Bad Repetition*

Geoffroy de Clippel† and Kareen Rozen‡

July 2020

Abstract

This paper experimentally explores the often-overlooked negative implications of repeated interactions, whereby promises of rewards and threats of punishments may lead people to take actions that are detrimental to themselves and/or others. For both finitely and infinitely repeated games, we demonstrate the prevalence of negative equilibria (which are worse than repetition of the one-shot Nash equilibrium) by departing from the stage games most commonly studied in the experimental literature. For the simple games we study, repetition can open the door to equilibria that capture stylized aspects of peer pressure and bystander complacency.

1 Introduction

In a repeated game, players face the same strategic problem – the *stage game* – multiple times, with payoffs accruing in each round. A player can, at minimum, achieve her 'minmax' utility in each repetition: this is the best payoff she can guarantee herself when everyone else acts in concert against her. A player's payoff is *individually rational* if it dominates this lower bound. What can be said beyond this about equilibrium payoffs in the repeated game?

*We are grateful to Pedro Dal Bó 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.

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

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

A variety of 'folk theorems' suggest that *any* individually rational payoff profile achievable in the stage game can be sustained by some equilibrium of the repeated game. 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. Of course, the minmax coincides with the stage-game Nash equilibrium (NE) payoffs when players have dominant strategies. In such cases, players' welfare can only improve if repeating the stage game has any impact at all.

Yet there are important repeated settings where the minmax is worse than NE play, and where plausible strategic considerations can lead participants to take detrimental actions. Thus repetition can also be bad. Beyond what casual observation and theoretical arguments teach us, it is an empirical question whether people would abide by such strategies and how detrimental conventions (norms, cultures) crystallize. Our experimental study sheds light on these questions using two different stage games. As seen below, these are simple stage games whose repetition opens the door to negative equilibrium outcomes. While these games capture some stylized aspects of bullying, peer pressure and bystander complacency, we do not view them as models of those harmful phenomena.

Stage Game I: Peer-pressure game. We use suggestive language here to describe the stage game G_I and the underlying intuition more effectively, but player names and strategic options in the laboratory experiment were described using neutral terms. There is a group V of potential *victims* and a group B of *bullies*, each of whom selects a player in V to bully. These individuals play independently of each other, that is, G_I is a normal-form game. A bully's payoff increases by $\$x$ when he is part of a strict majority of players in B targeting the same player in V. One can think of this as bullies seeking each other's social acceptance by coordinating on a common victim. A victim incurs a loss of $\$\ell$ if picked by at least one bully. Simultaneously, each potential victim decides whether to pay a cost $\$c$ or not, which reduces their own payoff with no other immediate implication. One can think of paying this cost as engaging in conspicuous consumption that brings no utility (e.g., purchasing the trendy pair of shoes). Alternatively, one can think of paying the cost as engaging in observable consumption behaviors that the individual finds unappealing, such as smoking or binge drinking. Of course, there are many variants of this game one could study. Our objective is to establish stylized facts when this simple, normal-form stage game is repeated.

We are particularly interested in the behavior of potential victims, that is, players in V. Consider, for instance, a schoolyard where bullies pick on children who don't wear the trendy pair of shoes. Clearly, it is a strictly-dominated strategy for them to pay $\$c$ when G_I is played only once: paying the cost just reduces utility, no matter the bullies' strategies.[1] Once the game is repeated, however, one can easily construct self-enforcing conventions where a large number of potential victims engage in conspicuous consumption, even though just one victim is ultimately bullied. Under such norms, the children feel pressured to purchase the trendy items, so as to avoid becoming a target of the bullies in the future. With victims effectively competing with each other to stay under the bullies' radar, bullying has wider consequences than the number of individuals actually targeted.

As is well-known, if a stage game has a unique NE, then the only subgame-perfect equilibrium (SPE) of the *finitely* repeated game consists of playing that unique NE in each period. The reason is simple: there is no room for credible threats of future punishment or reward based on past behavior. Whatever happened in the past, the stage-game NE must prevail in the last round. By backward induction, the same holds true in all rounds. Things are dramatically different, however, when the stage game admits multiple (non-payoff equivalent) NE.[2] In G_I, for instance, bullies face a pure coordination problem, and coordinating on any given victim forms an NE of the stage game. Even though a stage-game NE must prevail in the last round, which one gets selected can be conditioned on past-rounds' outcomes. Hence for this stage game, even the chance of a single repetition can lead to bad outcomes. To make the starkest possible point, we will examine the game through that lens (the argument immediately generalizes to any finite repetition).

Consider the following convention for bullies, with $V = \{v_1, \ldots, v_n\}$: in Round 1, the bullies coordinate on targeting victim v_1; but in Round 2, they coordinate on the smallest-indexed victim among those that did not pay the cost in Round 1 (if any exist), and coordinate on victim v_n otherwise. One can think of the index

[1] Throughout most of the paper, we assume players maximize their own monetary payoffs. We will argue, however, that our analysis is robust to many forms of other-regarding preferences.

[2] Benoit and Krishna (1985) prove a limit folk-theorem for a large class of stage games: "any feasible and individually rational payoff vector of the one-shot game can be approximated by the average payoff in a perfect equilibrium of a repeated game with a sufficiently long horizon." Such result is not applicable to our analysis, as G_I will be played at most twice.

as encapsulating a victim's salience to the bullies, which can be reduced through conspicuous consumption (except[3] for v_n). Clearly, bullies play their part of a NE of the stage game by following this convention. Paying the cost is not part of a NE of the stage game for potential victims, but may make sense in the repeated game, so long as $\ell > c$ and the probability of a second round is high enough. Consider victim v_1: she feels pressured to pay the cost, as it shelters her from being bullied in the next period. Her doing so, however, puts pressure on higher-indexed players to pay the cost themselves, else they will be the next target. Indeed, an unraveling argument implies that all potential victims other than v_n also prefer to pay c to stay under the bullies' radar in the next period (v_n has no way to avoid the bullies). Thus, no matter how large the set of potential victims is relative to the set of bullies, $n-1$ victims paying comprises an SPE of the repeated game when the probability of a second round is high enough. In that case, repetition induces a deadweight loss, with potential victims other than v_n paying c to stay under the bullies' radar in the event of a second round.

