Bad Repetition*

Geoffroy de Clippel† and Kareen Rozen‡

November 2022

Abstract

While most experimental papers on repeated games explore the benefits of repeated interactions, we explore the often-overlooked potential for negative implications. We demonstrate this possibility in the lab with a standard random-termination protocol applied to a new, simple and easily interpretable stage game, capturing stylized aspects of bystander complacency.

*We are grateful to Pedro Dal Bó and David Levine for valuable suggestions, and to Sarah Conlisk, Jeongbin Kim, Pierre Lipton, Zeky Murra, Shreya Ramayya, Ian Tarr, Yinuo Zhang, and Jiaxiu Zhong for excellent research assistance. We also thank the seminar audience at the Stony Brook International Conference on Game Theory. Brown University IRB Protocol #1703001732. First version July 2020.

†Department of Economics, Brown University. Email address: declippel@brown.edu.

‡Department of Economics, Brown University. Email address: kareen_rozen@brown.edu.

1 Introduction

In a repeated game, players face the same strategic problem–the *stage game*–multiple times, with payoffs accruing each round. Folk theorems illustrate how the equilibrium set typically expands substantially beyond the mere repetition of a stage-game Nash equilibrium (NE). This is most often presented as great news: the common good is compatible with the pursuit of selfish interests. The Prisoners' Dilemma provides a classic, textbook example: cooperation can arise at equilibrium even though defecting is the dominant strategy. Friedman (1971), a classic reference for an early folk theorem, studies only feasible payoffs that strictly Pareto dominate some stage-game NE, and most experiments in repeated games concentrate on the beneficial nature of repeated interaction. But the set of equilibrium outcomes can expand even further: Fudenberg and Maskin (1986) shows that players' utility can be brought down close to their minimax level, which is oftentimes much lower than the stage-game equilibrium payoff (although not for the Prisoners' Dilemma, where the two coincide). Thus repeated interaction can also potentially be detrimental. Whether it is actually so is an important, yet so far overlooked, empirical question. We test in the lab a new, simple and easily interpretable stage game to shed light on this topic.

Imagine an aggressor (a bully) who wants to take resources from a disempowered victim. A bystander observes the interaction, and finds such exploitation repugnant. Will the bystander intervene? If there is no exogenous cost to intervention (e.g., risk of injury), then he would certainly intervene in a one-time interaction. Bystander complacency can potentially arise in equilibrium, however, if repeated encounters are likely. The possibility of the bully retaliating against the bystander can be a credible threat in the repeated game, even if it is costly to the bully. Thus, repetition opens up the possibility of equilibria with an endogenous cost of intervention. The mere fact that the bully has the possibility of punishing the bystander at a cost may allow him to implicitly coerce the bystander to be complacent with the theft. Of course, other bullying-free equilibria are possible too.

We designed a simple stage game (the 'Bystander game') reflecting the above considerations to experimentally investigate the possibility of bad repetition. If played only once, or more generally if repetition is unlikely, then backward induction (and NE) leads to a unique prediction: the bystander will intervene to stop the bully should he try to take money from the victim. Thus, in this game the unique equilibrium outcome of the stage game is both egalitarian and Pareto efficient (even sum-efficient). This game would be of limited interest if repeated interaction could only make participants weakly better off, as the stage-game Nash equilibrium is already Pareto efficient, and would thus be the anticipated outcome in every period. This outcome is even more salient because it is egalitarian. Notice, though, that new conventions, beyond the scope of simple Nash-reversion strategies used in Friedman (1971), become self-enforcing when the continuation probability is large enough. For instance, it can be sustainable for the bystander to go along with the bully's theft (on the equilibrium path), even if this is costly, due to fear of the bully punishing him in future rounds otherwise (off path). Outcomes in this equilibrium with theft see increased inequality accompanied by a decrease in aggregate payoffs. The bully may want such equilibria to be focal, leading to departure from an efficient and egalitarian outcome, but it is an empirical question whether this occurs.

Our experiment tests behavior for two continuation probabilities: $\delta = 75\%$, under which it is theoretically possible in our benchmark model for the bully to sustain transfers from the victim in every period, and $\delta = 10\%$. To preview our results, transfer attempts are frequent in both treatments, occurring approximately two-thirds of the time in the first round of a match. Since requesting a transfer is weakly dominant in the stage game, this is not particularly surprising, and we observe only a small and statistically insignificant amount of learning by bullies: for instance, comparing the first five matches to the last five, there is a slight increase (decrease) in transfer requests in the later matches when δ is high (low). On the other hand, there does appear to be some learning by bystanders in the very early matches of both treatments, with behavior settling fairly quickly to significantly different levels as a function of δ. Bystanders are rarely complacent when $\delta = 10\%$, but we observe a 42% complacency rate when $\delta = 75\%$. We observe a decrease in total surplus and an increase in inequality relative to the egalitarian baseline of the game (with the bully the clear benefactor). With the rate of transfer requests relatively stable and similar across treatments, these differences are attributable to the difference in complacency rates as δ varies. Our observations are broadly consistent with the fact that successful transfers can occur in equilibrium when $\delta = 75\%$, but are not expected when $\delta = 10\%$. When $\delta = 75\%$, bystanders who intervene to protect the victim in the first round face an approximately 1/3 chance of punitive action by the bully in the second round (we observe that this rate varies significantly by gender, with women less likely to react punitively).

We discuss related literature below, before formalizing the game and laying out the experimental procedure in Section 2. Section 3 develops the theoretical benchmark. Section 4 discusses our experimental results.

