Semi-Grim strategy, i.e. that the Semi-Grim pattern discussed above (Table 1)

Table 3: Testing the predictions of limiting-Logit, belief-free equilibria

	Treatment parameters			$H_0 : \sigma_{cc} = \sigma_{dd}$			$H_0 : \sigma_{dc} = \sigma_{cd}$			$H_0 : \sigma_{cc} = \sigma_{\emptyset}$			Pooled estimate (using $\sigma_{cd,dc}$)				
Tr	b	a	δ	σ_{cc}		σ_{dd}	σ_{dc}		σ_{cd}	σ_{cc}		$\sigma_{\emptyset}$	σ_{cc}		$\sigma_{cd,dc}$		σ_{dd}
Blonski et al. (2011)																	
6	1.43	1.29	0.5														
3	2.5	1.5	0.5	0.923	$\gg$	0.011	0.186	$\approx$	0.09	0.923	$\gg$	0.217	0.93	$\gg$	0.142	$>$	0.01
2	1.25	1.12	0.75	0.97	$\gg$	0.004	0.33	$\approx$	0.218	0.97	$\gg$	0.044	0.976	$\gg$	0.276	$\gg$	0.004
7	1.43	1.29	0.75	0.994	$\gg$	0.002	0.373	$\approx$	0.289	0.994	$\gg$	0.15	1.001	$\gg$	0.333	$\gg$	0.001
4	2.5	1.5	0.75	0.886	$\gg$	0.038	0.189	$\approx$	0.154	0.886	$\gg$	0.243	0.888	$\gg$	0.172	$\gg$	0.037
8	1.43	1.29	0.88	0.914	$\gg$	0.032	0.237	$\approx$	0.207	0.914	$\gg$	0.372	0.916	$\gg$	0.222	$\gg$	0.031
5	2.5	1.5	0.88	0.915	$\gg$	0.053	0.318	$\approx$	0.236	0.915	$\gg$	0.393	0.922	$\gg$	0.279	$\gg$	0.051
9	2.4	1.8	0.75	0.913	$\gg$	0.03	0.206	$\approx$	0.172	0.913	$\gg$	0.553	0.915	$\gg$	0.189	$\gg$	0.029
1	3	2	0.75	0.801	$\gg$	0.051	0.318	$\approx$	0.266	0.801	$\gg$	0.35	0.807	$\gg$	0.295	$\gg$	0.049
10	4.67	3	0.75	0.872	$\gg$	0.053	0.28	$\approx$	0.163	0.872	$\gg$	0.595	0.874	$\gg$	0.223	$\gg$	0.051
Dal Bo and Fréchette (2011)																	
1	2.92	1.54	0.5	0.795	$\gg$	0.031	0.428	$>$	0.207	0.795	$\gg$	0.062	0.816	$\gg$	0.327	$\gg$	0.03
3	2.92	2.15	0.5	0.811	$\gg$	0.118	0.383	$\gg$	0.255	0.811	$\gg$	0.194	0.822	$\gg$	0.323	$\gg$	0.116
2	2.92	1.54	0.75	0.899	$\gg$	0.041	0.414	$>$	0.3	0.899	$\gg$	0.263	0.906	$\gg$	0.36	$\gg$	0.039
5	2.92	2.77	0.5	0.929	$\gg$	0.072	0.244	$\approx$	0.319	0.929	$\gg$	0.408	0.917	$\gg$	0.28	$\gg$	0.077
4	2.92	2.15	0.75	0.945	$\gg$	0.167	0.548	$\gg$	0.364	0.945	$\gg$	0.768	0.948	$\gg$	0.454	$\gg$	0.157
6	2.92	2.77	0.75	0.98	$\gg$	0.106	0.383	$\approx$	0.303	0.98	$>$	0.955	0.981	$\gg$	0.342	$\gg$	0.104
Duffy and Ochs (2009), "random rematching" treatment																	
	3	2	0.9	0.957	$\gg$	0.134	0.38	$\approx$	0.342	0.957	$\gg$	0.696	0.958	$\gg$	0.361	$\gg$	0.133
Fudenberg et al. (2012), "no-noise" treatment																	
6	5	4	0.88	0.949	$\gg$	0.171	0.483	$\approx$	0.466	0.949	$\gg$	0.825	0.95	$\gg$	0.474	$\gg$	0.169

Note: p-values are bootstrapped, ** indicates significance at $p < 10^{-4}$, * indicates significance at $p < 10^{-2}$

is econometrically robust to individual random effects. I test this in the same set of regressions used to test $\sigma_{cc} > \sigma_{cd,dc}$, again bootstrapping the *p*-values. The results, shown in the last column of Table 3, confirm the Semi-Grim hypothesis universally. The population plays $\sigma_{cc} > \sigma_{cd,dc} > \sigma_{dd}$ in all treatments of all experiments.

Result 3 (Semi-Grim)**.** *The average cooperation probabilities satisfy* $\sigma_{cc} > \sigma_{cd,dc} > \sigma_{dd}$ *universally.*

Note that subjects play Semi-Grim strategies even when the respective equilibrium, the belief-free, limiting-logit MPE, does not exist. This "puzzle" is further analyzed (and resolved) in the next section, when the individual strategies are considered. At this point, let me just clarify that Result 3 neither claims nor shows that every subject plays a Semi-Grim strategy (which is the second half of Hypothesis 3).

The related Hypothesis 4, that cooperation in round 1 correlates with the existence of limiting-logit, belief-free equilibria, can be confirmed as well. Figure 2 provides a first intuition. It plots the relative frequencies of observing at least one cooperative action (c) in round 1 against $\delta - \delta^\star$, for all treatments. The respective logistic regression controls for the weights of the various treatments and for correlation within experiments by random effects. If Hypothesis 4 is correct, the logistic curve crosses the 50%-line approximately at $\delta - \delta^\star = 0$ and the regression intercept is insignificant.[8] This is indeed the case, the curve crosses the line almost exactly at $\delta = \delta^\star$. The regression coefficients presented alongside Figure 2 show that the intercept is insignificant. That is, subjects respond to increasing δ, as the coefficient of $\delta - \delta^\star$ is significant, and they starting cooperating "systematically" (50%) once the Semi-Grim strategies turn into Markov-perfect equilibria, as the intercept is insignificant.[9]

Result 4 (Round-1 cooperation)**.** *Cooperation in round 1 is increasingly frequent as* δ *increases and can be observed systematically (in at least 50% of the games) if* $\delta \geq \delta^\star$.

[8] In a logistic regression $\Pr(\text{Coop}) = 1/\big(1+\exp\{-a-b\cdot(\delta-\delta^\star)\}\big)$, the intercept is a.

[9] Blonski et al. (2011) obtain a similar result on the hypothesis that cooperation occurs at least once in any round, while I focus on cooperation in round 1. BOS further refine their result by excluding paths of play that fail a rationalizability condition. Their "necessary" condition ("Filtering Rule 1") for rationalizability is that it must be possible to extend the observed path of play by a continuation equilibrium such that both players obtain at least their minimax payoffs. This condition is necessary for individual rationality if one assumes that subjects play pure strategies. With mixed strategies, such as Semi-Grim, there is a positive probability (which is high if the game is short) such that a path of play results that is not ex-post rationalizable in this way although it resulted from an equilibrium strategy. Since I explicitly allow for mixed strategies, I therefore abstain from such refinements.

Figure 2: Cooperation in round 1 in relation to $\delta - \delta^{\star}$

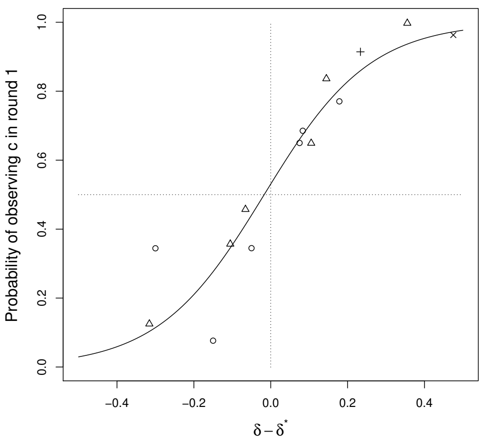

Regression coefficients

Intercept a	Coefficient of $\delta - \delta^{\star}$
0.123	$7.237^{\star\star}$
(0.124)	(0.838)

$\Pr(\text{Coop}) = 1/\big(1 + \exp\{-a - b \cdot (\delta - \delta^{\star})\}\big)$

In conjunction with Table 3, which shows that the average memory-1 strategy is Semi-Grim and fairly invariant across treatments, this suggests that subjects play Semi-Grim continuation strategies intuitively in all cases, but they trust these strategies enough to also cooperate in round 1 only if they are equilibrium strategies.

Finally, I evaluate to which degree the observed continuation strategies comply with the comparative statics of belief-free, limiting-logit equilibrium (Hypothesis 5). As discussed in relation to Eq. (12), σ_{cc} is decreasing in δ in the limiting-logit, belief-free equilibrium. A confirmation of this prediction would therefore contradict fairly robust previous experimental results (see Dal Bo, 2005). We may therefore expect the comparative statics to fail. This is confirmed by Figures 3 and 4, which plot the observed memory-1 strategies against $\delta - \delta^{\star}$ and against the predicted strategies (respectively). Table 4 provides the respective regression coefficients. Pooling states (c,d) and (d,c), we are left with three states. The cooperation probabilities in all three states are increasing in $\delta - \delta^{\star}$ (Table 4a), i.e. cooperation becomes increasingly robust as δ increases in relation to the threshold. This is intuitive and compatible with previous experiments, but it contradicts the comparative statics of limiting-logit, belief-free MPE. Indeed, Table 4b shows that the aggregate strategies are negatively correlated with the predictions in states (c,c) and (cd,dc).

