

4-8-2019

Do Economic Inequalities Affect Long-Run Cooperation & Prosperity?

Gabriele Camera

Chapman University, camera@chapman.edu

Cary Deck

Chapman University, deck@chapman.edu

David Porter

Chapman University, dporter@chapman.edu

Follow this and additional works at: https://digitalcommons.chapman.edu/esi_working_papers



Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Camera, G., Deck, C., & Porter, D. (2019). Do economic inequalities affect long-run cooperation & prosperity? ESI Working Paper 19-09. Retrieved from https://digitalcommons.chapman.edu/esi_working_papers/267/

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Do Economic Inequalities Affect Long-Run Cooperation & Prosperity?

Comments

Working Paper 19-09

This working paper was later published as:

Camera, G., Deck, C., & Porter, D. (2019). Do economic inequalities affect long-run cooperation and prosperity? *Experimental Economics*. <https://doi.org/10.1007/s10683-019-09610-5>

Do Economic Inequalities Affect Long-Run Cooperation & Prosperity?*

Gabriele Camera Cary Deck David Porter
Chapman University University of Alabama Chapman University
University of Bologna Chapman University

April 8, 2019

Abstract

We explore if fairness and inequality motivations affect cooperation in indefinitely repeated games. Each round, we randomly divided experimental participants into donor-recipient pairs. Donors could make a gift to recipients, and ex-ante earnings are highest when all donors give. Roles were randomly reassigned every period, which induced inequality in ex-post earnings. Theoretically, income-maximizing players do not have to condition on this inequality because it is payoff-irrelevant. Empirically, payoff-irrelevant inequality affected participants' ability to coordinate on efficient play: donors conditioned gifts on their own past roles and, with inequalities made visible, discriminated against those who were better off.

Keywords: cooperation, experiments, indefinitely repeated games, social dilemmas.
JEL codes: C70, C90, D03, E02

1 Introduction

There is a view that inequality—in opportunity, income, and wealth—undermines the long-run prosperity of a nation (Stiglitz, 2012). Two possible channels have been identified. One is purely economic: inequality creates distortions

* We thank Co-Editor M. C. Villeval and two anonymous referees for helpful suggestions that improved this study from an earlier version. We also thank N. Wilcox for helpful conversations, K. Bregu for help running the experiments and seminar participants at Chapman University, the Federal Reserve Bank of Cleveland, Humboldt University, the University of Basel, and the Theory and Experiments in Monetary Economics Conference at GMU. G. Camera acknowledges partial research support through the NSF grant CCF-1101627. Correspondence address: Gabriele Camera, Economic Science Institute, Chapman University, One University Dr., Orange, CA 92866; e-mail: camera@chapman.edu.

in capital and labor markets, which result in resource misallocation and efficiency declines (Aghion and Williamson, 1998; Piketty, 2014). The other is behavioral: inequality erodes cohesion and trust in society, which harms prosperity by undermining norms of cooperation and prosocial behavior (Putnam, 2000). We focus on this second channel, which has not been investigated very extensively; the impact of fairness and inequality motivations is still largely unexplored in the literature on long-run cooperation and social norms. We contribute to filling this gap through an experiment based on the theory of repeated games (or, supergames); see Table 1.

Table 1: Related experimental literature

	One-shot/Finite	Supergames
Equality	✓	✓
Inequality	✓	This study

In the experiment, a group of four subjects interacts as strangers. The game consists of an indefinite sequence of “helping games” in ever-changing pairs (Camera et al., 2013). In each pair, one party (the donor) is given the option to cooperate by providing a large benefit to the counterpart (the recipient), at a modest personal cost. The alternative is to defect by remaining idle, which grants modest and slightly unequal payoffs to both. This gives rise to a social dilemma because, though cooperation is socially efficient, a short-run temptation may exist to free ride by defecting.

A key design feature is that individuals switch from one role to the other at random. In each round, a virtual coin flip determines who is the donor in the pair. This random role assignment amounts to an uncontrollable shock to earning opportunities. It ensures equal opportunity *ex-ante* but not *ex-post* since, as the game progresses, someone may be a donor more often than others. These differences in *past* roles (or, opportunities) are the exogenous source of uncontrollable inequality. This kind of inequality is payoff-irrelevant because

it neither modifies the power structure in the game, nor alters the expected return from cooperation—which is independent of past opportunities.