Related Literature

There is a substantial experimental literature on infinitely-repeated games, which is typically implemented in the lab using a random termination rule, as in the present paper. These studies are often interested in when and how better outcomes prevail in the repeated Prisoners' Dilemma; see Table 3 in Dal Bó and Fréchette (2018) for a list of references. Though the literature on infinitely repeated games has most often focused on the prisoners' dilemma, other stage games have been tested as well. Even then, as discussed in Dal Bó and Fréchette (2018, Section 7), the focus has been on the potential benefits of cooperation, e.g. increasing cooperation by adding options to punish, sustaining tacit collusion, promoting trusting behavior, and contributing more in public good games. In typical experiments with the trust game and public good games, the minimax payoff levels and the Nash equilibrium payoffs of the stage game still coincide. They need not coincide in the quadratic-payoff structures specified in industrial organization experiments, which consider large action spaces and vary parameters determining whether actions are complements or substitutes, to analyze which is more conducive to collusion. Though the vast majority of such studies consider only finite repetition, Mermer et al. (2021) considers infinite repetition (the distribution of match payoffs there is highly skewed above Nash payoffs).

The bully's option to penalize the bystander is an instance of costly punishment (paying a cost to decrease another player's payoff). Important early works on this topic include Ostrom et al. (1992) and Fehr and Gächter (2000, 2002). An important difference is that the punishment in those works is a means to potentially sustain *better* outcomes, whether in a public goods game or the prisoner's dilemma. In our work, costly punishment only raises the possibility of bad equilibrium outcomes instead, and is wielded by the bully.

The stage game we study may call to mind classic entry-deterrence games, to the extent that a player may wield a punishment action to obtain a better payoff under repetition. One could think of the bully as being a challenger who decides whether to stay out of a market (leaving the victim alone) or enter it (attempting to steal from the victim). Under that reinterpretation, the bystander would be the incumbent being challenged in that market. The bystander intervening would amount to fighting the entrant. Otherwise, she is accommodating it. The analogy is imperfect though, as a third firm (the victim) is involved and, more importantly, intervention only brings payoff back to their baseline level when the challenger stays out, instead of the incumbent having an action to punish the challenger for entering the market. In our game, it is the bully (entrant) who can punish the bystander (incumbent). The strategic considerations are thus quite different in entry games and ours. Another main difference is that the experimental literature on entry games, such as Jung et al. (1994), has focused on finitely-repeated games and the role of reputation under incomplete information.

The possibility of perverse outcomes has been considered in more elaborate repeated interactions with incomplete information. While at first conjectured to always be beneficial, Ely and Välimäki (2003) point out that the reputational concerns of a long-run privately-informed player who repeatedly interacts with short-lived players can undermine his commitment power, erasing any surplus; see also Ely et al. (2008). However, Grosskopf and Sarin (2010) experimentally test situations where reputation can be either good or bad, and observe that reputation is rarely harmful, though benefits are weaker than predicted by the theory.

Fearon (1995) provides another interesting instance where dynamics may create inefficiency, but again for different reasons than those we explore. He explains inefficient wars as the outcome of a dynamic interaction with stochastically changing fundamentals. Two countries have claims over a resource providing a regular profit flow. Pareto efficiency is obtained when countries peacefully share profits in each period. But of course the countries' outside option (waging war in the hope of acquiring and keeping the resource forever after) impacts each period's profit split. Fearon observes that determining ownership through war can be unavoidable at equilibrium when countries' relative strengths (e.g., winning probability) vary over time. A country expecting to grow weaker may prefer war today, as its profit share will otherwise diminish in the future. The other country would prefer giving higher profit shares in the future to avoid war today, but such promises are not credible; see Tingley (2011) for an experimental test of this idea, and McBride and Skaperdas (2014) for a setting where current conflict impacts future strategic position. Of course, dynamic bargaining in the shadow of conflict is subtler than the mere repetition of a stage game,[1] and the possible necessity of preventive war is not the consequence of folk theorems.

The general combination of coercion, complacency and detrimental outcomes prompted by fear of future punishments, is somewhat reminiscent of Padró i Miquel (2007). His theoretical model shows how the presence of ethnic identities and the absence of institutionalized succession may allow for large, inefficient rent extraction by the rulers. In this case, supporters are disciplined by a fear of facing an equally inefficient ruler who favors another group. The logic of the argument, however, is unrelated to folk theorems in repeated games, as Padró i Miquel characterizes the unique Markov-perfect equilibrium of his model (a stochastic game). By contrast, we study the empirical prevalence of detrimental outcomes that are supported by non-stationary strategies in a repeated game, as is standard in folk theorems.

2 Bystander Game and Experimental Procedure

We start by describing the bystander game we tested, and its loose interpretation, before presenting its neutral implementation in the lab.

The bystander game comprises a bully, a bystander, and a helpless victim who has no strategic choices. Each player starts with an initial utility of 100. The bully can either do nothing, harm the bystander (at a cost of 2 utils to himself and 50 utils to the bystander), or attempt to steal from the victim. For theft to succeed, he needs the bystander to choose complacency over protecting the victim. Successful theft leads to a social loss: (i) the bystander incurs a loss of 10 utils, and (ii) the bully gains only 40 utils while the victim loses 50. Bullying fails when the bystander intervenes, leaving payoffs unchanged (100 each). Of course, one could imagine many alternative formulations, but our purpose is to offer a stylized model capturing some of the general forces at play, not to offer a precise model of a specific situation.

While we use suggestive language to describe the bystander game above, options and player names were described to experimental subjects in neutral language, without using terms like bully or theft. The extensive form of the game is presented in Figure 1, with the neutral language presented to subjects in the main labels and their interpretations in parentheses. At the beginning of the experiment, each subject is

[1] First, war offers an outside options that stops future bargaining altogether. Second, countries' relative strengths must change over time, which in turn changes payoffs.