Result 5 (Comparative statics)**.** *The cooperation probabilities in all states are increasing in $\delta - \delta^{\star}$. The cooperation probabilities in states (c,c) and (cd,dc) correlate*

Figure 3: The observed memory-1 strategies in relation to $\delta - \delta^{\star}$

(a) Cooperation after *cc* (b) Cooperation after cd, dc (c) Cooperation after *dd*

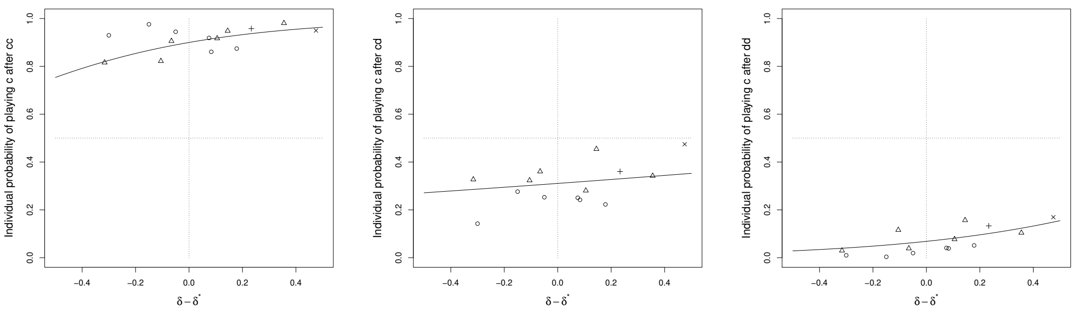

Figure 4: The observed memory-1 strategies in relation to the predictions

(a) Cooperation after *cc* (b) Cooperation after cd, dc (c) Cooperation after *dd*

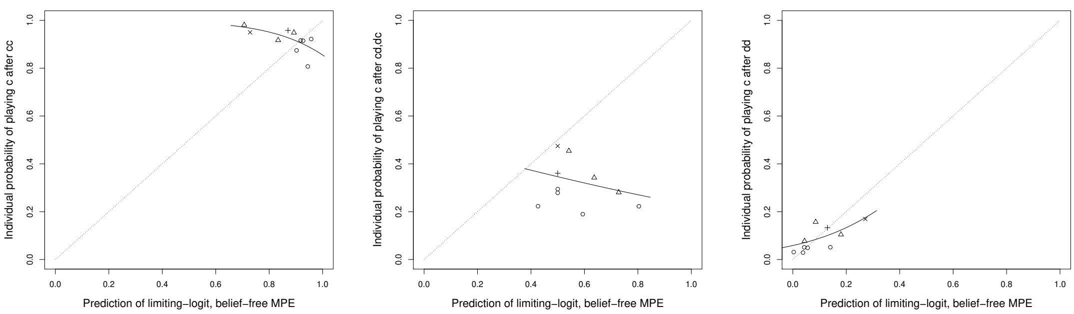

negatively with the predictions, while those in state (d,d) correlate positively. Overall, the comparative statics of limiting-logit, belief-free MPE are violated.

This suggests that subjects do not play limiting-logit, belief-free MPEs for all δ. They start cooperating when these equilibria start to exist, but they seem to maintain efficiency ($\sigma_{cc} = 1$) as δ increases. This is analyzed next, by inspecting individual strategies.

6 Analysis of individual strategies

The previous section shows that subjects play Semi-Grim strategies on average and that they start cooperating when MPEs inducing Semi-Grim strategies start to exist.

Table 4: Regression of observed strategies on $\delta - \delta^{\star}$ and the predicted MPE

(a) Logistic regression on $\delta - \delta^{\star}$

Cooperation	Intercept	Coefficient
after cc	$2.196^{\star\star}$ (0.035)	$2.156^{\star\star}$ (0.177)
after cd, dc	$-0.797^{\star\star}$ (0.022)	$0.38^{\star}$ (0.111)
after dd	$-2.611^{\star\star}$ (0.041)	$1.829^{\star\star}$ (0.198)

(b) Logistic regression on prediction

Cooperation	Intercept	Coefficient
after cc	$7.728^{\star\star}$ (0.488)	$-5.954^{\star\star}$ (0.551)
after cd, dc	-0.045 (0.115)	$-1.181^{\star\star}$ (0.197)
after dd	$-2.779^{\star\star}$ (0.062)	$4.529^{\star\star}$ (0.418)

These MPEs are limiting-logit and belief-free, i.e. robust to utility perturbations and imperfect monitoring—two ex-ante plausible selection criteria. The comparative statics of limiting-logit, belief-free MPE are not satisfied, however. Efficient cooperation ($\sigma_{cc} \approx 1$) results in all treatments, and if anything, σ_{cc} is increasing in δ, as opposed to being decreasing as limiting-logit, belief-free MPE predicts. A potential explanation is that subjects start cooperating if limiting-logit, belief-free MPE exist, at $\delta = \delta^{\star}$, and at this threshold, the equilibrium is even efficient (inducing $\sigma_{cc} = 1$). As δ increases, however, subjects have to pick two of these three attributes—MPEs that are efficient, limiting-logit *and* belief-free do not exist for high δ. Depending on which two of these three attributes they pick, they would end up with one of three Semi-Grim MPEs, and the hypotheses (3 and 6) to be analyzed in this section are that they do pick one of these MPEs.

I proceed in two steps, starting with a general classification of the individual strategies, addressing Hypothesis 3, and secondly estimating the weights of eight relevant MPEs, addressing Hypothesis 6. To analyze individual behavior, the experimental literature following Stahl and Wilson (1994, 1995), including Dal Bo and Fréchette (2011) and Fudenberg et al. (2012), usually estimates latent strategy weights by finite-mixture modeling (McLachlan and Peel, 2000). That is, assuming the population consists of a finite number of components (each component is associated with a particular strategy), the weights of these components are treated as latent and estimated by maximum likelihood. I adopt this approach as well, all details are provided in Appendix A, and following Biernacki et al. (2000), I identify insignificant components using the *ICL-BIC* information criterion. The weights of the remaining significant components are of interest in my analysis. As before, all standard errors are bootstrapped.

First, to classify individual strategies without imposing an equilibrium constraint, I estimate the weights of six prototypical strategies—only one of which is Semi-Grim. The other candidate strategies are (1) the constant strategy $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (\alpha_1, \alpha_1, \alpha_1, \alpha_1)$, $\alpha_1 \in [0,1]$, which contains Always-Defect and Always-Cooperate as special cases; (2) a generalized Grim strategy of the form $(1, \alpha_2, \alpha_2, \alpha_2)$, which contains Grim and Always cooperate as special cases; (3) a generalized TFT strategy $(1, 0, \alpha_3, 0)$; (4) a generalized Win-Stay-Lose-Shift strategy $(1, 0, 0, \alpha_4)$; (5) a generalized cooperative strategy $(1, \alpha_5, 0, 0)$, which is behaviorally equivalent to "Always Cooperate" for $\alpha_5 = 1$ and additionally allows to reconstruct a Semi-Grim population (in conjunction with the strategies 3 and 4) without any individual playing a Semi-Grim strategy; and (6) Semi-Grim $(1, \alpha_6, \alpha_6, 0)$ itself. In order to facilitate the identification of the strategies, I restrict the cooperation probabilities $\alpha_3, \alpha_4, \alpha_5$ of TFT, WSLS, and Coop (resp.) to be in $[0.5, 1]$. That is, a subject's behavior has to be sufficiently pronounced and distinctive from Grim to be classified as either TFT, WSLS, or Coop. Subject to this constraint, all strategy parameters $(\alpha_1, \ldots, \alpha_6)$ are treated as free parameters and estimated from the data (as described in the appendix).

The results, i.e. estimated parameters and weights, are provided in Table 5. First, let me summarize the results of Table 5 by simply counting which components have majority weight in the various treatments.

Result 6 (Semi-Grim strategies)**.** *In* $15/18$ *treatments, the majority of subjects plays Semi-Grim strategies, and in* $12/18$ *treatments, at least 80% of the subjects play Semi-Grim. It is the only strategy that is played systematically (in more than 50% of the treatments). The only other strategy that has more than 10% weight in more than* $3/18$ *treatments is Generalized Grim* $(1, \alpha, \alpha, \alpha)$*, which is borderline Semi-Grim.*

Subjects overwhelmingly play strategies with Semi-Grim structure. They do not play mixtures of say $(1, 0, \alpha_3, 0)$, $(1, 0, 0, \alpha_4)$, $(1, \alpha_5, 0, 0)$, which on average would yield something akin to Semi-Grim. Aside from Treatment 6 in BOS, where subjects do virtually not cooperate, Semi-Grim has positive weight in all cases. In 12 of these 17 treatments, its weight is above 80%. In these cases, alternative strategies can be considered residual. Of the five remaining treatments, where the Semi-Grim weight is below 80%, four can be found in the DF data set. This seems to relate to the cohort-effect suspected in relation to Result 2 and Table 3. Overall, I consider the Semi-Grim hypothesis (Hypothesis 3) therefore to be confirmed both on average and individually.

Table 5: Estimated weights and randomization parameters of prototypical strategies $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$

			Constant $(\alpha,\alpha,\alpha,\alpha)$		**Gen Grim** $(1,\alpha,\alpha,\alpha)$		**Gen Coop** $(1,\alpha,0,0)$		**Gen TFT** $(1,0,\alpha,0)$		**Gen WSLS** $(1,0,0,\alpha)$		**Semi-Grim** $(1,\alpha,\alpha,0)$		
Treat	$\delta-\delta^{\star}$	γ	Weight	α	Weight	α	Weight	α	Weight	α	Weight	α	Weight	α	LL
Blonski et al. (2011)															
6	-0.3	0.001 (0)	1 (–)	0.017 (0.013)	–	–	–	–	–	–	–	–	–	–	−15.3
3	-0.3	0.015 (0.012)	–	–	–	–	–	–	–	–	–	–	1 (–)	0.143 (0.094)	−32.8
2	-0.15	0.003 (0.003)	–	–	–	–	–	–	–	–	–	–	1 (–)	0.286 (0.082)	−52.2
7	-0.05	0.006 (0.005)	–	–	–	–	–	–	–	–	–	–	1 (–)	0.326 (0.072)	−72.7
4	-0.05	0.008 (0.007)	–	–	0.12 (0.049)	0.252 (0.054)	–	–	–	–	–	–	0.88 (–)	0.167 (0.048)	−264.1
8	0.075	0.018 (0.01)	–	–	0.025 (0.025)	0.8 (0)	–	–	–	–	–	–	0.975 (–)	0.18 (0.04)	−170.6
5	0.075	0.015 (0.007)	0.05 (0.042)	0.218 (0.077)	–	–	0.118 (0.096)	0.868 (0.224)	–	–	–	–	0.831 (–)	0.277 (0.071)	−265.6
1	0.083	0.017 (0.01)	0.053 (0.039)	0.244 (0.088)	0.361 (0.094)	0.016 (0.011)	0.055 (0.054)	0.985 (0.082)	0.031 (0.045)	0.985 (0.118)	–	–	0.5 (–)	0.43 (0.042)	−278.6
9	0.083	0.022 (0.018)	0.052 (0.042)	0.259 (0.167)	–	–	–	–	–	–	–	–	0.948 (–)	0.193 (0.047)	−200.3
10	0.179	0.059 (0.017)	0.037 (0.056)	0.412 (0.067)	–	–	–	–	–	–	–	–	0.963 (–)	0.221 (0.044)	−352.4
Dal Bo and Fréchette (2011)															
1	-0.316	0.006 (0.004)	–	–	0.045 (0.033)	0.553 (0.129)	–	–	0.069 (0.047)	0.997 (0.046)	–	–	0.885 (–)	0.276 (0.064)	−187
3	-0.105	0.033 (0.018)	0.355 (0.077)	0.008 (0)	0.14 (0.048)	0.647 (0.069)	–	–	0.157 (0.062)	0.88 (0.161)	–	–	0.348 (–)	0.446 (0.083)	−513
2	-0.066	0.027 (0.023)	–	–	0.325 (0.071)	0.005 (0)	–	–	0.141 (0.061)	0.952 (0.065)	–	–	0.533 (–)	0.508 (0.046)	−562.6
5	0.105	0.004 (0.006)	–	–	0.299 (0.105)	0.137 (0.093)	0.402 (0.119)	0.97 (0.241)	–	–	–	–	0.299 (–)	0.552 (0.117)	−337.3
4	0.145	0.026 (0.023)	–	–	0.226 (0.084)	0.703 (0.103)	0.146 (0.073)	0.873 (0.146)	–	–	–	–	0.628 (–)	0.358 (0.054)	−401.8
6	0.355	0.017 (0.007)	0.023 (0.025)	0.5 (0)	0.074 (0.071)	0.796 (0.109)	–	–	–	–	–	–	0.902 (–)	0.227 (0.041)	−322.2
Duffy and Ochs (2009), "random rematching" treatment															
1	0.233	0.028 (0.022)	0.144 (0.048)	0.34 (0.047)	–	–	–	–	–	–	–	–	0.856 (–)	0.385 (0.033)	−815.7
Fudenberg et al. (2012), "no-noise" treatment															
1	0.475	0.021 (0.015)	0.083 (0.05)	0.419 (0.164)	–	–	–	–	–	–	–	–	0.917 (–)	0.434 (0.065)	−354.7