In this game, efficient play is incentive-compatible if individuals adopt a norm of mutual assistance supported by a “grim” trigger. Every player should cooperate whenever she can, with anyone she meets, independent of the distribution of past roles. If anyone breaks this norm, then the players will cease all cooperation, effectively triggering the worst equilibrium (Abreu et al., 1990; Ellison, 1994; Kandori, 1992). This norm ensures 100% efficiency (ex-ante and ex-post) and guarantees equal and maximal ex-ante payoffs. Hence, under an efficiency criterion, full cooperation is a focal outcome. But efficient play cannot guarantee equal payoffs *ex-post* because some players may be frequent recipients who enjoyed more of the cooperation benefits while shouldering less of its costs (*advantaged* players), while the reverse holds true for others (*disadvantaged* players). According to the theory of repeated games, variation in past roles should not affect the structure of incentives, so income-maximizing players do not have to condition their choices on information about past roles or earnings. But if individuals are driven by a mix of motives, including fairness or inequality aversion, then they might act upon such information.

Our first contribution is to document that variation in past roles hindered coordination on efficient play. Donors conditioned their choices on *their own* past roles, acting less cooperatively when disadvantaged. This reveals that payoff-irrelevant inequalities interfered with the incentives for cooperation, undermining coordination on efficient play. Second, we show that disclosing information compounds this problem. In a treatment, prior to choosing, we informed donors of how their past opportunities compared to the recipient’s. This information made salient inequalities in past roles without theoretically altering the structure of economic incentives compared to the no-information treatment. It neither revealed past conduct or future intentions, nor expanded the payoff set, nor disclosed identities. Empirically, disclosing information on

past roles in the group affected behavior and cooperation. Donors conditioned their choices on the past roles of their counterparts, discriminating against recipients known to be advantaged in terms of past opportunities. We also see a greater inclination to unilaterally break a norm of mutual support by defecting even if no one else in the group had yet done so. This may reflect players' attempts to raise their own payoff to counteract unfavorable past opportunities. Yet, this behavior backfired because people reacted to defections by permanently lowering their cooperation rate. As a result, realized efficiency and coordination on efficient play declined as compared to the treatment where payoff-irrelevant inequalities remained hidden.

Since inequalities in past roles are positively correlated with inequalities in past earnings, as a robustness check we also investigate the case where past earnings inequalities are made explicit. In two additional treatments, before taking an action, donors could compare the running total of their earnings to the recipient's. Earnings information is crossed with role information resulting in a two-by-two design. Here, too, we see a decline in pro-social behavior as compared to the no-information treatment. Donors actively discriminated against more fortunate players, and were also less likely to stick to a cooperative norm when no one else had defected. However, we also see an increased desire to help the less fortunate, so that, although coordination on efficient play suffered, there is a minor decline in realized efficiency. The fact that payoff-irrelevant shocks affected behavior in our experiment suggests care is needed when introducing exogenous variation in laboratory supergames (e.g., see Camera and Casari, 2014; Camera et al., 2013). The message is that random elements of the design may affect the incentives for cooperation biasing the results in theoretically unanticipated ways. This has two distinct possible implications. The first one concerns economic policies that are inherently redistributive, such as fiscal policy. The experiment suggests such policies should be evaluated not only in terms of their effectiveness at meeting the pol-

ity objectives, but also their potential impact on society’s cohesion. A second potential implication concerns the social impact of inequality. Much of the economics literature has focused on inequality’s potential to lower society’s welfare by misshaping the structure of economic incentives. Our study suggests that inequality alters economic behavior even when it leaves the economic incentives unaltered, as it inhibits individuals’ pro-sociality and cooperative attitudes.

We proceed as follows. Section 2 discusses the related experimental literature. Section 3 describes the design. Section 4 presents theory and hypotheses. Section 5 reports the results. Section 6 offers some final considerations.

2 Related studies

Our design is mainly related to experimental studies of cooperation with repeated play and, in particular, to indefinitely repeated social dilemmas, which support a richer set of equilibria compared to games with one-shot or a commonly known number of rounds. These experiments have mainly focused on fixed pairs facing a Prisoner’s Dilemma (Dal Bó and Fréchet, 2018). This setup allows for reciprocity and it also ensures that earnings are equal under full cooperation, ex-ante and ex-post, because subjects have identical opportunities to give and receive benefits. By contrast, we adopt a design as in Camera and Casari (2014); Camera et al. (2013), which rules out reciprocity through random re-matching, and where efficient play cannot guarantee equal earnings ex-post because subjects have random opportunities to give or receive benefits. We extend this earlier design by informing subjects about differences in past opportunities or realized incomes in the group, to determine if and how such information about inequalities impacts cooperation and realized efficiency.