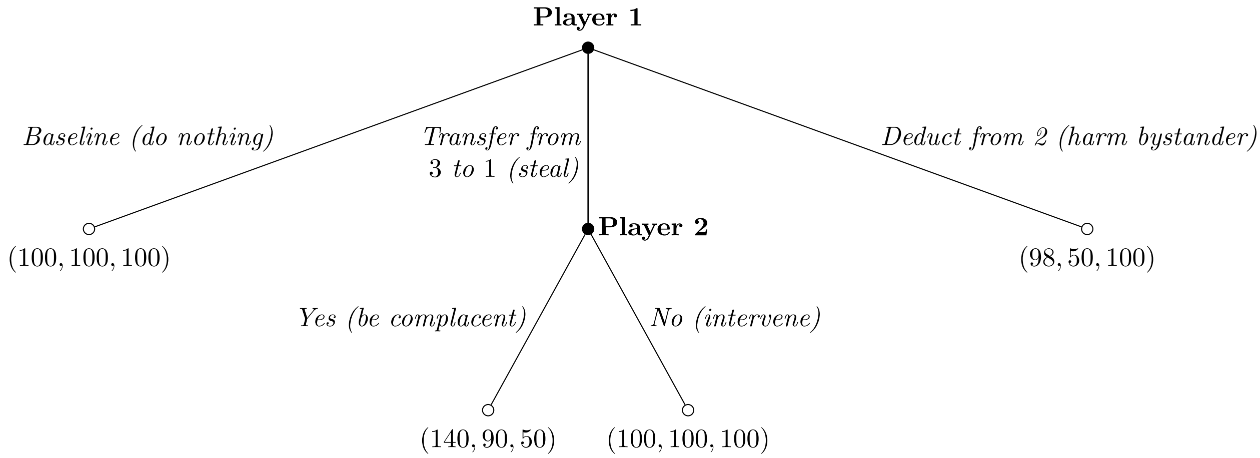

Figure 1: Bystander game tree. The first payoff is for Player 1 (the "bully"), the second is for Player 2 (the "bystander"), and the third is for Player 3 (the "victim"). The main labels in the figure reflect the neutral action and player labels used during the experiment; the suggestive interpretations in parentheses were not used.

randomly assigned to one of three possible roles: Player 1, Player 2, or Player 3 (corresponding to the bully, bystander, and victim, respectively). The assigned role is held for the duration of the session. Player 1's options are called *Baseline*, *Deduct from 2*, and *Transfer from 3 to 1*; these correspond to doing nothing, attacking the bystander, and attempting theft, respectively. Player 2's options are called *Yes* and *No*, which correspond to agreeing to theft and stopping it, respectively. We provide subjects with a table describing how many experimental points each player gets from each of the four possible combinations of choices in the stage game (i.e., *Baseline*, *Deduct from 2*, *Transfer from 3 to 1/Yes*, *Transfer from 3 to 1/No*), using the numbers shown in Figure 1. The instructions clarify that Player 1 moves first.

Each session has 15 matches. At the start of every match, each subject is matched into a three-person group, constructed by randomly drawing one Player 1, one Player 2, and one Player 3. Players stay in their three-person group for the duration of the match. Within a match, each group plays exactly one supergame. A supergame may consist of multiple rounds. Each round was played sequentially, in accordance with the extensive form in Figure 1; and Player 2, if required to make a decision due to a transfer request, was expressly informed that a transfer request had been made by Player 1. At the end of each round, subjects learn whether their supergame ends or continues to the next round. The probability of another round is held constant throughout a session, at either 75% or 10%. A subject's payoff in a match is the sum of payoff(s) accrued in the round(s) of that match. In each round of the supergame, group members play the above bystander game. At the end of each round, the three players in each group are told the choices made in that round, and the resulting payoffs, within their group. At the end of the experiment, each subject is paid their dollar payoff from one randomly chosen match in addition to the $8 show-up fee. Experimental points are converted to dollars at the rate of $0.05 per EP. Player identities remain anonymous. Subjects are not told what roles other subjects play, or the payoffs others receive.

All sessions were conducted at the Brown University Social Sciences Experimental Laboratory, which is equipped with sunken terminals and vertical privacy panels between desks. Subjects could participate in at most one session. Subjects were recruited via the BUSSEL website. A total of 225 subjects participated in the sessions, which took place in 2017 and 2018.[2] There were 120 subjects for $\delta = 75\%$ and 105 subjects for $\delta = 10\%$. The instructions, available in the Online Appendix, were provided in paper format at the start of the experiment and read out loud by the experimenter. The experiment was programmed using the z-Tree software (Fischbacher, 2007). Prior to starting the first match, subjects were asked to complete a comprehension quiz through the z-Tree interface. After the final match, subjects were given an optional exit survey. Subjects were paid in cash before leaving the laboratory.

3 Theoretical Benchmark

As is common in the experimental literature on repeated games, we conduct our analysis while assuming that players maximize their own monetary payoffs. We discuss the possibility and implications of more complex preferences further below. We will formulate our hypothesis after establishing four observations. The proof of the first two are left to the reader.

[2] A referee inquired as to why the study was not pre-registered. This was not an active decision, and reflects the time period during which the experiment took place. As noted by Abrams et al. (2020) (pertaining to data from 2017 through part of 2019): "...perhaps more telling is that only 8% of the working papers in the NEP report on experimental economics are registered. Interestingly, none of the RCTs published in the premier field journal Experimental Economics are registered. These RCTs consist entirely of lab experiments, suggesting that the norm is such that researchers do not register lab experiments." We also note that (a) the experimental design is driven by a precise theoretical question, with the predictions in our main hypothesis guided by Observations 1 through 4 below, and (b) we follow the well-established experimental literature on infinitely-repeated games in mostly focusing on first-round choices.

Observation 1 *Each player getting* 100 *EPs is the unique SPE outcome (even the unique NE outcome) of the bystander game played once. This outcome is egalitarian, Pareto efficient, and even sum efficient.*

Observation 2 *Whatever the continuation probability, each player getting* 100 *EPs in each round is the unique SPE outcome (even the unique NE outcome) that weakly Pareto dominates the stage-game equilibrium.*[3]

Observation 3 *Each player getting* 100 *EPs in each round is the unique SPE outcome (even the unique NE outcome) when* $\delta = 10\%$.