Note: Bootstrapped standard errors are provided in parentheses. Irrelevant components are identified (and then eliminated, as indicated by "–") based on the *ICL-BIC* information criterion, as described in the appendix. The right-most weight, usually that of the Semi-Grim component, is simply the difference of the remaining weights to 1. Thus, it is not a model parameter and is not assigned a standard error.

This raises the question how these individual strategies relate to Semi-Grim equilibria. To address it, the population weights of the relevant candidate MPEs are estimated in a finite mixture similar to the one above. The composition will be similar, too, while the prototypical strategies are now replaced by their MPE counterparts. In addition to the Always-Defect equilibrium $(\sigma_{cc},\sigma_{cd},\sigma_{dc},\sigma_{dd}) = (0,0,0,0)$ and the three Semi-Grim MPEs hypothesized above, I include four candidate MPEs that are efficient ($\sigma_{cc} = 1$) but not Semi-Grim—and their conjunction will again allow to reconstruct a Semi-Grim population. On the one hand, the three Semi-Grim MPEs are those that are obtained if players pick Semi-Grim equilibria satisfying two of the three criteria: efficiency, robustness to utility perturbations (RUP), and robustness to imperfect monitoring (RIM). The *limiting-logit, belief-free* MPE characterized in Prop. 4.3 is RIM and RUP, the *efficient belief-free Semi-Grim* MPE in Eq. (13) is efficient and RIM, and the *efficient limiting-logit Semi-Grim* MPE in Eq. (14) is efficient and RUP.

On the other hand, the four efficient MPEs that are not Semi-Grim are Grim, $(\sigma_{cc},\sigma_{cd},\sigma_{dc},\sigma_{dd}) = (1,0,0,0)$, and three equilibrium strategies related most closely to the prototypical strategies considered before: (1) Tit-for-tat $(\sigma_{cc},\sigma_{cd},\sigma_{dc},\sigma_{dd}) = (1,0,1,0)$, which of course is not an MPE but included as it represents the conventional notion of reciprocity most explicitly; (2) Weak Win-Stay-Lose-Shift (*W-WSLS*), which is an MPE of the form $(\sigma_{cc},\sigma_{cd},\sigma_{dc},\sigma_{dd}) = (1,0,0,\alpha)$, namely with[10]

$$\sigma_{dd} = \frac{bd^2-2d-b+a+1-\sqrt{b^2d^4-2\left(b^2-ab+b+2a-2\right)d^2+4(b-1)d+(b-a-1)^2}}{2d(bd^2-a)}; \tag{15}$$

and (3) the *asymmetric belief-free* MPE that obtains for $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ in Eq. (9),

$$\sigma_{cc} = 1 \qquad \sigma_{cd} = \frac{(b-1)d-b+a}{(a-1)d} \qquad \sigma_{dc} = \frac{1-d}{(a-1)d} \qquad \sigma_{cc} = 0. \tag{16}$$

The asymmetric belief-free MPE is particularly interesting, as it is maximizes the difference $\sigma_{cd} - \sigma_{dc}$ subject to efficiency ($\sigma_{cc} = 1$) and to $\sigma_{cc} > \sigma_{dd}$ as observed universally. It is therefore most suitable to reconstruct the Semi-Grim population without any individual playing a Semi-Grim strategy. To see this, recall that Proposition 4.1 shows

[10]This Weak WSLS strategy is a more promising candidate strategy than the pure WSLS $(1,0,0,1)$ for two reasons. First, pure WSLS is an MPE only if one round of punishment suffices to retaliate unilateral deviations, which is the case only in few treatments. Secondly, pure WSLS attracted virtually no weight in the analyses of Dal Bo and Fréchette (2011) and Fudenberg et al. (2012). Thus, replacing it by the less extreme "Weak WSLS" strategy improves its chances of attracting weight.

Table 6: The mixed MPEs $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$ in the analysis and bootstrapped confidence intervals of the empirical strategies (The MPEs are set bold face type when they exist, and plain when they would exist for some $\delta' \leq \delta^{2/3}$)

	Confidence Intervals				Mixed MPEs $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$				
Tr	$\hat{\sigma}_{cc}$	$\hat{\sigma}_{cd}$	$\hat{\sigma}_{dc}$	$\hat{\sigma}_{dd}$	LimLog BF	Eff BF SG	Eff LimLog SG	Asymm BF	W-WSLS
Blonski et al. (2011)									
6									$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{1})$
3	$[0.63, 1]$	$[0, 0.27]$	$[0, 0.4]$	$[0, 0.05]$					
2	$[0.89, 1]$	$[0, 0.48]$	$[0.08, 0.62]$	$[0, 0.01]$					
7	$[0.87, 1]$	$[0.1, 0.46]$	$[0.19, 0.56]$	$[0, 0.02]$	$(1, 0.87, 0.87, 0)$	$(1, 0.88, 0.88, 0)$	$(\mathbf{1}, \mathbf{0.67}, \mathbf{0.67}, \mathbf{0})$	$(1, 0.86, 1, 0)$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.84})$
4	$[0.81, 0.95]$	$[0.07, 0.24]$	$[0.11, 0.28]$	$[0.02, 0.05]$	$(1, 0.5, 0.5, 0)$	$(1, 0.5, 0.5, 0)$	$(1, 0.5, 0.5, 0)$	$(\mathbf{1}, \mathbf{0.33}, \mathbf{0.67}, \mathbf{0})$	
8	$[0.85, 0.97]$	$[0.11, 0.31]$	$[0.13, 0.35]$	$[0.02, 0.05]$	$(\mathbf{0.92}, \mathbf{0.8}, \mathbf{0.8}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.89}, \mathbf{0.89}, \mathbf{0.09})$	$(\mathbf{1}, \mathbf{0.33}, \mathbf{0.33}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.93}, \mathbf{0.5}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.57})$
5	$[0.87, 0.96]$	$[0.15, 0.34]$	$[0.22, 0.41]$	$[0.04, 0.07]$	$(\mathbf{0.96}, \mathbf{0.5}, \mathbf{0.5}, \mathbf{0.04})$	$(\mathbf{1}, \mathbf{0.54}, \mathbf{0.54}, \mathbf{0.09})$	$(\mathbf{1}, \mathbf{0.23}, \mathbf{0.23}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.71}, \mathbf{0.29}, \mathbf{0})$	
9	$[0.86, 0.96]$	$[0.07, 0.27]$	$[0.1, 0.31]$	$[0.01, 0.06]$	$(\mathbf{0.93}, \mathbf{0.59}, \mathbf{0.59}, \mathbf{0.04})$	$(\mathbf{1}, \mathbf{0.67}, \mathbf{0.67}, \mathbf{0.11})$	$(\mathbf{1}, \mathbf{0.3}, \mathbf{0.3}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.75}, \mathbf{0.42}, \mathbf{0})$	$(1, 0, 0, 1)$
1	$[0.7, 0.88]$	$[0.18, 0.36]$	$[0.23, 0.41]$	$[0.03, 0.07]$	$(\mathbf{0.94}, \mathbf{0.5}, \mathbf{0.5}, \mathbf{0.06})$	$(\mathbf{1}, \mathbf{0.56}, \mathbf{0.56}, \mathbf{0.11})$	$(\mathbf{1}, \mathbf{0.26}, \mathbf{0.26}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.67}, \mathbf{0.33}, \mathbf{0})$	
10	$[0.81, 0.93]$	$[0.07, 0.25]$	$[0.18, 0.38]$	$[0.03, 0.09]$	$(\mathbf{0.9}, \mathbf{0.43}, \mathbf{0.43}, \mathbf{0.14})$	$(\mathbf{1}, \mathbf{0.52}, \mathbf{0.52}, \mathbf{0.24})$	$(\mathbf{1}, \mathbf{0.13}, \mathbf{0.13}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.72}, \mathbf{0.17}, \mathbf{0})$	$(1, 0, 0, 0.61)$
Dal Bo and Fréchette (2011)									
1	$[0.7, 0.87]$	$[0.08, 0.34]$	$[0.32, 0.54]$	$[0.02, 0.04]$					
3	$[0.72, 0.89]$	$[0.17, 0.34]$	$[0.31, 0.45]$	$[0.1, 0.14]$	$(1, 0.57, 0.57, 0)$	$(1, 0.57, 0.57, 0)$	$(1, 0.53, 0.53, 0)$	$(\mathbf{1}, \mathbf{0.33}, \mathbf{0.87}, \mathbf{0})$	
2	$[0.86, 0.93]$	$[0.22, 0.38]$	$[0.34, 0.49]$	$[0.03, 0.05]$	$(1, 0.42, 0.42, 0)$	$(1, 0.42, 0.42, 0)$	$(1, 0.46, 0.46, 0)$	$(\mathbf{1}, \mathbf{0.14}, \mathbf{0.62}, \mathbf{0})$	
5	$[0.9, 0.96]$	$[0.25, 0.4]$	$[0.18, 0.31]$	$[0.05, 0.1]$	$(\mathbf{0.83}, \mathbf{0.73}, \mathbf{0.73}, \mathbf{0.04})$	$(\mathbf{1}, \mathbf{0.89}, \mathbf{0.89}, \mathbf{0.21})$	$(\mathbf{1}, \mathbf{0.33}, \mathbf{0.33}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.91}, \mathbf{0.57}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.54})$
4	$[0.92, 0.96]$	$[0.27, 0.46]$	$[0.46, 0.64]$	$[0.12, 0.22]$	$(\mathbf{0.89}, \mathbf{0.54}, \mathbf{0.54}, \mathbf{0.09})$	$(\mathbf{1}, \mathbf{0.65}, \mathbf{0.65}, \mathbf{0.19})$	$(\mathbf{1}, \mathbf{0.21}, \mathbf{0.21}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.78}, \mathbf{0.29}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.7})$
6	$[0.97, 0.99]$	$[0.18, 0.44]$	$[0.26, 0.51]$	$[0.06, 0.17]$	$(\mathbf{0.71}, \mathbf{0.64}, \mathbf{0.64}, \mathbf{0.18})$	$(\mathbf{1}, \mathbf{0.93}, \mathbf{0.93}, \mathbf{0.47})$	$(\mathbf{1}, \mathbf{0.14}, \mathbf{0.14}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.97}, \mathbf{0.19}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.23})$
Duffy and Ochs (2009), "random rematching" treatment									
	$[0.94, 0.97]$	$[0.27, 0.41]$	$[0.31, 0.45]$	$[0.11, 0.18]$	$(\mathbf{0.87}, \mathbf{0.5}, \mathbf{0.5}, \mathbf{0.13})$	$(\mathbf{1}, \mathbf{0.63}, \mathbf{0.63}, \mathbf{0.26})$	$(\mathbf{1}, \mathbf{0.09}, \mathbf{0.09}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.89}, \mathbf{0.11}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.2})$
Fudenberg et al. (2012), "no-noise" treatment									
6	$[0.93, 0.97]$	$[0.35, 0.58]$	$[0.38, 0.6]$	$[0.11, 0.24]$	$(\mathbf{0.73}, \mathbf{0.5}, \mathbf{0.5}, \mathbf{0.27})$	$(\mathbf{1}, \mathbf{0.77}, \mathbf{0.77}, \mathbf{0.54})$	$(\mathbf{1}, \mathbf{0.04}, \mathbf{0.04}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0.95}, \mathbf{0.05}, \mathbf{0})$	$(\mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{0.06})$