Stage Game 2: Bystander game. Imagine a school bully trying to take the lunch money of a younger student. One of the bully's classmates sees this scene unfold. Will the classmate, who is a bystander to the bullying and finds it difficult to tolerate such behavior, intervene? Suppose there is no exogenous cost to intervention (e.g., risk of injury). Then, the bystander would certainly intervene if this interaction occurs only once. Bystander complacency can potentially arise in equilibrium, however, if repeated encounters are likely. Retaliating against the bystander can be a credible threat in the repeated game, even if retaliating is costly to the bully. Thus, there can be an endogenous cost of intervention. We designed a simple stage game to test for such an effect.

As illustrated in the game tree in Figure 1, the stage-game G_{II} consists of a bully, a bystander, and a helpless victim who has no strategic choices. The players are each initially endowed with $\$m$. The bully can try to steal $\$t$ from the victim, but needs the bystander to be complacent in order to succeed. The bully can choose to harm the

[3]The game where conspicuous consumption also reduces v_n's salience has similar properties but will have players in V randomizing at equilibrium. Given that our primary interest is to learn about repeated games (instead of modeling a precise situation), we opted for a case where v_n's decision has no impact on his salience.

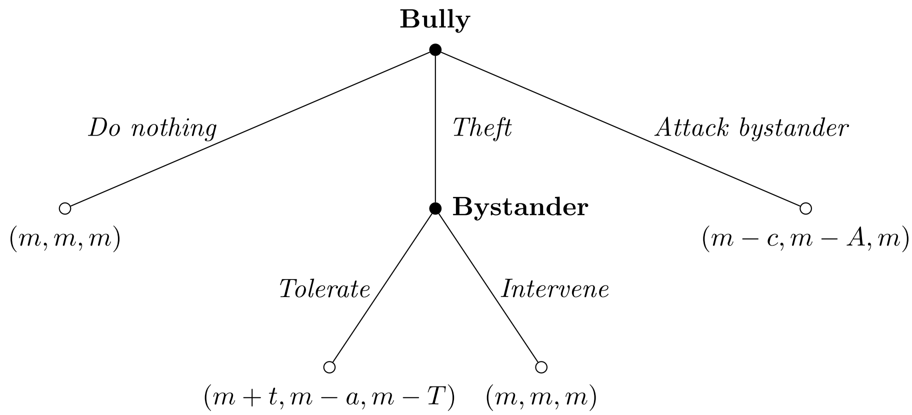

Figure 1: G_{II}-game tree, where $c > 0$, $0 < t < T$ and $0 < a < A$. The first payoff is for the bully, the second for the bystander, and the third for the victim.

bystander, but only by incurring a cost \$$c$.[4] Theft leads to a social loss, as captured by the loss of $T > t$ to the victim and $a > 0$ to the bystander, who finds it difficult to quietly witness aggression.[5] However, tolerating the bullying is less painful for the bystander than being attacked directly, as captured by $a < A$. Our purpose is to offer a stylized model capturing some of the general forces at play; one could imagine alternative formulations where the bystander directly benefits from protecting the victim (e.g., there is a psychological benefit or monetary reward for heroism), but the formal analysis is robust to such considerations. Of course, options and player names were described to experimental subjects in neutral language, without using terms like bully or theft.

If the game is played only once, then backward induction (and NE) leads to a unique prediction: the bystander will intervene to stop the bully should he aggress anyone, resulting in an equal payment of \$$m$ to each player. Hence, as explained earlier, this will remain the SPE outcome if G_{II} is repeated with a known deadline. Consider instead the case of an uncertain deadline, another classic scenario in the repeated-game literature. Each time G_{II} is played, there is a fixed probability that the interaction carries on for another period (which is essentially equivalent to an infinitely-repeated game with exponential discounting).

[4]This ensures the bully would only want to harm the bystander because of strategic considerations in the repeated game.

[5]For example, complacent bystanders may feel guilt, or have distributional preferences and dislike the increased inequality. While subjects in our experiment may incur a utility loss from such sources, we introduced the monetary cost a to ensure a tangible loss to the bystander from complacency.

Theoretically, new conventions (norms, cultures) become self-enforcing when the continuation probability is large enough. Some of these have the bystander paying the cost for going along with the bully's theft from the victim (on the equilibrium path), because the bully would punish the bystander in future rounds otherwise (off the equilibrium path). In that case, both the bystander and the victim end up with a lower payoff, and equilibrium payoffs become unfair. While the bully is better off, inequality increases and the total sum of payoffs is lower in this equilibrium with theft: complacency is costly and the transfer from the victim to the bully happens at a loss. In particular, the outcome will be Pareto inefficient if this convention is followed in richer games where players can, in turn, be bully, victim or bystander.

Two Main Empirical Questions. Game theory highlights that repeated games can admit a multitude of equilibria, but mostly remains silent as to which might become focal in each specific circumstance. There are likely many contextual details in the real world affecting which norms arise, but purely strategic considerations do not reveal which aspects are relevant to replicate in a controlled laboratory setting (if replicating them is at all possible). Interaction may be too short and artificial to see subjects' strategies settle on a strongly inefficient norm among those that are self-enforcing. Prior to addressing any question of equilibrium selection, one must ask the following: *Can behavior resulting in bad outcomes be self-enforcing due to repeated interactions?* We propose the following methodology to test this: using neutral terms, a negative convention for how some players could choose is suggested (which subjects are free to follow or not); and we examine whether subjects freely play their part in the associated 'bad equilibrium' when it is self-enforcing due to the likelihood of repetition, but not otherwise.

Consider the finitely-repeated game associated to G_I. The convention mentioned earlier is obviously Pareto inferior to the natural alternative NE where the bullying target remains unchanged while victims do not pay the cost (a dominant strategy in the stage game). It is, furthermore, Pareto inferior to all one-shot NEs when comparing anonymous payoff vectors (that is, ranking players' payoffs in increasing order without paying attention to identities). Indeed, all NEs of the stage game have one bullied victim and no one paying the cost. While casual observation suggests that, unfortunately, such detrimental conventions may crystallize in real-world situations, we suspect it may be hard to see them arise spontaneously in a short, abstract encounter of a standard lab experiment. Instead, we followed the new methodology above, by simply offering the above convention as a non-binding suggestion for which victim bullies could pick as a function of the history. We describe the experiment, theoretical predictions, and empirical results in Section 3. As seen there, by the end of a session virtually all subjects play their part in the peer-pressure equilibrium when the probability of a second encounter is almost certain (99%). By contrast, hardly any subjects pay the cost when that probability is minimal (1%). Thus, a simple suggestion becomes quickly entrenched into a culture of peer pressure when the behavior is an equilibrium (i.e., with high continuation probability), but not otherwise. The strong adherence of subjects to the negative equilibrium in the 99% treatment occurs despite evidence from other works that a significant portion of subjects are altruistic (Andreoni et al., 2008); and despite the fact that subjects who pay the cost are not just shielding themselves from the bullies, but effectively deflecting them onto others.