Proof. Indeed, successful transfers cannot be sustained in that case: even if Player 1 were to impose the harshest possible punishment on Player 2 by choosing *Deduct from 2* forever after (which is not incentive compatible for Player 1), deviating to *No* nets $10 - 40\frac{\delta}{1-\delta}$ for Player 2, which is strictly positive if $\delta = 10\%$. ■

Intuitively, when considering SPE of the repeated game, Player 2's disutility from selecting *Yes* in response to a transfer request makes it more difficult to sustain a non-egalitarian outcome. For this to be possible, δ must be sufficiently high.

Observation 4 *The equilibrium set is much larger when* $\delta = 75\%$. *In particular, Player 1's average discounted payoff can range at equilibrium from 100 to 140 (the upper bound corresponding to successful transfers in every period), while Player 3's average discounted payoff can range at equilibrium from 50 to 100.*

Proof. Consider the following strategies:

> **Transfer Requests and Retaliation** (**TRR**): Player 1 chooses *Transfer from 3 to 1*, except for choosing *Deduct from 2* once immediately after any round where Player 2 was supposed (according to the strategy below) to

[3] This remains true when Pareto comparisons apply only to Players 1 and 2, who act in the game. Therefore the standard logic of Nash-reverting punishments would still not sustain any outcomes beyond repetition of the stage NE in a variant of the game that excludes Player 3. We note that one could create a two-player variant where sum inefficiency also arises under transfers (e.g., having the payoffs $(110, 80)$ instead of $(150, 90)$). The equilibrium logic we highlight remains unchanged in such a variant, and we would expect similar qualitative results. One benefit of having a third player is the natural interpretation of the game in terms of bystander intervention. Moreover, if pro-social preferences play any role here at all (the literature suggests their impact on behavior is limited in repeated games when cooperation is sustainable as part of a selfish equilibrium), then baseline payoffs are only more appealing to all players. Thus, seeing successful transfers, with outcomes away from the $(100, 100, 100)$ equal split, is only more surprising and further evidence towards the repeated-game considerations we put forth.

say *Yes* but says *No* instead;

Complacency for Fear of Retaliation (**CFR**): Player 2 chooses *Yes* unless Player 1 has ever failed to deduct him immediately after choosing *No* (in which case he continues choosing *No*.)

The only non-trivial incentive conditions to check, which depend on the probability of continuation, are for Player 2 to choose *Yes* when he is expected to do so, and for Player 1 to choose *Deduct from 2* following a deviation to *No* by Player 2. The former means Player 2 must be willing to lose 10 EP in the current round, to avoid losing 40 EP if there is a next round: $-10 + 40\delta > 0$, or $\delta > 1/4$. The latter means Player 1 must be willing to lose 2 EP in the current round, to gain 40 EP in all future rounds: $-2 + 40\frac{\delta}{1-\delta} > 0$, or $\delta > 1/21$. These conditions are satisfied when $\delta = 75\%$. Given that Player 1 can also guarantee herself a payoff of 100 EP's in each period (e.g., by picking *Baseline* in each round), which of course remains an equilibrium outcome, we conclude that the range of discounted average payoffs is now the interval $[100, 140]$ for Player 1. The equilibrium with succesful transfers gives Player 3 his minimal average discounted payoff of 50. Alternatively, repeating the the stage-game equilibrium gives him his maximal average discounted payoff of 100, which shows that the whole range of payoffs, $[50, 100]$, is possible at equilibrium. ■

The existence of an intuitive equilibrium where Player 2 is complicit due to fear of retaliation from Player 1, makes it reasonable to expect some subjects in the role of Player 1 to receive a discounted average payoff that is strictly larger than in the low-δ treatment, at the detriment of Player 3. Of course, this goes hand-in-hand with greater inequality and less sum-efficiency. We can thus formulate the following testable hypothesis:

Hypothesis. *Complacency and successful transfers should be rare when δ is low, but significantly more frequent when δ is large. Thus, in the latter case we expect an increase in Player 1's payoff at the expense of the other players, greater inequality and lower total surplus.*

4 Experimental Results

There were 525 supergames played in the 10% treatment, with the mean number of rounds being 1.13 and the maximum being 2. For the 75% treatment, there were 600

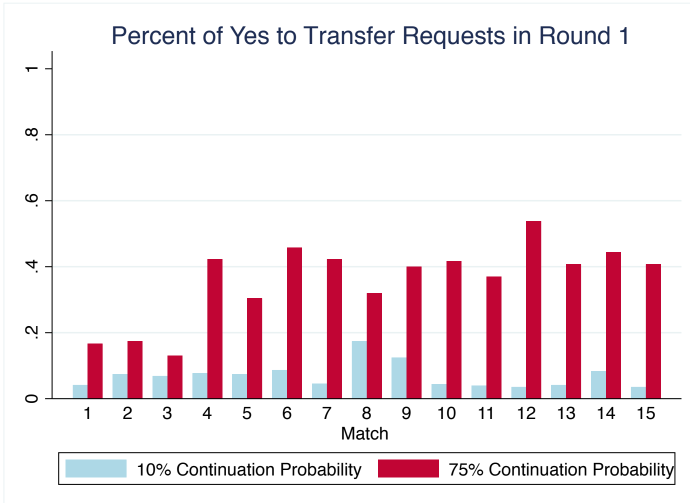

Figure 2: Response to transfer requests in Round 1, per match.

supergames played, with the mean number of rounds being 3.58 and the maximum being 19. The cleanest comparison between the treatments occurs when comparing the first round of matches: there are few second rounds in the 10% treatment, and later rounds in the 75% treatment may be polluted with dynamic considerations within the game.

A. Checking the Hypothesis

We start by looking at complacency rates. Figure 2 offers a match-by-match look at Round-1 choices of Player 2 following transfer requests. Broadly speaking, we see that bystanders agree to about 40% of transfer requests in the 75% treatment, compared to well under 10% of them in the 10% treatment. However, there appears to be some learning in the first few matches, as we anticipated. We tested 15 matches to account for this possibility and be able to focus on the last ten. Henceforth, we drop Matches 1-5 when performing statistical tests.