Note: LimLog BF is the limiting-logit, belief-free MPE characterized in Prop. 4.3, the *Eff BF SG* is the efficient belief-free Semi-Grim MPE in Eq. (13), *Eff LimLog SG* is the efficient limiting-logit Semi-Grim MPE in Eq. (14), *Asymm BF* is the asymmetric belief-free MPE in Eq. (16), and W-WSLS is the Weak WSLS MPE in Eq. (15).

that there are no limiting-logit MPEs satisfying $\sigma_{cd} > \sigma_{dc}$ and $\sigma_{cc} > \sigma_{dd}$. Hence, no such MPEs are regular in the sense of Doraszelski and Escobar (2010), and in turn, any MPE satisfying both $\sigma_{cc} > \sigma_{dd}$ and $\sigma_{cd} > \sigma_{dc}$ must be belief-free. Of those, the asymmetric belief-free MPE maximizes $\sigma_{cd} - \sigma_{dc}$, and since $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ holds true as well, it is the most suitable MPE to reconstruct the Semi-Grim population. The cooperation probabilities predicted by the mixed MPEs are provided in Table 6.

As for the analytical set up, one choice is left to be made: Which strategy should I use (if any) when an MPE does not exist? For example, the limiting-logit, belief-free equilibrium exists only if $\delta \geq \delta^\star$. We will see that subjects indeed play Semi-Grim strategies if $\delta \geq \delta^\star$. Now, setting the Semi-Grim weights to zero whenever $\delta < \delta^\star$ would thus effectively assume the result that subjects start playing Semi-Grim at $\delta = \delta^\star$. To avoid this effect, I relax the existence conditions slightly and consider the respective MPE also if it does not exist for the δ in question, but if δ is increased slightly (up to $\delta^{2/3}$).[11] In addition to allowing me to analyze the transition to Semi-Grim MPEs more effectively, this approach levels the playing field between the various MPEs and TFT (which is never an MPE but always considered) and it benefits Weak WSLS, which otherwise exists in the smallest number of treatments. The results are robust to varying the threshold of $\delta^{2/3}$, as I show in the supplementary material.

Table 7 presents the results. Accounting for the noise in experimental data, there is a surprisingly stable overall pattern. If δ is far below the threshold, roughly if $\delta - \delta^\star < -0.2$, subjects do not cooperate and play always defect. Around the threshold, roughly if $-0.2 \leq \delta - \delta^\star \leq 0.1$, subjects switch to Grim and Semi-Grim strategies, with Semi-Grim MPEs attracting around 50% on aggregate. The particular distribution of weight between the Semi-Grim MPEs is rather uninformative if $\delta \approx \delta^\star$, since the three Semi-Grim MPEs are quantitatively fairly similar around the threshold (as shown in Table 6). Above the threshold, roughly if $\delta - \delta^\star > 0.1$, the subjects have switched (almost) completely to Semi-Grim MPEs and distribute fairly evenly across the three Semi-Grim MPEs (these are the treatments BOS-10, DF-4 and 6, as well as DO and FRD). Thus, above the threshold, where subjects have to pick two of the three selection criteria (efficiency, RIM, and RUP), they consistently sort into either of the three cases. Alternative MPEs are only residual then, confirming Hypothesis 6. In particular, Weak

[11]To be precise, if a particular mixed MPE does not exist for the treatment parameters in question, I verify if it exists for $\delta^{2/3}$ (i.e. if the period length was shortened by 33% and δ gets adapted correspondingly), and in case it does, I use the mixed MPE at the threshold of existence.

Table 7: Estimated weights of equilibrium strategies

			Efficient MPEs that are not Semi-Grim				Semi-Grim MPEs				
Treat	$\delta - \delta^{\star}$	A-Def	Grim	TFT	W-WSLS	Asymm BF	LimLog BF	Eff Symm BF	Eff LimLog SG	γ	LL
Blonski et al. (2011)											
6	-0.3	1 (–)	–	–	–					0.02 (0.013)	−15.3
3	-0.3	1 (–)								0.08 (0.03)	−66.4
2	-0.15	–	1 (–)	–						0.01 (0.006)	−86.5
7	-0.05	–	–	–	–	–	–	–	1 (–)	0.01 (0.005)	−93.5
4	-0.05	0.537 (0.079)	–	0.161 (0.062)		–	–	–	0.303 (–)	0.026 (0.023)	−279.6
8	0.075	–	–	–	–	–	–	0.029 (0.03)	0.971 (–)	0.018 (0.009)	−180
5	0.075	–	–	–		–	–	0.24 (0.083)	0.76 (–)	0.012 (0.006)	−269
1	0.083	–	0.434 (0.084)	–		–	0.566 (–)	–	–	0.009 (0.008)	−300.1
9	0.083	0.116 (0.068)	–	–	–	–	–	–	0.884 (–)	0.033 (0.023)	−218.1
10	0.179	0.089 (0.057)	–	–	–	–	0.323 (0.094)	–	0.588 (–)	0.028 (0.027)	−346.9
Dal Bo and Fréchette (2011)											
1	-0.316	1 (–)								0.1 (0.038)	−487.1
3	-0.105	–	0.458 (0.077)	0.196 (0.073)		–	0.346 (–)	–	–	0.076 (0.023)	−604.5
2	-0.066	–	0.315 (0.072)	0.139 (0.064)		–	–	–	0.546 (–)	0.027 (0.024)	−564.3
5	0.105	–	0.448 (0.077)	–	–	–	–	0.15 (0.057)	0.403 (–)	0.01 (0.009)	−337.9
4	0.145	–	–	–	–	–	0.199 (0.08)	0.415 (0.103)	0.386 (–)	0.006 (0.005)	−398.9
6	0.355	–	–	–	–	–	0.092 (0.045)	–	0.908 (–)	0.012 (0.006)	−324.2
Duffy and Ochs (2009), "random rematching" treatment											
1	0.233	–	–	0.095 (0.049)	–	–	0.312 (0.075)	0.298 (0.081)	0.296 (–)	0.011 (0.005)	−822.2
Fudenberg et al. (2012), "no-noise" treatment											
1	0.475	–	–	0.137 (0.064)	–	–	0.165 (0.057)	0.423 (0.082)	0.275 (–)	0.005 (0.004)	−334.3

Note: The mixed MPEs are as described in Table 6. Bootstrapped standard errors are provided in parentheses. Empty cells indicate that the respective MPE does not exist even after inflating δ up to $\delta^{2/3}$. Hyphens ("–") indicate that the MPE exists but attracts insignificant weight according to *ICL-BIC*. Since the right-most weight is not a parameter but a (usually sizeable) residual, it is not assigned a standard error.

Figure 5: The aggregate share of subjects playing Semi-Grim MPEs

(a) ... in relation to $\delta - \delta^\star$ (b) ... in relation to Pr(Coop) in round 1

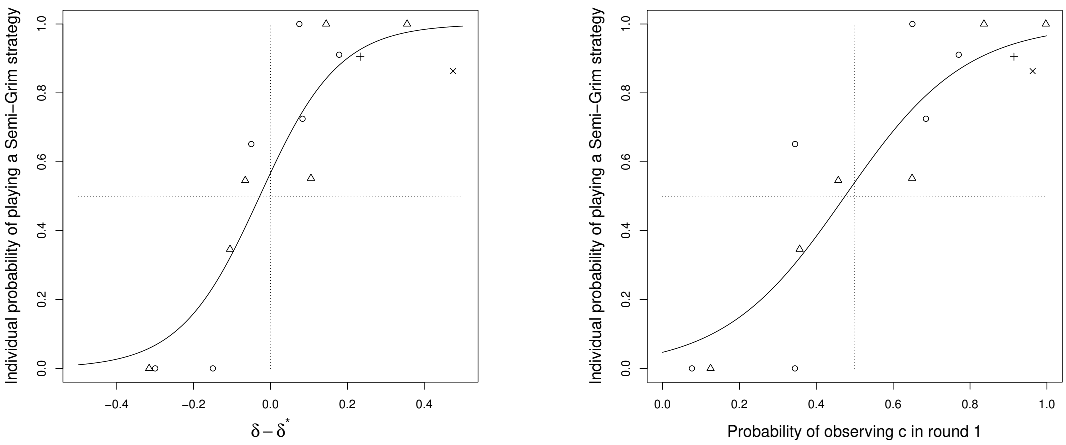

Table 8: Logistic regression of Pr(Semi-Grim) on $\delta - \delta^\star$ and cooperation

	Intercept	Coefficient
Pr(Semi-Grim) | $\delta - \delta^\star$	0.271 (0.302)	9.662* (2.319)
Pr(Semi-Grim) | Pr(Coop)	0.161 (0.275)	6.357* (1.333)

WSLS and the asymmetric belief-free MPE never attract any weight, and TFT is used only unsystematically (in 5/18 treatments), never attracting more than 20% weight.