In addition to testing the sustainability of some detrimental norms due to repeated interaction, we also explore the possibility of them arising spontaneously. This is our second question of interest: *Can detrimental conventions, that become self-enforcing when future encounters are likely, spontaneously crystallize?* As discussed earlier, not seeing a norm spontaneously crystallize in the lab does not mean that it wouldn't arise in richer contexts. For instance, we doubt that the above Pareto inferior-convention would spontaneously occur in the lab when G_I is repeated, and we did not test behavior in the absence of a suggested convention for that game. Instead, we turned our attention to G_{II}. Rather than sustaining a Pareto-inferior equilibrium outcome, repetition now creates new self-enforcing conventions that are better for the bully. We conjectured this may be a reason why such conventions might spontaneously crystallize, even in the lab (and a fortiori in the real-world), as it is in the bully's advantage to use aggression strategically to make that equilibrium focal. Though not comparable in terms of the Pareto criterion, bullying equilibria in the repeated game are still bad in the sense of generating unequal payoffs, and decreasing the total sum of payoffs. Here, we consider two treatments: a continuation probability of 75%, under which it is theoretically possible for the bully to sustain transfers in every period, and a continuation probability of 10%, under which it is theoretically impossible for the bully to sustain transfers in any period. In Section 4, we describe the experiment, theoretical predictions, and results. As seen there, we find that the bystander almost always intervenes in the 10% treatment, but that around 42% of bystanders are complacent in the 75% treatment. Those bystanders who do intervene often face punitive actions. We do find interesting gender differences in this game. Men and women act nearly identically in the 10% treatment, and also act similarly as bystanders in the 75% treatment. However, we find that women are significantly less likely to react punitively towards bystanders who intervene; and they are significantly less likely to request the bystander's complacency in the first place, when that request is enforceable in equilibrium.

2 Related Literature

There is a substantial experimental literature on repeated games. Most effort has been devoted to infinitely repeated games (or with a random termination rule), and more specifically in determining 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. When alternative games (e.g., investment in a public good, Cournot competition, gift-exchange, etc.) are considered, the main focus is still on the possibility of sustaining cooperative outcomes that improve upon the stage game NE. While reaching more efficient outcomes is obviously desirable whenever possible, we collect data for an entirely different stage-game (the bystander game G_{II}) to test whether repetition may induce less desirable outcomes instead. We not only find that such undesirable outcomes occur, but observe players retaliating against those who try to prevent them. In a further contrast to the previous literature, punishments occur through a minimax action (i.e., attacking the bystander) that is not part of a one-shot Nash equilibrium, as opposed to defection in the prisoner's dilemma. This demonstrates experimentally the force of Fudenberg and Maskin (1986)'s folk theorem, which uses minimax punishments.

The experimental literature on finitely repeated games is much smaller, and focuses mostly on why cooperation may occur, especially in earlier rounds, when the prisoners' dilemma is repeated a fixed number of times. Conducting a meta-analysis of prior experiments and collecting new data, Embrey et al. (2018) provide the best account so far. In a nutshell, participants' play get closer to backwards induction, and cooperation lasts for fewer rounds, as they gain experience, but unraveling can take a very long time. By contrast, we study a twice-repeated game (G_I, the peer-pressure game) that admits multiple stage-game NEs, and explore the stability of a norm that leads to a

Pareto inferior outcome.

The use of a suggestion on how to play the game to test the self-enforcing nature of specific strategy profiles is not new in general,[6] but (as far as we know) has not been applied before in repeated games. Though studying the spontaneous emergence of norms is a natural subsequent question, we find the method particularly appealing in such games. First, it provides a conservative way to trim the huge set of equilibrium strategies that is characteristic of repeated games. Second, the self-enforcing nature of a norm can be tested while controlling for experimenter's demand effects. Indeed, one can compare subjects' obedience to the same norm described in exactly the same terms, while varying the continuation probability or the number of repetitions.

Though the repeated-games literature has focused on the potential for positive equilibrium outcomes, the possibility of perverse outcomes has been considered in more elaborate dynamic games. In the context of dynamic incomplete-information games with a single long-run player, Ely and Välimäki (2003) also point to the potential for outcomes that overturn conventional wisdom. With a simple model, they show that a long-lived agent's reputational concern undermines his commitment power, and erase any surplus. Ely et al. (2008) theoretically characterizes when such 'bad reputation' is possible, and Grosskopf and Sarin (2010) examines the possibility in an experimental study where reputation can be either good or bad. They observe that while the benefits of reputation are weaker than predicted by the theory, reputation is rarely harmful. By contrast, we examine the more basic setting of repeated games and make a simpler point: for stage games beyond those commonly studied, repeated-game equilibria worse than repetition of the one-shot Nash equilibrium can be prevalent.

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

[6] See Brandts and Holt (1992), van Huyck et al. (1992) and Brandts and Macleod (1995) for early uses of this method.

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,[7] and the possible necessity of preventive war is not the consequence of folk theorems.

3 Game I: Peer-Pressure Game

We begin by describing the Game I-experiment (in the neutral language used for subjects), along with the precise payoffs used.[8] Subjects were told that at the start of the experiment, the computer randomly assigns 4 out of every 7 participants to a *low-index role*, and assigns the others to a *high-index role*; and that subjects retain those roles throughout the session.

There are 30 matches in the experiment. At the start of each match, each subject is randomly matched into a group of seven, with each group comprising 4 low-index and 3 high-index members. For each group, the low-index members are randomly assigned a distinct index from 1 to 4, and the high-index members are randomly assigned a distinct index from 5 to 7. Each subject is told the index he or she is assigned for that match, but is never told which role or index is assigned to another participant. Each group in a match plays a supergame that has at most two rounds, which are identical in structure. Hence the supergame is a finitely-repeated game. The probability of having a second round is held constant throughout the 30 matches of a session, at either 99% or 1%. Subjects learn whether Round 2 will be held only after Round 1 ends. A subject's payoff from a match is the sum of payoff(s) accrued in the round(s) of that match.