We use logistic regression to estimate the probability of *Yes* in Round 1 of Matches 6-15, conditional on a transfer request (using heteroskedasticity-robust standard errors clustered by participant). The estimates appear in Table 1. The 0.419 probability of agreeing to a transfer in Round 1 of the 75% treatment is significantly different from zero (p-value 0.0000), and is significantly larger than the 0.070 probability in the 10% treatment (p-value 0.0000). Though small, the latter is statistically different from zero at the 5% level (p-value 0.0248).[4] Of course, the practical importance of complacency rates hinges on the prevalence of transfer requests themselves. Without conditioning on a transfer request, the overall estimated percentage of successful transfers in Round 1 is 4.9% when $\delta = 10\%$ and 27.0% when $\delta = 75\%$. These are statistically different at all standard levels of significance, as hypothesized (p-value 0.0000). A more detailed analysis of Player 1 choices can be found in the next subsection; as will be seen there, the majority of Player 1 choices in Round 1 are transfer requests, and the percentage of such requests is very similar across the treatments. Thus the difference in the rate of successful transfers is driven entirely by the difference in complacency rates.

	$\delta = 10\%$	$\delta = 75\%$
Complacency Rate	0.070* (0.031)	0.419*** (0.075)
Observations	244	258

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Estimated complacency rates (Player 2 picking *Yes* following a transfer request) in Round 1, Matches 6 to 15. Robust standard errors clustered by participant.

With an increase of more than twenty percentage points in successful transfers when $\delta = 75\%$, one expects a surplus loss of 4 EPs compared to the baseline, along with increased inequality (8 extra EPs for Player 1, a loss of 2 EPs for Player 2 and a loss of 10 EPs for Player 3). Round 1 payoffs confirm this overall trend. Average payoffs over Matches 6-15 for Players 1, 2 and 3 are 101.82, 96.37, and 97.57 EP, respectively in the $\delta = 10\%$ treatment; while they are 110.76, 96.30, and 86.50 EP, respectively, in the $\delta = 75\%$ treatment. As hypothesized, there is a decrease in total surplus ($p = 0.0000$, Wilcoxon rank-sum test) and a clear increase in inequality due to a significant transfer of EPs from 3 to 1. Magnitudes are a bit different from what the theory predicts. In particular, average payoffs when $\delta = 10\%$ are somewhat different from the equilibrium prediction of 100 EPs each. The primary cause is the presence of a small fraction of deductions and successful transfers in that treatment. We come back to these observations as we explore players' choices in greater detail in the next two subsections. Small variations in magnitude such as these are to be expected given that our benchmark analysis assumes that participants only care about maximizing

[4]One can think of many explanations why. Perhaps subjects in the role of player 2 see *Deduct* as part of an equilibrium because they misperceive payoffs and continuation probabilities. They may also be influenceable (experiencing a cost when not doing as requested), or may fear that their matched subject has an intrinsic propensity to retaliate after a denial (intention-based preference).

their monetary payoffs (similarly, a non-negligible fraction of participants cooperate in prisoner's dilemma experiments with low continuation probabilities).

B. Taking a Closer Look at Choices

Table 2 below provides further information on Player 1's choices in the first two rounds of Matches 6 to 15.[5] We first focus on Round 1. Remember that Player 2 should always pick *No* in our benchmark analysis when $\delta = 10\%$. Given this, it is dominated for Player 1 to pick *Deduct* when $\delta = 10\%$. Still, 6.3% of subjects did pick it (p-value 0.0134). We see a 2% probability of first-round deductions when $\delta = 75\%$ (p-value 0.0363). Picking *Deduct* is part of a Pareto-dominated SPE when $\delta = 75\%$: for example, expecting to play the (TRR,CFR) equilibrium after *Deduct* in Round 1, and the (*Baseline*,*No*) stage-game equilibrium forever after Round 1 otherwise.

	Round 1		Round 2 After No		Round 2 After Yes
	$\delta = 10\%$	$\delta = 75\%$	$\delta = 10\%$	$\delta = 75\%$	$\delta = 75\%$
Baseline	0.240***	0.335***	0.031	0.058	0.213**
	(0.063)	(0.062)	(0.032)	(0.031)	(0.075)
Deduct	0.063*	0.020*	0.250**	0.337***	0.000
	(0.025)	(0.010)	(0.089)	(0.091)	(–)
Transfer	0.697***	0.645***	0.719***	0.606***	0.787***
	(0.065)	(0.063)	(0.092)	(0.095)	(0.075)
Observations	350	400	32	104	75

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Estimated choice probabilities for Player 1 (and robust standard errors, clustered by participant) in Round 1, Matches 6-15; and in Round 2, conditional on reply to Round-1 transfer request (excluding $\delta = 10\%$ after *Yes*, due to only 3 observations).

The theory does not make an irrefutable prediction about which of *Baseline* or *Transfer from 3 to 1* will be the choice of Player 1 in the first round. Picking the latter is weakly dominant in our benchmark analysis when $\delta = 10\%$, but picking *Baseline* while expecting player 2 to refuse transfer requests is also in equilibrium. And *Baseline* is weakly dominant instead for subjects who dislike inefficiency or inequality. No strategy is weakly dominant when $\delta = 75\%$, but the equilibrium set does expand,

[5]The conditional behaviors discussed in this section are the most relevant for identifying behavior, but the interested reader may also be curious about how common the different outcomes themselves are in later rounds. We include in the Online Appendix a table showing the percentages of each type of outcome per round when pooling all Matches 6-15 in the $\delta = 75\%$ treatment in which that round occurred.

and hence both *Baseline* and *Transfer from 3 to 1* can occur. Whether a significant number of transfers occurs in each treatment is thus an empirical question. We find it is more likely than not that a subject in the role of Player 1 requests a transfer: the estimates from Table 2 are both significantly larger than 1/2 (p-values of 0.0217 for $\delta = 75\%$ and 0.0024 for $\delta = 10\%$). Moreover, there is no statistically significant difference found when testing the joint null hypothesis that the Round-1 choices of Player 1 do not differ across treatments (p-value 0.1825).[6]

In our benchmark analysis with $\delta = 75\%$, the potential for deductions in histories where some transfers were denied is critical to make successful transfers possible. Table 2 confirms that the use of deductions when $\delta = 75\%$ is highly contingent on whether the past transfer request was granted or denied. Subjects in the role of Player 2 were complacent in the first-round of 75 supergames (focusing on Matches 6-15), and *Deduct from 2* was never used against those subjects in Round 2. By contrast, Player 2 is deducted in 33.7% of cases following a *No* response, which is statistically different from zero (p-value 0.0002).