Figure 5 illustrates this overall description by plotting the share of subjects playing either of the Semi-Grim MPEs in relation to both $\delta - \delta^\star$ and the probability of observing cooperation in round 1. The coefficients of the respective logistic regressions are provided in Table 8. The plot in relation to cooperation in round 1 is particularly illustrative. Subjects switch to Semi-Grim almost exactly as they start cooperating in round 1, both visually in Figure 5b and econometrically as the intercept is insignificant (Table 8), and the relationship is virtually linear.

Result 7 (Semi-Grim MPEs)**.** *There is a stable pattern across experiments that subjects switch to Semi-Grim MPEs as they start cooperating in round 1, and both occurs at the threshold* $\delta = \delta^\star$ *where the Semi-Grim strategies turn into equilibrium strategies.*

7 Discussion

The purpose of the paper was to analyze the strategies underlying the observation of Blonski, Ockenfels, and Spagnolo (2011, BOS) and Dal Bo and Fréchette (2011, DF) that experimental subjects start to cooperate when the discount rate δ reaches the seemingly abstract threshold $\delta^{\star} = (p_{dc} + p_{dd} - p_{cd} - p_{cc})/(p_{dc} - p_{cd})$. To this end, I first computed the average memory-1 strategies in a variety of recent experiments. This revealed a stable Semi-Grim pattern $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (1, \alpha, \alpha, 0)$ with $\alpha \in [0.2, 0.5]$, across all 18 treatments of four recent experiments.

In aggregate, my results explain both the Semi-Grim pattern and the predictiveness of the BOS-threshold. Semi-Grim equilibria are implied by requiring robustness to imperfect monitoring (RIM) and utility perturbations (RUP), they exist once the discount rate exceeds the BOS-threshold $\delta^{\star}$, that is when subjects start to cooperate, and subjects indeed play Semi-Grim strategies in most treatments. Three cases need to be distinguished here. If the discount rate is far below the threshold $\delta^{\star}$, players hardly cooperate in round 1, and Semi-Grim thus predicts rare cooperation in subsequent rounds. In this case, subjects are most parsimoniously categorized as playing Always Defect. Around the threshold $\delta^{\star}$, the population is a mixture of Grim and Semi-Grim MPEs, which have the structure $(\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}) = (1, 0, 0, 0)$ and $(1, \alpha, \alpha, 0)$, respectively. This yields a Semi-Grim structure $(1, \alpha', \alpha', 0)$ on average, just as observed. Above the threshold, finally, most subjects play Semi-Grim MPEs, but in this case they have to pick two of the three selection criteria (efficiency, RIM, and RUP). Subjects then distribute fairly evenly across the three resulting MPEs, and their mixture of course yields Semi-Grim on average, again. Thus, it is not a single selection principle, but a mixture of three, and if the respective predictions are sufficiently different (far above the threshold $\delta^{\star}$), subjects distribute evenly in all experiments.

In addition, the Semi-Grim equilibria explain that individual choices may sometimes be lenient (Grim only after two defections) or forgiving (Tit-for-Two-Tats, Two-Tits-for-Tat), as observed by Fudenberg et al. (2012). As subjects cooperate with symmetric, intermediate probabilities $\alpha \in [0.2, 0.5]$ after mixed memory-1 histories, (c, d) and (d, c), they do not reciprocate systematically. This yields either lenient or forgiving behavior depending on how the coin falls.

The absence of direct reciprocity, $\sigma_{cd} = \sigma_{dc}$, which was predicted and observed, marks a departure from the literature following Axelrod (1980a,b). This literature emphasizes the theoretical effectiveness of strategies such as Tit-for-Tat (TFT) and Win-Stay-Lose-Shift (WSLS) in evolutionary tournaments (Nowak et al., 1993; Imhof et al., 2007). The initial inspection of the average choices, which revealed $\sigma_{cd} = \sigma_{dc}$ as well as $\sigma_{dd} \leq 0.1$, suggested that neither TFT nor WSLS would have substantial weight, and the analysis of individual strategies confirmed this hypothesis. This observation accompanies Press and Dyson (2012), who recently showed that the "zero-determinant" strategies generalize TFT in that they are unbeatable by any opponent strategy (see also Duersch et al., 2013). These zero-determinant strategies contain the "belief-free" equilibria (Ely and Välimäki, 2002) considered here. The finding that these equilibria explain both the Semi-Grim pattern across experiments and the BOS-threshold thus relates neatly to the recent discussion of zero-determinant strategies. Finally, the existence of belief-free equilibria is guaranteed in general repeated games (Ely et al., 2005), which suggests that the results obtained in the present paper will generalize as well. This appears to be a promising avenue for further research.

References

Abreu, D., Pearce, D., and Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063.

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.

Bhaskar, V., Mailath, G., and Morris, S. (2008). Purification in the infinitely-repeated prisoners' dilemma. *Review of Economic Dynamics*, 11(3):515–528.

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.

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.

Breitmoser, Y. (2012). Strategic reasoning in p-beauty contests. *Games and Economic Behavior*, 75(2):555–569.

Breitmoser, Y., Tan, J., and Zizzo, D. (2010). Understanding perpetual r&d races. *Economic Theory*, 44(3):445–467.

Bruttel, L. and Kamecke, U. (2012). Infinity in the lab. how do people play repeated games? *Theory and Decision*, 72(2):205–219.

Capra, C., Goeree, J., Gomez, R., and Holt, C. (1999). Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3):678–690.

Choi, S., Gale, D., and Kariv, S. (2008). Sequential equilibrium in monotone games: A theory-based analysis of experimental data. *Journal of Economic Theory*, 143(1):302–330.

Dal Bo, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 95(5):1591–1604.

Dal Bo, P. and Fréchette, G. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429.

Doraszelski, U. and Escobar, J. F. (2010). A theory of regular markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoretical Economics*, 5:369–402.

Duersch, P., Oechssler, J., and Schipper, B. C. (2013). When is tit-for-tat unbeatable? *International Journal of Game Theory*.

Duffy, J. and Ochs, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, 66(2):785–812.

Ely, J., Hörner, J., and Olszewski, W. (2005). Belief-free equilibria in repeated games. *Econometrica*, 73(2):377–415.

Ely, J. and Välimäki, J. (2002). A robust folk theorem for the prisoner's dilemma. *Journal of Economic Theory*, 102(1):84–105.

Engle-Warnick, J. and Slonim, R. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior and Organization*, 55(4):553–573.

Engle-Warnick, J. and Slonim, R. (2006). Inferring repeated-game strategies from actions: evidence from trust game experiments. *Economic Theory*, 28(3):603–632.

Erev, I. and Roth, A. (2001). Simple reinforcement learning models and reciprocation in the prisoner's dilemma game. In Gigerenzer, G. and Selten, R., editors, *Bounded rationality: The adaptive toolbox*, pages 215–231. MIT Press.

Evans, R. and Maskin, E. (1989). Efficient renegotiation-proof equilibria in repeated games. *Games and Economic Behavior*, 1:361–369.

Fey, M., McKelvey, R., and Palfrey, T. (1996). An experimental study of constant-sum centipede games. *International Journal of Game Theory*, 25(3):269–287.

Fudenberg, D. and Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554.

Fudenberg, D., Rand, D., and Dreber, A. (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*.

Goeree, J., Holt, C., and Laury, S. (2002a). Private costs and public benefits: Unraveling the effects of altruism and noisy behavior. *Journal of Public Economics*, 83(2):255–276.

Goeree, J., Holt, C., and Palfrey, T. (2002b). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104(1):247–272.

Imhof, L., Fudenberg, D., and Nowak, M. (2007). Tit-for-tat or win-stay, lose-shift? *Journal of Theoretical Biology*, 247(3):574–580.

Kandori, M. (2002). Introduction to repeated games with private monitoring. *Journal of Economic Theory*, 102(1):1–15.

Luce, R. (1959). *Individual choice behavior*. Wiley New York.

McKelvey, R. and Palfrey, T. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

McLachlan, G. and Peel, D. (2000). *Finite Mixture Models*. Wiley series in probability and statistics.

Murnighan, J. and Roth, A. (1983). Expecting continued play in prisoner's dilemma games. *Journal of Conflict Resolution*, 27(2):279–300.

Nowak, M., Sigmund, K., et al. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, 364(6432):56–58.

Powell, M. (2006). The newuoa software for unconstrained optimization without derivatives. *Large-Scale Nonlinear Optimization*, pages 255–297.

Press, W. H. and Dyson, F. J. (2012). Iterated prisonerâĂŹs dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences*, 109(26):10409–10413.

Rapoport, A. and Mowshowitz, A. (1966). Experimental studies of stochastic models for the prisoner's dilemma. *Behavioral Science*, 11(6):444–458.

Roth, A. and Murnighan, J. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17(2):189–198.

Stahl, D. et al. (1991). The graph of prisoners' dilemma supergame payoffs as a function of the discount factor. *Games and Economic Behavior*, 3(3):368–384.

Stahl, D. and Wilson, P. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327.

Stahl, D. and Wilson, P. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.

Stewart, A. J. and Plotkin, J. B. (2012). Extortion and cooperation in the prisonerâĂŹs dilemma. *Proceedings of the National Academy of Sciences*, 109(26):10134–10135.

Van Damme, E. (1989). Renegotiation-proof equilibria in repeated prisoners' dilemma. *Journal of Economic Theory*, 47(1):206–217.