We now describe what happens in each supergame. All groups members see a non-binding *suggestion* at the beginning of each round, described further below. Group

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

[8] We note, however, that the experimental instructions for Games I and II use the term 'stage' for each time the stage game is played within the supergame, and the term 'round' for the supergame itself. The discussion in the paper departs from our instructions in order to avoid confusion with some other papers' use of the word 'round' to describe each play of the stage game.

members make their choices simultaneously, with each high-index member choosing an index between 1 and 4, and each low-index member deciding whether to pay a cost of \$2. The round payoff of a high-index member is \$5 if their choice matches that of at least one other high-index member in their group, and \$0 otherwise. The baseline round-payoff of a low-index member is \$13, from which there is a \$2 deduction if the low-index member pays the cost in that round, and there is a separate \$10 deduction if their own index is chosen by some high-index member of their group in that round. At the end of the round, subjects are informed of the choices made by the members of their own group, as well as the suggestion for the next round, if applicable.

The suggestion for each group is constructed as follows. In Round 1, the suggestion is always for high-index members to select index 1. In Round 2 (if it occurs), the suggestion depends on Round-1 choices in that group. The computer displays the smallest index among those in the group who didn't pay the cost in Round 1, and suggests for high-index members to pick that index in Round 2. If all low-index members in the group paid the cost in Round 1, then the group's suggestion will be for high-index members to pick index 4 in Round 2. Subjects are free to follow the suggestion or not. Payoffs are constructed as above, regardless of whether suggestions are followed.

At the end of the experiment, the computer randomly chooses one match, and each subject is paid the \$8 show-up fee plus their payoff in that match. Subjects' roles, choices, the groups and indices they were assigned, and final payoffs all remain private. All subjects were paid in cash before leaving the room.

All sessions were conducted at the Brown University Social Sciences Experimental Laboratory (BUSSEL). The laboratory is equipped with sunken terminals and vertical privacy panels between desks. Subjects were allowed to participate in at most one session, and could not participate in both Game I and Game II sessions. Subjects were recruited via the BUSSEL website.[9] A total of 161 subjects participated in the Game-I experiments, which took place in 2017 and 2018. There were four sessions with a 99% continuation probability (with 112 subjects) and two sessions with a 1% continuation probability (with 49 subjects). Subjects could participate in at most one

[9]This site, available at bussel.brown.edu, offers an interface to register in the system and sign up for economic experiments. To do so, the information requested from subjects is their name and email address and, if applicable, their school and student ID number. The vast majority of subjects registered through the site are Brown University and RISD graduate and undergraduate students, but participation is open to all interested individuals of at least 18 years of age without discrimination regarding gender, race, religious beliefs, sexual orientation or any other personal characteristics.

session of this experiment; and were not allowed to participate in both the Game-I and Game-II experiments. The experiment was programmed using the z-Tree software (Fischbacher, 2007). The instructions, available in the Appendix, were provided in paper format at the start of the experiment and read out loud by the experimenter. Prior to starting the first match, subjects were asked to complete a comprehension quiz through the z-Tree interface to check their understanding of indices, payoffs, and suggestions. After the final match, subjects were given the option to complete an exit survey.[10]

3.1 Theoretical Predictions

Consider how subjects might behave in Game I. A first observation is that without a suggestion, the high-index members (the bullies) would have no obvious device to coordinate on the same low-index member and receive the \$5 payment. If they were to choose uniformly at random, then for any index $i \in \{1, 2, 3, 4\}$ chosen by a bully, there would be a $(3/4)^2 = 56.25\%$ probability of receiving \$0 in that round. Nor do low-index members have any way to impact high-index members' payoffs. Any common suggestion thus provides a potentially valuable coordination device for the bullies, regardless of the round being played and the continuation probability. Moreover, if high-index members expect each other to follow the suggestion, then not following it surely leads to a payoff of \$0. For these reasons, we would expect high-index members to follow the suggestion. This prediction remains valid even if high-index individuals have altruistic preferences. Choosing only *which* index to select, the suggested coordination device has the altruistic benefit of reducing the number of indices selected, and therefore the number of \$10 deductions among low-index members.

Assuming that bullies follow the suggestion, subjects with an index $i \in \{1, 2, 3\}$ would most likely pay the cost in any Round 1 of the 99% treatment, as the \$2 cost in Round 1 is small compared to the 99% chance of a \$10 benefit in Round 2. Of course, paying the cost amounts to deflecting the bullies from themselves *onto someone else in Round 2*, which may induce an extra cost for altruistic subjects. But the conclusion remains unchanged so long as the utility function places sufficient weight on one's own payoff. By contrast, the future \$10 benefit is quite unlikely in the 1% treatment.

[10] In one of the two 1%-sessions, a technical error prevented the recording of survey data only.

This is especially true for indices $i \in \{2, 3\}$, who are only at risk of being targeted by the bullies if index 1 pays (assuming the suggestion is followed). Index 4 cannot impact the Round-2 suggestion, as explained further below, and has no incentive to pay. Clearly, for any low-index member there is zero benefit from paying the cost in Round 2; this contrasts with Round 1 in the 99% treatment.

Notice that this game requires some iterated reasoning in the 99% treatment. The benefit of paying the cost in Round 1 of that treatment is immediately obvious for index 1. But index 2 needs to recognize that index 1 will pay, in order to recognize the benefit of paying himself. Requiring yet another step of reasoning, index 3 needs to recognize that index 1 will pay, and that as a consequence index 2 will pay, in order to recognize the benefit of paying himself. Non-trivial reasoning is also required by index 4, to recognize that he cannot impact the suggestion by paying the cost: if some others in his group don't pay the cost then the suggestion will target one of them, and if everyone in his group pays then the suggestion will target him regardless of his current choice. The subjects play 30 repeated games, and those in the low-index role do have the opportunity to experience different low indices, because of the random reassignments at the start of each match.

We summarize our theoretical analysis in the following hypothesis.

Hypothesis 1. *High-index members follow the suggestion. After perhaps some learning in earlier matches, members with an index lower than or equal to* 3 *pay the cost in Round 1 of the 99% treatment. Low-index members do not pay the cost otherwise. As a result, outcomes in the 99% treatment should be Pareto inferior to those in the 1% treatment.*

3.2 Experimental Results

We find strong empirical support in favor of high-index members following suggestions: 99.01% of the 2,832 choices in the 99% treatment (including 99.07% in Round-2), and 97.08% of the 651 choices in the 1% treatment (including 100% in Round-2, which occurred in only one match of that treatment).