Deduct from 2 remains a dominated action in the second round when $\delta = 10\%$. Yet 25% of subjects pick it after being denied a transfer in the first round (see Table 2). There are only 32 instances where a second round occurs after a *No* when $\delta = 10\%$, so the numbers are small, but still different from zero (p-value 0.0047). Though 25% is certainly smaller than the estimated 33.7% who deduct when $\delta = 75\%$, the difference is not statistically significant (p-value 0.4951) given rather large standard errors. These are due to the small number of second rounds in the $\delta = 10\%$ case, and potentially due to high variance in behavior in the $\delta = 75\%$ case (which we conjecture is attributable to gender differences, as discussed further below). We cannot provide estimated probabilities of Player 1's Round-2 choices after *Yes* when $\delta = 10\%$, as there are only three such instances in our dataset. It is safe though to conjecture the deduction rate would be close to zero.

Why do we observe deductions following denied transfer requests when $\delta = 10\%$? We cannot exclude the possibility that subjects misperceive payoffs and/or continuation probabilities, prompting them to follow strategies that form an equilibrium in

[6] We note that we do not see significant evidence of learning by Player 1 in either treatment. For instance, if one compares the rate of transfer requests in the first five matches versus the last five, there is about 7.5% less requests at the beginning when $\delta = 75\%$ and about 2% more requests at the beginning when $\delta = 10\%$. Neither of these differences is statistically significant, even at the 10% level (the smallest p-value, which is for $\delta = 75\%$, is 0.161).

our benchmark analysis only for larger δ's. Another possibility is that some subjects have intention-based preferences, feeling unkind towards a Player 2 who denies a profitable transfer request. Better understanding Player 1's motives is an interesting topic to explore in future work. Earlier papers on the prisoner's dilemma suggest various directions to pursue. In the spirit of Dal Bó (2005), one could contrast behavior when the bystander game is played exactly four times, instead of being repeated with a 75% continuation probability. One could also examine possible correlations between behavior in the repeated games and elicited individual characteristics of other-regarding preferences (ORPs). Interestingly, the consensus from past studies on the prisoner's dilemma is that there is no robust evidence of individual characteristics (including ORPs) having a systematic effect on cooperation when δ is high enough to support cooperation as an equilibrium; see Dal Bó and Fréchette (2018, Result 7). By contrast, cooperation can be significantly correlated with ORPs when it is not part of an equilibrium because δ is low (Dreber et al., 2014). Given this, one might conjecture that deductions in the bystander game are driven primarily by strategic considerations when δ is high, and primarily by ORPs when δ is low. Whatever progress is made regarding Player 1's motives, our conclusion that repetition can be detrimental is only reinforced, given that ORPs provide an additional rationale here for bad repetitions. What matters for complacency is Player 2's fear of a likely retaliation if he turns down a transfer request, which occurs only when the continuation probability is high and does not hinge on Player 1's motive for retaliating.

Finally, though our study was not designed to isolate the effect of gender on behavior, we do note some gender differences that may be worth examining in future studies to evaluate their robustness.[7] These differences mainly appear in the role of Player 1; we find no significant gender differences in first-round choices in the bystander role (p-values 0.2265 for $\delta = 10\%$ and 0.9929 for $\delta = 75\%$). As seen from Round-1 behavior in Table 3, even though men and women act nearly identically as Player 1 when $\delta = 10\%$ (p-value 0.9307), behavior differs when $\delta = 75\%$ (p-value 0.0351). While men are much more likely to request *Transfer from 3 to 1* in the first round than select *Baseline*, women are more evenly split between these options (see Table 3). Moreover, behavior differs following a Round-1 denial of a transfer request.

[7] The game studied here, with the possibility of equilibrium strategies beyond Nash reversion, is somewhat different than those considered in the survey of Dal Bó and Fréchette (2018). They note that "there is no robust evidence" that gender has a "systematic effect on cooperation in infinitely repeated games in which cooperation can be supported in equilibrium."

The rightmost columns of Table 3 show that behaviors are nearly reversed, with the difference in choices statistically significant at all levels (p-value 0.0000). Men choose *Deduct from 2* about 73% of the time following a *No*, and repeat the transfer request in most other cases, while women repeat the transfer request over 75% of the time, and deduct in only 16% of cases. By contrast, we find no difference in behavior by gender following a *Yes* from Player 2 (p-value 0.8117). We note that among subjects in the roles of Player 1 and 2, 50% and 56.6% of those who provided their gender identified as female in the $\delta = 10\%$ and $\delta = 75\%$ treatments, respectively.

	Round 1 $\delta = 10\%$		Round 1 $\delta = 75\%$		Round 2 After No $\delta = 75\%$	
	Men	Women	Men	Women	Men	Women
Baseline	0.230^{**} (0.084)	0.253^{**} (0.096)	0.179^{*} (0.074)	0.474^{***} (0.087)	0.033 (0.034)	0.081 (0.051)
Deduct	0.070 (0.039)	0.053 (0.028)	0.021 (0.011)	0.022 (0.015)	0.733^{***} (0.088)	0.161^{*} (0.077)
Transfer	0.700^{***} (0.084)	0.693^{***} (0.102)	0.800^{***} (0.075)	0.504^{***} (0.088)	0.233^{***} (0.071)	0.758^{***} (0.093)
Observations	200	150	140	230	30	62

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Estimates of Player 1's choices by reported gender, with robust standard errors clustered by participant.