A Econometric analysis of individual behavior

Essentially, my approach adapts Dal Bo and Fréchette (2011) to enable the inclusion of mixed equilibria, but is otherwise similar. Thus, subjects (of a given type) play their equilibrium action with probability $1-\gamma$, $\gamma \in (0,1)$, and they randomize uniformly with probability γ (in each round, taking independent draws). This approach estimates the same population weights as Dal Bo and Fréchette's approach if only pure equilibria are considered, while it provides a straightforward generalization to mixed MPEs. The population is a mixture of a finite number of components $k \in K$, while $o_{s,t} \in \{0,1\}$ and $\omega_{s,t} = \{\emptyset, cc, cd, dc, dd\}$ denote choice and state of the decision number t of subject $s \in S$, respectively ($o_{s,t} = 1$ denotes cooperation and $o_{s,t} = 0$ denotes defection). Let $\sigma_\omega(k)$ denote the (perturbed) probability that members of component k cooperate in state ω and let $\rho(k)$ denote the component weight. Thus, the log-likelihood of the model $\langle \sigma, \rho, K \rangle$ is

$$LL = \sum_{s \in S} \log \sum_{k \in K} \rho(k)\, L(s,k) \qquad \text{with } L(s,k) = \prod_t \left(\sigma_{\omega_{s,t}}(k)\right)^{o_{s,t}} \cdot \left(1 - \sigma_{\omega_{s,t}}(k)\right)^{1-o_{s,t}}.$$

The likelihood is maximized jointly over all parameters (σ, ρ), first using the robust, gradient-free NEWUOA algorithm (Powell, 2006) and secondly, verifying convergence using a Newton-Raphson algorithm. Standard errors are bootstrapped.

The model dimensionality K is estimated using *ICL-BIC* (integrated classification likelihood-Bayes information criterion), which is an entropy-based generalization of Bayes information criterion appropriate to discriminate finite-mixture models (Biernacki et al., 2000). Models with poorly distinguished components have high entropy, and in such cases, *ICL-BIC* recommends to eliminate a component. It is defined as

$$\begin{aligned} \textit{ICL-BIC} &= -LL + D/2 \cdot \ln n + \mathrm{En}(\hat{\tau}) \\ &\qquad \text{with } \mathrm{En}(\hat{\tau}) = -\sum_{s \in S} \sum_{k \in K} \hat{\tau}_{sk} \ln \hat{\tau}_{sk} \quad \text{with} \quad \hat{\tau}_{sk} = \frac{\rho_k L(s,k)}{\sum_{k' \in K} \rho_{k'} L(s,k')}, \end{aligned} \tag{17}$$

with n as number of subjects and D as number of parameters. Note that $\hat{\tau}_{sk}$ is the posterior probability that subject s belongs to component k. The precise procedure to estimate K is as follows. I start with the complete model containing all components and estimate all nested models where exactly one component is eliminated. I rank all the nested models according to their log-likelihood (LL) and start with the model with the highest LL. If eliminating the respective component improves *ICL-BIC*, it is eliminated and the procedure restarts with the correspondingly reduced number of components. Otherwise, the next-ranked component is considered, and so on. The procedure stops if no component can be eliminated. Details on the intermediate outcomes and on the elimination order are provided as supplementary material.

B Proofs

Proof of Proposition 4.1. Recall $\tilde{\pi}_{s',s''}$ as defined in Eqs. (7)–(8) and that $\log\left[(1-\sigma_{s',s''})/\sigma_{s',s''}\right] = -\lambda\tilde{\pi}_{s',s''}$ in all (s',s'') if (σ,σ) is an MLE.

First, by $\tilde{\pi}_{cc}-\tilde{\pi}_{dd} = (\sigma_{dd}-\sigma_{cc})\cdot\mu$, $\sigma_{cc}\neq\sigma_{dd}$ implies $\mu<0$. For, $\mu=0$ implies $\sigma_{cc}=\sigma_{dd}$, and in case $\mu>0$, $\sigma_{cc}\gtrless\sigma_{dd}$ implies $\tilde{\pi}_{cc}\lessgtr\tilde{\pi}_{dd}$ and thus $\sigma_{cc}\lessgtr\sigma_{dd}$ (contradiction). In turn, $\mu<0$ implies $\sigma_{cd}=\sigma_{dc}$, as in case $\mu<0$, $\sigma_{cd}\gtrless\sigma_{dc}$ implies $\tilde{\pi}_{cd}\lessgtr\tilde{\pi}_{dc}$ and thus $\sigma_{cd}\lessgtr\sigma_{dc}$, a contradiction to $\sigma_{cd}\neq\sigma_{dc}$. It remains to show that $\sigma_{dc}<\sigma_{cc}$, or equivalently $\tilde{\pi}_{cc}<\tilde{\pi}_{dc}$. Using $\sigma_{cd}=\sigma_{dc}$, $\tilde{\pi}_{cc}-\tilde{\pi}_{dc}$ simplifies toward

$$\tilde{\pi}_{cc}-\tilde{\pi}_{dc} = \frac{(d-1)(\sigma_{dc}-\sigma_{cc})(d(b\sigma_{dd}-2a\sigma_{dd}+2a\sigma_{dc}-2\sigma_{dc}-b\sigma_{cc}+2\sigma_{cc})+b-a-1)}{d\left(\sigma_{dd}^2-2\sigma_{dd}-2\sigma_{dc}^2+2\sigma_{dc}+\sigma_{cc}^2\right)+2d^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-1}.$$

For contradiction assume $\sigma_{dc}\geq\sigma_{cc}$. The denominator of the fraction is generally decreasing in σ_{dd}, and in the limiting case $\sigma_{dd}=0$ it is

$$2d\,(d\sigma_{cc}-1)\,\sigma_{dc}^2-2d\,\left(d\sigma_{cc}^2-1\right)\sigma_{dc}+d\sigma_{cc}^2-1<0.$$

Thus, it is generally negative. The numerator of the right-hand side is negative if

$$\sigma_{dd}<\frac{(2a-2)\,d\sigma_{dc}+(2-b)\,d\sigma_{cc}+b-a-1}{(2a-b)\,d}=:\tilde{\sigma}_{dd}.$$

Thus, in case $\sigma_{dd}<\tilde{\sigma}_{dd}$, $\tilde{\pi}_{cc}-\tilde{\pi}_{dc}$ is positive, contradicting the initial assumption $\sigma_{cc}<\sigma_{dc}$. Alternatively, in case $\sigma_{dd}\geq\tilde{\sigma}_{dd}$, the cooperation incentive $\tilde{\pi}_{dc}$ is decreasing in σ_{dd}, and in the limiting case $\sigma_{dd}=\tilde{\sigma}_{dd}$,

$$\tilde{\pi}_{dc}=\frac{(d-1)\,(bd\sigma_{dc}-bd\sigma_{cc}+b-a)}{2d\sigma_{dc}-2d\sigma_{cc}+1}.$$

Thus, $\tilde{\pi}_{dc} < 0$ follows if $\sigma_{dc} > \sigma_{cc}$. Since $\tilde{\pi}_{dc} < 0$ also implies $\sigma_{dc} < 0.5$, this contradicts $\max\{\sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd}\} > 0.5$. $\square$

Proof of Proposition 4.2. Eqs. (7)–(8) hold equivalently here, now with $\mu = r_1/r_2$ where

$$\begin{aligned} r_1 = \delta\,(p_{dc}+p_{cd})\,(\sigma_{dd}-\sigma_{cc}) - 2\,\delta\,p_{cc}\,(\sigma_{dd}-\sigma_{dc}) \\ - 2\,\delta\,p_{dd}\,(\sigma_{dc}-\sigma_{cc}) - p_{dd} + p_{dc} + p_{cd} - p_{cc} \end{aligned}$$

and $r_2 \neq 0$. Thus, $r_1 = 0$ again yields $\tilde{\pi}_{cc} = \tilde{\pi}_{cd} = \tilde{\pi}_{dc} = \tilde{\pi}_{dd}$. Solving $r_1 = 0$ for σ_{dc},

$$\sigma_{dc} = \frac{2\,\delta\,(p_{cc}\,\sigma_{dd} - p_{dd}\,\sigma_{cc}) - \delta\,(p_{dc}+p_{cd})\,(\sigma_{dd}-\sigma_{cc}) + p_{dd} - p_{dc} - p_{cd} + p_{cc}}{2\,\delta\,(p_{cc}-p_{dd})},$$

and substituting this into $\tilde{\pi}_{cc} = 0$ yields

$$\begin{aligned} \delta^2\,(p_{dc}-p_{cd})\,(\sigma_{dd}-\sigma_{cc}) - \delta\,(p_{dc}-p_{cd})\,(\sigma_{dd}-\sigma_{cc}-1) \\ + \delta\,(p_{dd}-p_{cc}) - p_{dd} - p_{dc} + p_{cd} + p_{cc} = 0. \end{aligned}$$

Solving these two conditions for $(\sigma_{dd}, \sigma_{dc})$ yields

$$\begin{aligned} \sigma_{dd} &= \frac{(p_{dc}-p_{cd})\,\delta\,\sigma_{cc} - p_{dd} - p_{dc} + p_{cd} + p_{cc}}{\delta\,(p_{dc}-p_{cd})}, \\ \sigma_{dc} = \sigma_{cd} &= \frac{(p_{dc}-p_{cd})\,\delta\,\sigma_{cc} - p_{dc} + p_{cc}}{\delta\,(p_{dc}-p_{cd})}, \end{aligned}$$

and rearranging yields Eqs. (11). As for existence, $\sigma_{dd} \geq 0$ holds true (at $\sigma_{cc} = 1$) iff $\delta \geq (p_{dc}+p_{dd}-p_{cd}-p_{cc})/(p_{dc}-p_{cd})$, while $\sigma_{dd} < \sigma_{dc} \leq 1$ is satisfied for all $\sigma_{cc} \in [0,1]$. $\square$

Proof of Proposition 4.3. First, I show that any limiting-logit, belief-free MPE at $\delta \approx \delta_{\mathrm{BOS}}$ must satisfy $\sigma_{cd} = \sigma_{dc}$. By Prop. 4.1, any limiting-logit MPE with $\sigma_{cd} \neq \sigma_{dc}$ satisfies $\sigma_{cc} = \sigma_{dd}$. Since belief-free MPEs satisfy Eq. (9), $\sigma_{cc} = \sigma_{dd}$ implies

$$\sigma_{cd} = \frac{(a-1)\,\delta\sigma_{cc} + a - b}{(a-1)\,\delta} \qquad \sigma_{dc} = \frac{(a-1)\,\delta\sigma_{cc} + 1}{(a-1)\,\delta}, \tag{18}$$

which is a strategy profile only if $\delta > (b-a+1)/(a-1)$. Thus, no such strategy profile exists at $\delta \approx \delta_{\mathrm{BOS}}$.