The data also confirms the general pattern of choices anticipated for low-index members. As noted, we would expect some learning to occur. Consider first the bar graphs in Figure 2, which depict the proportion of times each index $i \in \{1, 2, 3, 4\}$ pays the cost, aggregated over successive groups of five matches, in the three different

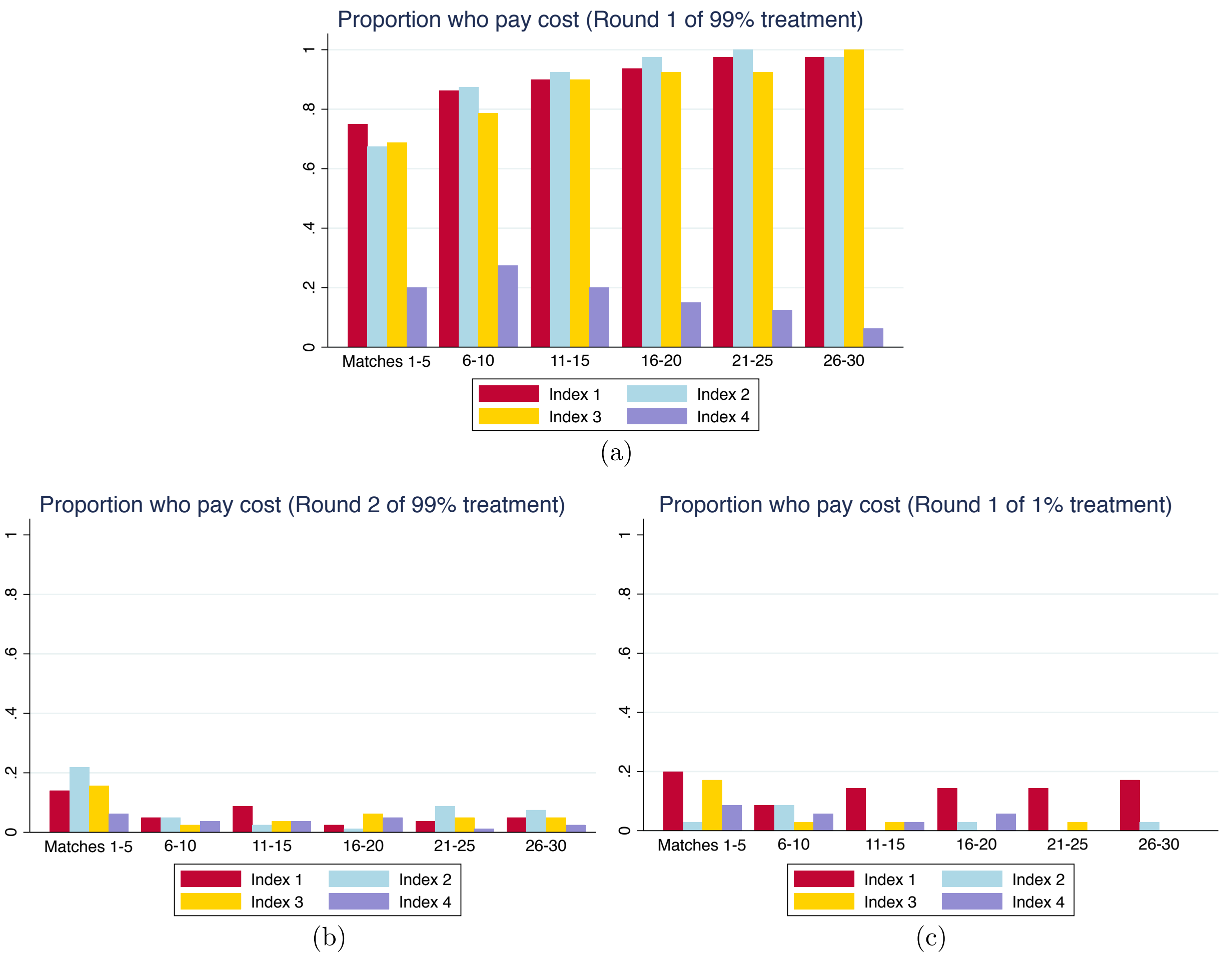

Figure 2: Proportion of times index $i \in \{1, 2, 3, 4\}$ pays the cost, over five matches.

scenarios (we do not include Round 2 of the 1% treatment, as only one match had a second round). In Figure 2(a), depicting Round 1 of the 99% treatment, there is a steady increase in payment from all indices except for index 4. By the last five matches, the cost is paid by 97.5% of indices 1 and 2, and 100% of index 3; while it is paid by only 6.3% of index 4. Figures 2(b) and 2(c), depicting Round 2 of the 99% treatment and Round 1 of the 1% treatment, show a markedly different picture. Both figures exhibit low levels of payment in all but the first five matches, with the exception that around 17% of index 1 members pay the cost in Round 1 of the 1% treatment. This may be consistent with some index-1 members acting out of an abundance of caution: if they do not pay, they would be the bully's target in the (highly unlikely) second round.

As is customary, we drop earlier matches (Matches 1-10, or the first third of all matches that each subject played) to discount transitory behavior that disappears after learning.[11] Focusing on Matches 11-30 and using logistic regression with heteroskedasticity-robust errors (clustered by session), we estimate and compare the probabilities each index $i \in \{1, 2, 3, 4\}$ pays in the three main circumstances.

	99% treatment		*1% treatment*
Index	Round 1	Round 2	Round 1
1	0.951^{***} (0.014)	0.049^{**} (0.016)	0.143^{**} (0.045)
2	0.974^{***} (0.009)	0.053 (0.037)	0.015 (0.010)
3	0.938^{***} (0.029)	0.053^{*} (0.026)	0.015 (0.010)
4	0.128^{***} (0.032)	0.030^{*} (0.014)	0.023^{***} (0.003)
Observations	1216	1216	532

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

Table 1: Estimated payment probabilities over Matches 11-30, with robust standard errors (clustered by session) in parentheses.