5 Concluding Remarks

Despite its size, the experimental literature on repeated games focus on a relatively small set of classic stage games, where the main question is the benefit of repeated interaction. We find it worthwhile to widen the scope of the analysis: repetition opens the door to negative conventions that do not Pareto-dominate repetition of one-stage Nash (which of course necessitate the consideration of games where minimax utilities are distinct from NE payoffs). We establish the possibility of such bad repetition in the lab, by applying a standard random-termination protocol to a new stage game.

Our stage game admits weakly dominant strategies, a unique SPE outcome, and a unique NE outcome. The equilibrium outcome is both egalitarian and Pareto efficient. Yet other SPE outcomes, all of which are favorable to Player 1, arise when the continuation probability is sufficiently high, because the possibility to punish Player 2 induces complacency at equilibrium. It remains an open question whether there are situations where high continuation probabilities result in Pareto-inferior outcomes. More generally, work is needed to understand both the positive and negative implications of repeated interaction, and both how and why such conventions crystallize in practice. For instance, investigating the potential role of risk dominance in selecting conventions would be a natural step in this direction (see Result 3 in Dal Bó and Fréchette (2018), and the discussion around it, for results and references on the role of risk dominance as a determinant of cooperation in the infinitely-repeated prisoners' dilemma).

Our game captures some stylized features of bullying and bystander complacency. In particular, it focuses on aspects related to repetition, as we suspect that the likelihood of repetition may play a role in that harmful phenomenon. While definitions of bullying differ, a commonly used one states that "an individual is a victim of bullying when he or she is exposed repeatedly over time to negative actions by one or more individuals and is unable to defend him or herself" (Hamburger et al., 2011). However, our game does not incorporate other important aspects that could be studied, such as how roles as bully or victim arise (roles are exogenously assigned in the experiment, and players are anonymous). Hence it should not be viewed as a full-fledged model of bullying.

References

Abrams, E., J. Libgober, and J. A. List (2020). Research registries: Facts, myths, and possible improvements. *NBER Working paper 27250*.

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

Dal Bó, P. and G. Fréchette (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature 56*(1), 60–114.

Dreber, A., D. Fudenberg, and D. Rand (2014). Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics. *Journal of Economic Behavior and Organization 98*(C), 41–55.

Ely, J., D. Fudenberg, and D. K. Levine (2008). When is reputation bad? *Games and Economic Behavior 63*(2), 498 – 526.

Ely, J. C. and J. Välimäki (2003). Bad reputation. *The Quarterly Journal of Economics 118*(3), 785–814.

Fearon, J. D. (1995). Rationalist explanations for war. *International Organization 49*(3), 379–414.

Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review 90*, 980–994.

Fehr, E. and S. Gächter (2002). Altruistic punishment in humans. *Nature 415*, 137–140.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies 38*(1), 1–12.

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

Grosskopf, B. and R. Sarin (2010). Is reputation good or bad? an experiment. *American Economic Review 100*(5), 2187–2204.

Hamburger, M., K. Basile, and A. Vivolo (2011). *Measuring Bullying Victimization, Perpetration, and Bystander Experiences: A Compendium of Assessment Tools*. Centers for Disease Control and Prevention, National Center for Injury Prevention and Control.

Jung, Y. J., J. H. Kagel, and D. Levin (1994). On the existence of predatory pricing: An experimental study of reputation and entry deterrence in the chain-store game. *The RAND Journal of Economics 25*(1), 72–93.

McBride, M. and S. Skaperdas (2014). Conflict, settlement, and the shadow of the future. *Journal of Economic Behavior and Organization 105*, 75 – 89.

Mermer, A. G., W. Muëller, and S. Suetens (2021). *Journal of Economic Behavior and Organization 188*, 1191–1205.

Ostrom, E., J. Walker, and R. Gardner (1992). Covenants with and without a sword: Self-governance is possible. *The American Political Science Review 86*, 404–417.

Padró i Miquel, G. (2007). The control of politicians in divided societies: The politics of fear. *The Review of Economic Studies 74*(4), 1259–1274.

Tingley, D. (2011). The dark side of the future: An experimental test of commitment problems in bargaining. *International Studies Quarterly 55*, 521 – 544.

Online Appendix for *Bad Repetition* by de Clippel and Rozen

Experiment instructions attached on the next pages. See below for table of the percentages of each type of outcome per round, when pooling all Matches 6-15 in which that round occurred, in the $\delta = 75\%$ treatment (# groups is the number of trios of Players 1, 2 and 3 that are observed playing the game in that round).

Round	Baseline	Deduct	Transfer/No	Transfer/Yes	# Groups
1	33.50%	2.00%	37.50%	27.00%	400
2	29.64%	12.86%	35.00%	22.50%	280
3	32.51%	5.91%	38.42%	23.15%	203
4	31.88%	15.22%	30.43%	22.46%	138
5	34.41%	9.68%	32.26%	23.66%	93
6	26.79%	17.86%	35.71%	19.64%	56
7	32.00%	18.00%	32.00%	18.00%	50
8	22.73%	9.09%	45.45%	22.73%	44
9	31.25%	15.62%	40.62%	12.50%	32
10	38.46%	7.69%	34.62%	19.23%	26
11	26.32%	10.53%	47.37%	15.79%	19
12	30.77%	0%	46.15%	23.08%	13
13	33.33%	16.67%	0%	50.00%	6
14	50.00%	0%	16.67%	33.33%	6
15	33.33%	33.33%	16.67%	16.67%	6
16	16.67%	16.67%	50.00%	16.67%	6
17	16.67%	33.33%	16.67%	33.33%	6
18	16.67%	33.33%	0%	50.00%	6
19	33.33%	0%	33.33%	33.33%	6

Welcome to this decision-making experiment!

Please silence and put away electronic devices.

Instructions

You will receive an $8 show-up fee, and will be able to earn more. The exact amount earned will depend on chance and choices made during the experiment.