In turn, if a limiting-logit, belief-free equilibrium exists at $\delta \approx \delta_{\text{BOS}}$, then it satisfies $\sigma_{cd} = \sigma_{dc}$. By Prop. 4.2, $\delta \geq \delta_{\text{BOS}}$ is necessary for a belief-free MPE satisfying $\sigma_{cd} = \sigma_{dc}$ to exist, and the following shows that $\delta \geq \delta_{\text{BOS}}$ is sufficient for a limiting-logit, belief-free equilibrium to exist. Since $\sigma_{dc} = \sigma_{cd}$ is also necessary, $\delta \geq \delta_{\text{BOS}}$ will thus be established as a necessary and sufficient condition.

Since it must be an MPE as characterized in Prop. 4.2, Eq. (11) applies for some $r \in [0,1]$. By definition of MLE, $\log\left[(1-\sigma_{s',s''})/\sigma_{s',s''}\right] \cdot f = -\tilde{\pi}_{s',s''}$ for all (s',s'') and some $f > 0$, where $\tilde{\pi}_{s',s''}$ as defined in Eqs. (7)–(8) expand to

$$\tilde{\pi}_{cc} = \frac{\begin{gathered}(\delta-1)\Big(b\delta\left(\sigma_{dd}^2-\sigma_{cc}\sigma_{dd}-\sigma_{dd}-\sigma_{dc}^2+\sigma_{dc}+\sigma_{cc}^2\right)-a\delta(\sigma_{dd}-\sigma_{dc})(\sigma_{dd}+\sigma_{dc}-2\sigma_{cc})\\ +b\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-\delta(\sigma_{dc}-\sigma_{cc})^2-b\sigma_{cc}+a\sigma_{cc}+\sigma_{cc}-1\Big)\end{gathered}}{\delta\left(\sigma_{dd}^2-2\sigma_{dd}-2\sigma_{dc}^2+2\sigma_{dc}+\sigma_{cc}^2\right)+2\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-1}$$

$$\tilde{\pi}_{dc} = \frac{\begin{gathered}(\delta-1)\Big(b\delta\left(\sigma_{dd}^2-\sigma_{dc}\sigma_{dd}-\sigma_{dd}-\sigma_{dc}^2+\sigma_{cc}\sigma_{dc}+\sigma_{dc}\right)-a\delta(\sigma_{dd}-\sigma_{dc})^2\\ +b\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})+\delta(\sigma_{dc}-\sigma_{cc})^2-b\sigma_{dc}+a\sigma_{dc}+\sigma_{dc}-1\Big)\end{gathered}}{\delta\left(\sigma_{dd}^2-2\sigma_{dd}-2\sigma_{dc}^2+2\sigma_{dc}+\sigma_{cc}^2\right)+2\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-1}$$

$$\tilde{\pi}_{dd} = \frac{\begin{gathered}(\delta-1)\Big(b\delta\left(\sigma_{cc}\sigma_{dd}-\sigma_{dd}-\sigma_{dc}^2+\sigma_{dc}\right)+a\delta(\sigma_{dd}-\sigma_{dc})^2\\ +\delta(\sigma_{dc}-\sigma_{cc})(2\sigma_{dd}-\sigma_{dc}-\sigma_{cc})+b\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-b\sigma_{dd}+a\sigma_{dd}+\sigma_{dd}-1\Big)\end{gathered}}{\delta\left(\sigma_{dd}^2-2\sigma_{dd}-2\sigma_{dc}^2+2\sigma_{dc}+\sigma_{cc}^2\right)+2\delta^2(\sigma_{dc}-\sigma_{cc})(\sigma_{dd}-\sigma_{cc})(\sigma_{dd}-\sigma_{dc})-1}.$$

Let $\sigma(r)$ denote the MPEs in Eq. (11) and substitute σ_{cc}, σ_{dd} in all expressions by

$$\sigma_{cc} = \sigma_{cc}(r) + w_1\big(\sigma_{dc} - \sigma_{dc}(r)\big) \qquad \sigma_{dd} = \sigma_{dd}(r) + w_2\big(\sigma_{dc} - \sigma_{dc}(r)\big).$$

Let $f_{s',s''} = \log\left[(1-\sigma_{s',s''})/\sigma_{s',s''}\right] \cdot f + \tilde{\pi}_{s',s''}$ for all (s',s'') denote the MLE conditions and totally differentiate the MLE condition for state (d,c), i.e. f_{dc}, with respect to (σ_{dc}, f). This yields an expression for $d\sigma_{dc}/df$. Next, define

$$d_{cc} := \frac{\partial f_{cc}}{\partial \sigma_{dc}} \cdot \frac{d\sigma_{dc}}{df} + \frac{\partial f_{cc}}{\partial f} \qquad d_{dd} := \frac{\partial f_{dd}}{\partial \sigma_{dc}} \cdot \frac{d\sigma_{dc}}{df} + \frac{\partial f_{dd}}{\partial f}$$

and substitute $\sigma_{dc}(r)$ as defined in Eq. (11) for σ_{dc}.

Thus, by the implicit function theorem, $\sigma(r)$ is a limiting-logit equilibrium if $d_{cc} = d_{dd} = 0$ along some direction $(w_1, w_2) \neq 1$ at $f = 0$. Note that $d_{cc} = d_{dd} = 0$ holds trivially for the direction $(w_1, w_2) = 1$ (at $f = 0$), as we then move along the line of belief-free Semi-Grim MPEs where the MPE conditions hold trivially for $f = 0$. Thus, the objective is to find $r \in (0,1)$ such that $d_{cc} = d_{dd} = 0$ for $(w_1, w_2) \neq 1$. To this end, solve $[d_{cc} = 0, d_{dd} = 0]$ for (w_1, w_2). The solution is of the form

$$w_1 = \frac{n_1 * f + x}{m_1 * f + x} \qquad\qquad w_2 = \frac{n_2 * f + x}{m_2 * f + x},$$

and thus it satisfies $(w_1, w_2) \neq 1$ if $x = 0$, where $x = x' * x''$ with

$$x' = \frac{(\delta-1)^2\,(b\delta-b+a-1)^2\,(r-1)\,r\,((b\delta-b+a-1)\,r+1)}{((b\delta-b+a-1)\,r+b-a+1)\,((b\delta-b+a-1)\,r-b\delta)\,((b\delta-b+a-1)\,r-b\delta+1)}$$

$$x'' = \Big(\log\Big(-\tfrac{(b\delta-b+a-1)\,(r-1)}{(b\delta-b+a-1)\,r+b-a+1}\Big)-(b-a+1)\,\log\Big(-\tfrac{(b\delta-b+a-1)\,r-b\delta+1}{b\delta}\Big)+b\,\log\Big(-\tfrac{(b\delta-b+a-1)\,r-b\delta}{(b\delta-b+a-1)\,r}\Big)$$
$$-a\,\log\Big(-\tfrac{(b\delta-b+a-1)\,r-b\delta}{(b\delta-b+a-1)\,r}\Big)+(b-a+1)\,\log\Big(\tfrac{(b\delta-b+a-1)\,r+1}{b\delta}\Big)\Big).$$

Note that $x' \neq 0$ for interior r follows from $\delta \geq \delta_{BOS}$, which implies $b\delta - b + a - 1 \geq 0$. Thus, $x = 0$ only if $x'' = 0$, which applies if

$$\left(\frac{((-b\delta+b-a+1)\,r+b\delta)\,((b\delta-b+a-1)\,r+1)}{(b\delta-b+a-1)\,r\,((-b\delta+b-a+1)\,r+b\delta-1)}\right)^{b-a} = \frac{((b\delta-b+a-1)\,r+b-a+1)\,((b\delta-b+a-1)\,r-b\delta+1)}{(b\delta-b+a-1)\,(r-1)\,((b\delta-b+a-1)\,r+1)}.$$

The left-hand side is finite for $r = 1$ and infinite for $r = 0$, while the right-hand side is infinite for $r = 1$ and finite for $r = 0$. Both are continuous in r, and thus an interior solution $r \in (0,1)$ exists. □

Lemma B.1. *A regular Semi-Grim MPE exists for all $\delta > \delta_{BOS}$ in general, and if $p_{cc} + p_{dd} > p_{dc} + p_{cd}$, then also for all*

$$\delta > 1 - \frac{\sqrt{2\sqrt{p_{cc}-p_{cd}}\,\sqrt{p_{dc}-p_{cc}}\,\sqrt{p_{dd}-p_{cd}}\,\sqrt{p_{dc}-p_{dd}}+(p_{dc}+p_{cd}-2\,p_{cc})\,p_{dd}+(p_{cc}-2\,p_{cd})\,p_{dc}+p_{cc}\,p_{cd}}}{p_{dc} - p_{cd}}. \tag{19}$$

Proof of Lemma B.1. If $\sigma_{cc} = 1$, $\sigma_{cd} = \sigma_{dc}$, and $\sigma_{dd} = 0$, the cooperation incentive in state (d,c), $\tilde{\pi}_{dc} := \pi_{dc}(c) - \tilde{\pi}_{dc}(d)$, is

$$\tilde{\pi}_{dc} = \frac{\delta\,(p_{dd}+\delta\,(p_{dc}-p_{cd})-p_{dc}+p_{cd}-p_{cc})\,\sigma_{dc}^2-\big(\delta^2\,(p_{dc}-p_{cd})+2\,\delta\,(p_{dd}-p_{dc})+p_{dc}+p_{cd}-p_{dd}-p_{cc}\big)\,\sigma_{dc}-(1-\delta)\,(p_{dd}-p_{cd})}{2\,\delta\,(\sigma_{dc}-1)\,\sigma_{dc}+1}. \tag{20}$$

First, I show that the two conditions $\tilde{\pi}_{cc} > \tilde{\pi}_{dc}$ and $\tilde{\pi}_{dc} = 0$ imply that σ is a mixed MPE. By $\sigma_{dc} = \sigma_{cd}$ and Eq. (8), $\tilde{\pi}_{dc} = \tilde{\pi}_{cd}$, i.e. $\tilde{\pi}_{dc} = 0$ implies $\tilde{\pi}_{cd} = 0$. Further, by

$\tilde{\pi}_{cc} > \tilde{\pi}_{dc}$ and Eqs. (7)–(8), $\sigma_{cc} = 1 > \sigma_{dc}$ implies $\mu < 0$, and by $\sigma_{dd} = 0 < \sigma_{dc}$ this implies $\tilde{\pi}_{dd} < \tilde{\pi}_{dc} = 0$. Hence, any strategy profile satisfying $\tilde{\pi}_{cc} > \tilde{\pi}_{dc} = 0$ (besides $\sigma_{cc} = 1, \sigma_{dd} = 0$) is mixed MPE with the claimed incentive structure.