The table overall shows the behavioral patterns predicted in Hypothesis 1. Each index $i \in \{1, 2, 3\}$ pays with very high probability in Round 1 of the 99% treatment. For each $i \in \{1, 2, 3\}$, that probability is strictly smaller than one (p-value ≤ 0.0286) but much larger than that same index's payment probabilities in the other two scenarios from Table 1 (p-value 0.0000). Moreover, index 4 indeed pays with a significantly smaller probability than any $i \in \{1, 2, 3\}$ (all p-values 0.0000), though that probability is statistically different from zero. For each $i \in \{2, 3, 4\}$, the null hypothesis that index i pays with the same probability in Round 2 of the 99% treatment and Round 1 of the 1% treatment cannot be rejected (p-values 0.3320, 0.1840, and 0.6204, respectively). We marginally reject the null that index 1 pays with the same probability in Round 2 of the 99% treatment and Round 1 of the 1% treatment (p-value 0.0493); the mere possibility of a second round in the 1% treatment thus appears to impact index 1's behavior.

Finally, first-round average payoffs for each index are provided in Table 2. We observe a Pareto ranking as expected: high-index members' average payoffs do not vary much across treatments and the difference is indeed not statistically different

[11] For the 99% sessions, each of matches 11-30 had a second round. For the 1% sessions, only one of the matches 11-30 had a second round.

	Index 1	Index 2	Index 3	Index 4	Index 5	Index 6	Index 7
$\delta = 1\%$	2.70	12.54	12.83	12.96	4.89	4.89	4.89
$\delta = 99\%$	1.11	10.88	11.13	12.67	4.95	4.98	4.94

Table 2: Observed first-round average payoffs over Matches 11-30.

(though marginal for index 6: p-values 0.2950 for index 5, 0.0520 for index 6 and 0.4721 for index 7 using a Wilcoxon rank-sum test). Payoffs for indices 1 to 3 in the 99% treatment are strictly lower than in the 1% treatment, with p-values 0.0000. For index 4, one would expect comparable payoffs in both treatments, provided that subjects understand that paying the cost does not shield themselves from being targeted when in that role. As noted earlier, index 4 paid in a small and statistically significant fraction of choices in the 99% treatment, a mistake which only decreases that index's average payoff too (p-value 0.0001).

4 Game II: Bystander Game

We start by describing the Game II-experiment (in the neutral language used for subjects), along with the precise payoffs used.

At the beginning of the experiment, each subject is randomly assigned to one of three possible roles: Player 1, Player 2, or Player 3. The assigned role is held for the duration of the session. The experiment has 15 matches. At the start of every match, each subject is matched into a three-person group. Each group is 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 one supergame. A supergame may consist of multiple rounds, all of which are identical in structure (it is a repeated game). 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.

Each round unfolds as follows. The baseline payoff in every round is 100 experimental points (EP) for each player. Within a group, the current round payoffs can be modified from these baselines by Player 1's and 2's decisions. Player 1 moves first and has three options: *Baseline*, *Deduct from 2*, and *Transfer from 3 to 1*. Choosing the last option requires Player 2 to choose either *Yes* or *No*. The possible round payoffs, in EP, are summarized as follows:

Outcome	*Player 1*	*Player 2*	*Player 3*
Baseline	100	100	100
Deduct from 2	98	50	100
Transfer from 3 to 1 / Yes	140	90	50
Transfer from 3 to 1 / No	100	100	100

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. All player identities remain anonymous. Subjects were not told what roles other subjects played, or what payoffs others received.

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 were allowed to participate in at most one session, and could not participate in both Game I and Game II sessions. Subjects were recruited via the BUSSEL website. A total of 225 subjects participated in the Game-II experiments, which took place in 2017 and 2018. There were 120 subjects in the sessions with a 75% continuation probability, and 105 subjects in the sessions with a 10% continuation probability.

The instructions, available in the 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 the option to complete an exit survey. Subjects were paid in cash before leaving the laboratory.

4.1 Theoretical Predictions

As is standard 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. Each player getting 100 EPs in each round is the unique SPE outcome with the low continuation probability ($\delta = 10\%$). 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 $\delta = 10\%$.

The equilibrium set is much larger when $\delta = 75\%$. In fact, Player 1 can even achieve her maximal payoff of 140 EPs in each round. To see this, 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 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 as for $\delta = 10\%$, we conclude that the range of discounted average payoffs for Player 1 is now the interval $[100, 140]$. 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. Of course, this goes hand-in-hand with greater inequality and less sum-efficiency. We can thus formulate the following testable hypothesis:

Hypothesis 2. *Complacency and successful transfers should be rare when δ is low,*

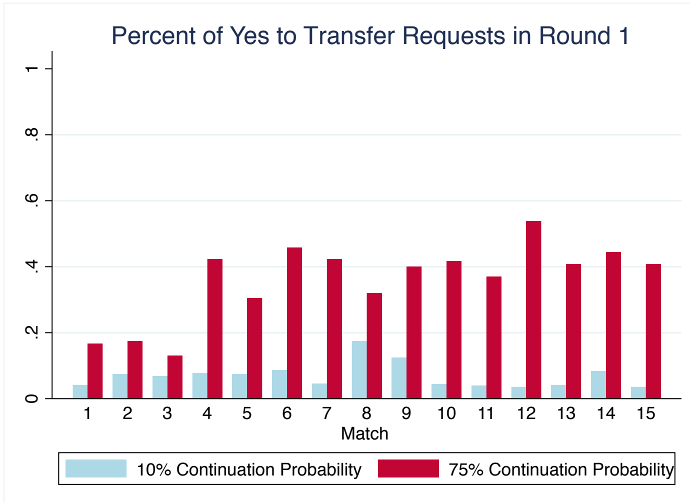

Figure 3: Response to transfer requests in Round 1, per match.

but significantly more frequent when δ is large. As a result, we expect higher inequality and lower total surplus in the latter case.

4.2 Experimental Results

Over the course of multiple sessions with 15 matches each, 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 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 Hypothesis 2

We start by looking at complacency rates. Figure 3 offers a match-by-match look at Round-1 choices of Player 2 following transfer requests. There appears to be some learning in the first few matches. 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. As before, we focus on the last two thirds of our data (dropping Matches 1-5) to account for learning. We then use logistic regression to estimate the probability of *Yes*, along with heteroskedasticity-robust standard errors (clustered by session) in Round 1, conditional on a transfer request. The estimates appear in Table 3.