The number of participants in this room is a multiple of 3, as there are 3 possible Player roles in this experiment: Player 1, Player 2 and Player 3. At the beginning of the experiment, the computer will randomly assign roles, making sure there are equally many Player 1's, Player 2's and Player 3's. The computer will inform you of your role, which will remain fixed throughout the experiment.

There are **15** rounds in this experiment. At the start of each round, the computer randomly places every participant into a three-person group. Each three-person group is constructed by drawing one Player 1, one Player 2, and one Player 3. Players stay in their three-person group for the duration of the round.

How does a round unfold?

Each round can consist of multiple stages, all of which are identical in structure. Once a stage is over, the computer throws a (simulated) 100-sided die and another stage occurs if a number lower than or equal to 75 comes up. Thus after every stage, there is a 75% chance that the round continues to another stage and a 25% chance that the round ends. Remember that everyone's player role and matched group is the same across all the stages of a round.

In each stage of a round, you will receive some experimental points (denoted EP) based on the decisions of players in your group. Your running total for a round is the sum of the points you received in the stages of that round.

What happens in each stage of a round?

The baseline payoff in every stage is 100 EP for each player; but within a group, the current stage payoffs can be modified from these baselines by Player 1's and 2's decisions as follows. Player 1 moves first and has three options:

○ *Baseline* means the baseline payoffs are received in this stage.

○ *Deduct from 2* means Player 1 pays 2 EP, and 50 EP are deducted from Player 2's baseline payoff in this stage.

○ *Transfer from 3 to 1* requires Player 2 to choose either *Yes* or *No*. If Player 2 picks *No*, baseline payoffs are received this stage. If Player 2 picks *Yes*, then both Player 1 and Player 2 pay 10 EP, and there is a transfer of 50 EP from Player 3 to Player 1 in this stage.

The following table summarizes the possible stage payoffs as a function of choices made:

	Player 1 stage payoff (EP)	Player 2 stage payoff (EP)	Player 3 stage payoff (EP)
Player 1: Baseline	100	100	100
Player 1: Deduct from 2	98	50	100
Player 1: Transfer from 3 to 1 **Player 2: Yes**	140	90	50
Player 1: Transfer from 3 to 1 **Player 2: No**	100	100	100

At the end of each stage, the stage's outcome will be displayed to the three players.

What happens at the end of the experiment?

Once all rounds have been completed, there will be a short and optional exit survey. Your participation is voluntary and does not affect your payoff.

At the end of the experiment, the computer randomly chooses one round for your payment. Your experimental-point total for that round will be converted to dollars at the rate of $0.05 per experimental point. You will be paid the dollar payoff from your selected round in addition to the $8 show-up fee. All player identities remain anonymous. No one will learn what role you played or what payoff you earned.

We are almost ready to start the experiment. Before doing so, there will be a short quiz to check your understanding of some key features of the experiment, as well as a chance to examine how the player decision screens work.

Welcome to this decision-making experiment!

Please silence and put away electronic devices.

<u>Instructions</u>

You will receive an $8 show-up fee, and will be able to earn more. The exact amount earned will depend on chance and choices made during the experiment.

The number of participants in this room is a multiple of 3, as there are 3 possible Player roles in this experiment: Player 1, Player 2 and Player 3. At the beginning of the experiment, the computer will randomly assign roles, making sure there are equally many Player 1's, Player 2's and Player 3's. The computer will inform you of your role, which will remain fixed throughout the experiment.

There are **15** rounds in this experiment. At the start of each round, the computer randomly places every participant into a three-person group. Each three-person group is constructed by drawing one Player 1, one Player 2, and one Player 3. Players stay in their three-person group for the duration of the round.

<u>How does a round unfold?</u>

Each round can consist of multiple stages, all of which are identical in structure. Once a stage is over, the computer throws a (simulated) 100-sided die and another stage occurs if a number lower than or equal to 10 comes up. Thus after every stage, there is a 10% chance that the round continues to another stage and a 90% chance that the round ends. Remember that everyone's player role and matched group is the same across all the stages of a round.

In each stage of a round, you will receive some experimental points (denoted EP) based on the decisions of players in your group. Your running total for a round is the sum of the points you received in the stages of that round.

<u>What happens in each stage of a round?</u>

The baseline payoff in every stage is 100 EP for each player; but within a group, the current stage payoffs can be modified from these baselines by Player 1's and 2's decisions as follows. Player 1 moves first and has three options:

○ *Baseline* means the baseline payoffs are received in this stage.

○ *Deduct from 2* means Player 1 pays 2 EP, and 50 EP are deducted from Player 2's baseline payoff in this stage.

○ *Transfer from 3 to 1* requires Player 2 to choose either *Yes* or *No*. If Player 2 picks *No*, baseline payoffs are received this stage. If Player 2 picks *Yes*, then both Player 1 and Player 2 pay 10 EP, and there is a transfer of 50 EP from Player 3 to Player 1 in this stage.

The following table summarizes the possible stage payoffs as a function of choices made:

	Player 1 stage payoff (EP)	Player 2 stage payoff (EP)	Player 3 stage payoff (EP)
Player 1: Baseline	100	100	100
Player 1: Deduct from 2	98	50	100
Player 1: Transfer from 3 to 1 Player 2: Yes	140	90	50
Player 1: Transfer from 3 to 1 Player 2: No	100	100	100

At the end of each stage, the stage's outcome will be displayed to the three players.

<u>What happens at the end of the experiment?</u>

Once all rounds have been completed, there will be a short and optional exit survey. Your participation is voluntary and does not affect your payoff.

At the end of the experiment, the computer randomly chooses one round for your payment. Your experimental-point total for that round will be converted to dollars at the rate of $0.05 per experimental point. You will be paid the dollar payoff from your selected round in addition to the $8 show-up fee. All player identities remain anonymous. No one will learn what role you played or what payoff you earned.

We are almost ready to start the experiment. Before doing so, there will be a short quiz to check your understanding of some key features of the experiment, as well as a chance to examine how the player decision screens work.