Second, I derive the existence condition. $\tilde{\pi}_{dc} = 0$ obtains if

$$\sigma_{dc} = \frac{(2\delta - 1)\, p_{dd} + (1-\delta)^2\, p_{dc} + \left(1-\delta^2\right)\, p_{cd} - p_{cc} \pm \sqrt{r}}{(2\delta^2 - 2\delta)\,(p_{dc} - p_{cd}) - 2\delta\,(p_{cc} - p_{dd})} \tag{21}$$

with

$$\begin{aligned} r = {} & (p_{dd} - p_{dc} - p_{cd} + p_{cc})^2 + 4\delta\left((p_{dc} + p_{cd})\, p_{dd} + p_{cc}\,(p_{dc} + p_{cd} - 2\,p_{dd}) - p_{dc}^2 - p_{cd}^2\right) \\ & - 2\delta^2\left((p_{dc} + p_{cd})\, p_{dd} + p_{cc}\,(p_{dc} + p_{cd} - 2\,p_{dd}) - 3\,p_{dc}^2 + 4\,p_{cd}\,p_{dc} - 3\,p_{cd}^2\right) \\ & + \delta^4\,(p_{dc} - p_{cd})^2 - 4\delta^3\,(p_{dc} - p_{cd})^2 \end{aligned}$$

These strategy profiles exist if $r \geq 0$, and solving $r = 0$ for δ, this yields the lower bound claimed in Eq. (19). Now, evaluating $\tilde{\pi}_{cc} - \tilde{\pi}_{dc}$ at $\sigma_{cc} = 1$, $\sigma_{cd} = \sigma_{dc}$, $\sigma_{dd} = 0$ yields

$$\tilde{\pi}_{cc} - \tilde{\pi}_{dc} = \frac{(1-\sigma_{dc})\,(\delta\,(2\,p_{dd}\,\sigma_{dc} - 2\,p_{cc}\,\sigma_{dc} - 2\,p_{dd} + p_{dc} + p_{cd}) + p_{dd} - p_{dc} - p_{cd} + p_{cc})}{2\delta\,(\sigma_{dc} - 1)\,\sigma_{dc} + 1} \tag{22}$$

and at the limiting strategy $\sigma_{dc}\big|_{r=0}$, it is positive if and only if

$$\frac{(d-1)^2\,(p_{dc} - p_{cd})\,(p_{dd} - p_{dc} - p_{cd} + p_{cc})}{p_{cc} - p_{dd} + (1-\delta)\,(p_{dc} - p_{cd})} > 0. \tag{23}$$

This is satisfied if and only if $p_{cc} + p_{dd} > p_{dc} + p_{cd}$. Otherwise, the limiting strategy σ_{dc} does not solve $r = 0$. Instead, it solves $\tilde{\pi}_{cc} = \tilde{\pi}_{dc}$, which yields

$$\sigma_{dc} = \frac{(2\delta - 1)\, p_{dd} + (1-\delta)\, p_{dc} + (1-\delta)\, p_{cd} - p_{cc}}{2\delta\,(p_{dd} - p_{cc})}. \tag{24}$$

Substituting it into $\tilde{\pi}_{dc} = 0$, and solving for δ yields $\delta > \delta_{\mathrm{BOS}}$. □

Table 9: Testing the Markov assumption: The weights of the memory-1 and memory-2 histories

	Treatment parameters			Memory-1 histories				Memory-2 histories $(t-1)\times(t-2)$										
Tr	b	a	δ	(c,c)	(d,c)	(c,d)	(d,d)	$cc\times cc$	$cc\times cd$	$cc\times dc$	$cd\times cc$	$cd\times dc$	$cd\times dd$	$dc\times cc$	$dc\times cd$	$dc\times dd$	$dd\times cd$	$dd\times dc$
Blonski et al. (2011)																		
6	1.43	1.29	0.5															
3	2.5	1.5	0.5	1.005** (0.028)	0.163 (0.079)	0.042 (0.063)	0.011 (0.011)	−0.153 (0.168)	0 (0)	0 (0)	0 (0)	0.061 (0.316)	0 (0.072)	0 (0)	−0.172 (0.1)	0.791 (0.41)	−0.014 (0.031)	0.071 (0.074)
2	1.25	1.12	0.75	0.984** (0.017)	0.234 (0.131)	0.144 (0.108)	0.003* (0.002)	0 (0)	0 (0)	0 (0)	0 (0)	0.434 (0.269)	−0.154 (0.111)	0 (0)	0.134 (0.282)	0.241 (0.306)	−0.013 (0.013)	0.001 (0.006)
7	1.43	1.29	0.75	0.983** (0.019)	0.313** (0.08)	0.139* (0.063)	0.006 (0.005)	−0.028 (0.033)	0 (0)	0 (0)	0 (0)	0.519* (0.174)	0.261 (0.373)	0 (0)	0.248 (0.196)	0.135 (0.364)	−0.006 (0.015)	0.002 (0.008)
4	2.5	1.5	0.75	0.832** (0.062)	0.182** (0.039)	0.106* (0.036)	0.043** (0.006)	0.061 (0.063)	−0.007 (0.082)	0.098 (0.07)	−0.134 (0.047)	0.386 (0.166)	0.041 (0.1)	−0.137 (0.081)	0.025 (0.142)	0.048 (0.087)	−0.085** (0.019)	0.009 (0.022)
8	1.43	1.29	0.88	0.905** (0.06)	0.211** (0.048)	0.118** (0.034)	0.033** (0.006)	0.03 (0.059)	−0.014 (0.114)	−0.041 (0.156)	0.407 (0.204)	0.723 (0.161)	0.068 (0.159)	−0.189 (0.072)	0.646 (0.17)	−0.082 (0.109)	−0.054* (0.019)	0.005 (0.012)
5	2.5	1.5	0.88	0.912** (0.046)	0.292** (0.049)	0.169** (0.035)	0.049** (0.007)	0.011 (0.043)	−0.096 (0.117)	0.053 (0.055)	−0.17 (0.079)	0.358* (0.117)	0.011 (0.097)	−0.34 (0.143)	0.256 (0.121)	−0.041 (0.088)	−0.008 (0.034)	0.071* (0.031)
9	2.4	1.8	0.75	0.923** (0.027)	0.207** (0.046)	0.103* (0.04)	0.029* (0.01)	−0.003 (0.032)	−0.489 (0.409)	−0.443 (0.35)	−0.078 (0.108)	0.487* (0.151)	0.193 (0.192)	−0.112 (0.116)	0.148 (0.163)	−0.078 (0.121)	−0.032 (0.025)	0.055 (0.035)
1	3	2	0.75	0.876** (0.063)	0.221** (0.037)	0.303** (0.047)	0.043** (0.008)	−0.077 (0.077)	−0.067 (0.104)	−0.019 (0.135)	−0.232 (0.126)	0.101 (0.102)	−0.271* (0.087)	0.277 (0.196)	0.468** (0.095)	−0.109 (0.071)	−0.032 (0.031)	0.014 (0.02)
10	4.67	3	0.75	0.885** (0.032)	0.301** (0.055)	0.09* (0.04)	0.051** (0.012)	−0.01 (0.038)	−0.341 (0.263)	−0.256 (0.242)	−0.121* (0.043)	0.399* (0.112)	0.139 (0.123)	−0.169 (0.08)	0.025 (0.111)	0.083 (0.135)	−0.041 (0.032)	0.093 (0.042)
Dal Bo and Fréchette (2011)																		
1	2.92	1.54	0.5	0.804** (0.036)	0.315** (0.056)	0.226* (0.061)	0.032** (0.004)	0 (0)	0 (0)	0 (0)	0 (0)	0.135 (0.119)	−0.231 (0.143)	0 (0)	0.416* (0.1)	0.128 (0.124)	−0.156* (0.048)	0.06 (0.06)
3	2.92	2.15	0.5	0.899** (0.058)	0.321** (0.031)	0.294** (0.04)	0.119** (0.009)	−0.072 (0.077)	−0.202 (0.105)	−0.091 (0.101)	−0.114 (0.204)	−0.167 (0.093)	−0.092 (0.082)	0.216 (0.199)	0.241 (0.106)	0.17 (0.07)	−0.065 (0.047)	0.033 (0.024)
2	2.92	1.54	0.75	0.937** (0.031)	0.307** (0.036)	0.339** (0.039)	0.034** (0.005)	−0.006 (0.032)	0 (0.032)	−0.179 (0.08)	−0.161 (0.129)	0.056 (0.079)	−0.163 (0.074)	0.432* (0.12)	0.437** (0.072)	−0.024 (0.077)	0.005 (0.03)	0.053 (0.026)
5	2.92	2.77	0.5	0.942** (0.011)	0.201** (0.024)	0.324** (0.031)	0.065** (0.01)	0 (0)	0 (0)	0 (0)	0 (0)	0.3 (0.133)	−0.138 (0.134)	0 (0)	0.455* (0.148)	0.273 (0.141)	−0.026 (0.034)	0.03 (0.016)
4	2.92	2.15	0.75	0.952** (0.012)	0.556** (0.048)	0.351** (0.054)	0.178** (0.021)	−0.008 (0.012)	−0.064* (0.025)	−0.059 (0.048)	−0.131 (0.101)	0.106 (0.093)	0.076 (0.139)	0.09 (0.106)	0.013 (0.096)	−0.224 (0.108)	−0.19** (0.03)	0.084 (0.059)
6	2.92	2.77	0.75	0.985** (0.005)	0.346** (0.079)	0.249** (0.073)	0.108** (0.027)	−0.006 (0.006)	−0.104 (0.113)	−0.003 (0.022)	−0.098 (0.105)	0.401 (0.155)	0.186 (0.163)	−0.086 (0.113)	0.232 (0.168)	0.244 (0.161)	−0.065 (0.043)	0.057 (0.069)
Duffy and Ochs (2009), "random rematching" treatment																		
	3	2	0.9	0.979** (0.014)	0.298** (0.042)	0.276** (0.043)	0.115** (0.014)	−0.014 (0.014)	−0.175* (0.06)	−0.123 (0.059)	0.064 (0.087)	0.235* (0.08)	0.021 (0.068)	0.212 (0.09)	0.149 (0.079)	0.074 (0.067)	−0.011 (0.033)	0.153** (0.042)
Fudenberg et al. (2012), "no-noise" treatment																		
6	5	4	0.88	0.944** (0.017)	0.502** (0.059)	0.417** (0.067)	0.166** (0.03)	0.011 (0.017)	−0.068 (0.067)	−0.086 (0.066)	0.062 (0.117)	0.128 (0.123)	0.077 (0.143)	−0.02 (0.109)	−0.067 (0.115)	0.006 (0.151)	−0.036 (0.068)	0.059 (0.056)

Note: ** indicates significance at $p < 10^{-4}$, * indicates significance at $p < 10^{-2}$