The 0.419 probability of agreeing to a transfer in Round 1 of the 75% treatment is significantly different from zero (p-value 0.0006), and is significantly larger than the 0.070 probability in the 10% treatment (p-value 0.0049). Though small, the latter is nevertheless statistically different from zero (p-value 0.0022).[12] Of course, the practical

	$\delta = 10\%$	$\delta = 75\%$
Complacency Rate	0.070** (0.023)	0.419*** (0.122)
Observations	244	258

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

Table 3: Estimated complacency rates (Player 2 picking *Yes* following a transfer request) in Round 1, Matches 6 to 15. Robust standard errors clustered by session.

importance of complacency rates hinges on the prevalence of transfer requests. Table 4 shows the estimated percentage of successful transfers (a more detailed analysis of Player 1 choices can be found in the next subsection). As hypothesized, it is very low when $\delta = 10\%$ and substantially larger when $\delta = 75\%$ (p-value 0.0051).

	$\delta = 10\%$	$\delta = 75\%$
Successful Transfer Rate	0.049** (0.017)	0.270*** (0.077)
Observations	350	400

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

Table 4: Estimated percent of successful transfers for Round 1, Matches 6-15.

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). Table 5 confirms this overall trend. As hypothesized, there is a decrease in total surplus ($p = 0.0000$ according to the 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

[12]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).

	Player 1	Player 2	Player 3
$\delta = 10\%$	101.82	96.37	97.57
$\delta = 75\%$	110.76	96.30	86.50

Table 5: Observed average payoffs for Round 1, Matches 6-15.

	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
	(0.041)	(0.013)	(0.031)	(0.031)	(0.111)
Deduct	0.063	0.020	0.250^{*}	0.337^{***}	0.000
	(0.049)	(0.011)	(0.125)	(0.099)	(−)
Transfer	0.697^{***}	0.645^{***}	0.719^{***}	0.606^{***}	0.787^{***}
	(0.068)	(0.010)	(0.133)	(0.102)	(0.111)
Observations	350	400	32	104	75

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

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

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 (though not statistically significant) 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 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 Action Choices

Table 6 provides further information on Player 1's choices in the first two rounds of Matches 6 to 15. 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 $\delta = 10\%$. Though 6.3% of subjects did pick it, the estimate is not statistically different from zero. Picking *Deduct* is part of an 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. But such Pareto-dominated equilibria are unlikely to occur. We see only a 2% probability of first-round deductions, which is not statistically different from zero.

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, 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 6 are significantly larger than 1/2 (p-values of 0.0000 for $\delta = 75\%$ and 0.0038 for $\delta = 10\%$). There is a small difference across treatments regarding the proportion of *Baseline* choices: 24% when $\delta = 10\%$ instead of 33.5% in the other case (p-value 0.0264).

We observe here some interesting gender differences in Player 1's first-round choices. As seen in Table 7, men and women choices are nearly identical when $\delta = 10\%$ (p-value 0.9332), but are rather different when $\delta = 75\%$ (p-value 0.0000). While men are at least four times more likely to request *Transfer from 3 to 1* in the first-round than select *Baseline*, women are split about evenly between these options (see Table 7). On the other hand, there are no significant gender differences in first-round choices in the role of Player 2 (p-values 0.1700 for $\delta = 10\%$ and 0.9915 for $\delta = 75\%$).

	Round 1 $\delta = 10\%$		Round 1 $\delta = 75\%$		Round 2 After No $\delta = 75\%$	
	Men	Women	Men	Women	Men	Women
Baseline	0.230^{**} (0.070)	0.253^{*} (0.118)	0.179^{*} (0.072)	0.474^{***} (0.050)	0.033 (0.033)	0.081 (0.055)
Deduct	0.070 (0.067)	0.053^{*} (0.026)	0.021^{*} (0.009)	0.022 (0.015)	0.733^{***} (0.070)	0.161 (0.088)
Transfer	0.700^{***} (0.079)	0.693^{***} (0.123)	0.800^{***} (0.073)	0.504^{***} (0.057)	0.233^{**} (0.071)	0.758^{***} (0.066)
Observations	200	150	140	230	30	62

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

Table 7: Estimates of Player 1's choices by reported gender, with robust standard errors, clustered by session.

In our benchmark analysis with $\delta = 75\%$, the use of deductions in histories where some transfers were denied was critical to make successful transfers possible. Table 6 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 Rounds 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 differs from zero at all standard levels of significance. There are some interesting gender differences here too. The rightmost columns of Table 7 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 tend to repeat the transfer request, deducting only 16% of the time (which is not different from zero at the 5% level). By contrast, there is no difference in behavior by gender following a *Yes* from Player 2 (p-value 0.8502).

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 6). With only 32 instances of such denials in the first place, the standard deviation is rather high, but the number is still marginally different from zero at the 5% level (p-value 0.0453). Though smaller than the estimated 33.7% who deduct when $\delta = 75\%$, the difference is not statistically significant (p-value 0.5877). Note that we cannot provide estimated probabilities of Player 1's Round-2 choices after an approved transfer when $\delta = 10\%$, as there are only three such instances in our dataset. It is safe though to conjecture that 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 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 different directions to pursue. In the spirit of Dal Bó (2005), one could contrast behavior when Game II 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). Interesingly, 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 cannot exclude the possibility that deductions in Game II 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.

5 Concluding Remarks

The experimental literature on repeated games focuses on a relatively small set of classic stage games, where the main question of interest is whether better outcomes are possible than repetition of one-stage Nash. We find it worthwhile to widen the scope of games studied. When the minmax utility in the stage game does not coincide with the NE payoffs, repetition opens the door to negative conventions that are worse than repetition of one-stage Nash. We establish that such bad equilibria are unfortunately sustainable in a laboratory setting, both in finite and infinitely repeated games. It is important to understand both the positive and negative implications of repetition, and both how and why conventions crystallize in practice. Of course, exploring the entire equilibrium set requires taking advantage of general folk theorems with minmax punishments, not just relying on Nash reversion.

A main challenge in repeated games (or opportunity, depending on the game) is the sheer size of the equilibrium set when the probability of future encounters is high. This amplifies the problem of equilibrium selection in the laboratory. The new, suggestion-based approach we use for Stage Game I has origins in other strands of game theory, but we believe it is especially natural in the context of repeated games. First, it offers a potential way to trim a hefty equilibrium set to focus on a question of interest. Second, repeated games provide a natural treatment and control. One can hold the suggestion constant across treatments, while varying the continuation probability (or fixed number of repetitions) to understand whether the suggested behavior is self-enforcing due to the repetition itself. As a methodological benefit, this keeps instructions and experimental procedures directly comparable across treatments.

Finally, we point out that while do not intend the two stage games introduced here as models of bullying, we suspect that the likelihood of repetition may play a role in that harmful phenomenon. While definitions of bullying differ, a commonly used definition would say 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). Bullying is an important public health issue, and its deleterious effects are well documented; see Hamburger et al. (2011) for a compendium of research. Continued research to better understand its origins, and how it can be prevented, is important.

References

Andreoni, J., W. Harbaugh, and L. Vesterlund (2008). Altruism in experiments. In *The New Palgrave Dictionary of Economics, 2nd edition.*

Benoit, J.-P. and V. Krishna (1985). Finitely repeated games. *Econometrica 53*(4), 905–22.

Brandts, J. and C. A. Holt (1992). An experimental test of equilibrium dominance in signaling games. *The American Economic Review 82*(5), 1350–1365.

Brandts, J. and W. B. Macleod (1995). Equilibrium selection in experimental games with recommended play. *Games and Economic Behavior 11*(1), 36–63.

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. G. 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.

Embrey, M., G. Fréchette, and S. Yuksel (2018). Cooperation in the finitely repeated prisoner's dilemma. *Quarterly Journal of Economics 133*, 509–551.

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

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

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.

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

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

van Huyck, J., A. B. Gillette, and R. Battalio (1992). Credible assignments in coordination games. *Games and Economic Behavior 4*(4), 606–626.

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.

Participants will act in one of two roles throughout this experiment, either as `low-index' or `high-index.' The computer will randomly determine and inform you of your assigned role at the beginning of the experiment, and you will stay in that role throughout. For reasons that will become clear shortly, 4 out every 7 participants will be assigned to the low-index role.

There are 30 rounds in this experiment. At the start of each round, participants are randomly rematched into groups of 7 members, 4 of which are low-index and 3 of which are high-index. Within each group of 7, low-index members will be randomly assigned a distinct index from 1 to 4, while high-index members will be randomly assigned a distinct index from 5 to 7. You will be told the index you have been assigned in each round.

You will only be identified to the other participants by your index in each round; your actual identity will never be revealed in the course of this experiment.

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

Each round will consist of at most two stages, Stage 1 and Stage 2, which are identical in structure. Once Stage 1 is over, the computer throws a (simulated) 100-sided die and Stage 2 occurs if a number lower than or equal to 99 comes up. Thus there is a 99% chance of a second stage. Everyone's index stays constant across the stages of a round. Your payoff in a round is the sum of the payoff(s) you accrued in the stage(s) of that round.

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

High-index group members will choose an index between 1 and 4, and get a $5 payment in that stage if their choice matches that of at least one other high-index member of their group.

Low-index group members will decide whether to pay a cost of $2 in that stage, which is deducted from their baseline stage-payoff of $13. Separately, their payoff is reduced by $10 in a stage where their index is chosen by at least one high-index member of their group.

At the end of each stage, all participants in a group will be presented with a table showing the choices each member of their group made.

<u>Suggestions on how to play</u>

At the start of each stage, all members of a group see a suggestion regarding how high-index members of that group could play in that stage. They are free to follow the suggestion or not. Payoffs only depend on the decisions made by group members, not on whether suggestions are being followed. We now explain how the suggestion is constructed.

In <u>Stage 1</u>, the suggestion is always for high-index group members to select index 1.

In Stage 2 (if it occurs), the suggestion depends on the Stage-1 choices. The computer will display the smallest index among those in the group who didn't pay the cost in Stage 1, and will suggest for all high-index members to pick that index in Stage 2. Note: if all low-index members paid the cost, the suggestion will be for high-index members to pick index 4.

What happens at the end of the experiment?

When 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. You will be paid the show-up fee plus your round payoff (that is, the sum of the payoff(s) you accrued in the stage(s) of that round). Again, all identities remain anonymous.

We are almost ready to start the experiment. First, there will be a short quiz to check your understanding of some key features of the experiment.

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.

Participants will act in one of two roles throughout this experiment, either as `low-index' or `high-index.' The computer will randomly determine and inform you of your assigned role at the beginning of the experiment, and you will stay in that role throughout. For reasons that will become clear shortly, 4 out every 7 participants will be assigned to the low-index role.

There are 30 rounds in this experiment. At the start of each round, participants are randomly rematched into groups of 7 members, 4 of which are low-index and 3 of which are high-index. Within each group of 7, low-index members will be randomly assigned a distinct index from 1 to 4, while high-index members will be randomly assigned a distinct index from 5 to 7. You will be told the index you have been assigned in each round.

You will only be identified to the other participants by your index in each round; your actual identity will never be revealed in the course of this experiment.

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

Each round will consist of at most two stages, Stage 1 and Stage 2, which are identical in structure. Once Stage 1 is over, the computer throws a (simulated) 100-sided die and Stage 2 occurs if the number 1 comes up. Thus there is a 1% chance of a second stage. Everyone's index stays constant across the stages of a round. Your payoff in a round is the sum of the payoff(s) you accrued in the stage(s) of that round.

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

High-index group members will choose an index between 1 and 4, and get a $5 payment in that stage if their choice matches that of at least one other high-index member of their group.

Low-index group members will decide whether to pay a cost of $2 in that stage, which is deducted from their baseline stage-payoff of $13. Separately, their payoff is reduced by $10 in a stage where their index is chosen by at least one high-index member of their group.

At the end of each stage, all participants in a group will be presented with a table showing the choices each member of their group made.

<u>Suggestions on how to play</u>

At the start of each stage, all members of a group see a suggestion regarding how high-index members of that group could play in that stage. They are free to follow the suggestion or not. Payoffs only depend on the decisions made by group members, not on whether suggestions are being followed. We now explain how the suggestion is constructed.

In <u>Stage 1</u>, the suggestion is always for high-index group members to select index 1.

In Stage 2 (if it occurs), the suggestion depends on the Stage-1 choices. The computer will display the smallest index among those in the group who didn't pay the cost in Stage 1, and will suggest for all high-index members to pick that index in Stage 2. Note: if all low-index members paid the cost, the suggestion will be for high-index members to pick index 4.

What happens at the end of the experiment?

When 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. You will be paid the show-up fee plus your round payoff (that is, the sum of the payoff(s) you accrued in the stage(s) of that round). Again, all identities remain anonymous.

We are almost ready to start the experiment. First, there will be a short quiz to check your understanding of some key features of the experiment.

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.

<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.

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 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.

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.

<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.