The experimental literature suggests that individuals do care about equality in outcomes, at least to some degree, in strategic settings where players

face a social dilemma.¹ Three features of our design differentiate it from those in earlier experiments. On the one hand, earlier designs focus on games that are one-shot or have a commonly known number of rounds. There, behavioral motivations cannot undermine the economic incentives to play efficiently, because—unless one explicitly considers heterogeneity or introduces social components in preferences—there are none, so efficient play is not part of a Nash equilibrium (e.g., Anderson et al., 2006; Greiner et al., 2012; Güth et al., 2003). By contrast, we adopt an indefinite-horizon design that supports a richer set of equilibria because—even if players are homogeneous and self-interested—they can exploit the dynamic structure of the game to support efficient play. Second, previous experiments focus on partners designs with perfect information about past conducts, or with communication facilitating coordination of play (e.g., Nishi et al., 2015; Tavoni et al., 2011). By contrast, we adopt a design where individual past conducts are opaque and there is no communication so efficient play cannot be incentivized through reciprocity, and individuals must tacitly coordinate on a community-wide norm of cooperation and of punishment. Third, inequality does not alter the return from efficient play (e.g., as in Gangadharan et al., 2015; Sadrieh and Verbon, 2006) because it is payoff-irrelevant. This allows us to decouple behavioral and economic effects of inequality, to investigate if inequality that should not theoretically alter behavior pushes subjects away from efficient play.²

Fairness and inequality have also been investigated in non-strategic distributive choice experiments where a disinterested third-party (a spectator) must select a division of resources between two others. This research finds support for “luck egalitarianism,” the notion that actions are meant to smooth

¹There is a large literature about how to incorporate fairness and inequality aversion into economic models (e.g., Bolton and Ockenfels, 2000; Fehr and Schmidt, 1999; Rabin, 1993) and although our study is related to it, it is not an experiment about testing these models in the lab (e.g., Blanco et al., 2011; Deck, 2001; Kagel and Wiley-Wolfe, 2001).

²In a bargaining experiment, Goeree and Holt (2000) find that differences in fixed payments that should not theoretically alter behavior induce offers inconsistent with Nash equilibrium but consistent with a fair division of final payments. See also Andreoni and Varian (1999).

out outcome differences due to uncontrollable factors (Konow, 2000; Mollerstrom et al., 2015). In our study the role assignment is a form of uncontrollable inequality and, although there are no disinterested spectators, we do find evidence consistent with luck egalitarianism.

Finally, by manipulating the amount of information across treatments we contribute to an experimental literature documenting how sometimes less, not more, information is beneficial in market experiments and strategic bargaining games (see Smith, 1994, p. 119) and also when the added information is theoretically irrelevant. For example, there is evidence that irrelevant information can influence decision-makers through information overload (O’Reilly, 1977), or by affecting inference judgments Troutman and Shanteau (1977). In our study, information is payoff-irrelevant, cannot reveal past conduct or individual valuations, nor can it be used to build reputations. We find that the less informed players are about the distribution of past earning opportunities, the easier it is for them to coordinate on the efficient equilibrium. This is not predicted by the standard application of folk theorem-type results to groups of strangers, as the efficient outcome is equally attainable in all informational settings.

3 Experimental design

In our experiment, four subjects face an indefinite sequence of “helping games.”

Interaction in a round: In each round of the experiment subjects face a helping game. Each game consists of a “donor” who is endowed with a good and a “recipient” who values the good more than the donor. The donor faces an individual decision problem: she can transfer the good to the recipient (cooperate, C), or she can consume the good (defect, D). The recipient has no endowment and no action to take. All framing in the experiment was neutral. Round payoffs are in Table 2; instructions are in Supplementary Information.

Table 2: Payoffs in a meeting (donor, recipient)

Donor's choice	
<i>C</i>	<i>D</i>
0, 25	6, 4

If the donor chooses C, then the recipient earns $g = 25$ points, while the donor earns nothing. Otherwise, both subjects earn a default payoff, which is $d = 6$ for the donor and $d - l = 4$ for the recipient, so $g > 2d - l > 0$. Given this, the donor's dominant action is to do nothing. Cooperation is not mutually beneficial but it maximizes surplus in the pair, so it is socially optimal. The surplus from cooperation is $g - (2d - l) = 15$ points. The cost of cooperation to a donor is the payoff difference $0 - d = -6$; the cooperation benefit to a recipient is her surplus $g - (d - l) = 21$; hence the benefit/cost ratio is 3.5.

Baseline session and supergame: A session involved 16 subjects in the lab at the same time, all exposed to the same treatment, which was chronologically divided into five distinct supergames. In a supergame, subjects interacted for an indefinite number of rounds in fixed matching groups of size four. Hence, there were four concurrent supergames being played in a session. In each group, subjects were randomly re-matched into pairs at the start of each round, so there was a 1/3 probability of meeting the same person in two consecutive rounds. Subjects did not know with whom they were paired nor did they know who was in their matching group in any supergame.

Every round, in each pair the computer randomly assigned the recipient role to one subject ("blue," in the experiment), and the donor role to the other ("red"), with equal probability. Hence, in every round half the subjects were recipients and half were donors. The random assignment of roles is a shock that affects the subject's earning potential for the round because recipients have a superior earning potential (25 points vs 6 points). This shock ensures equal economic opportunity going forward because the payoff matrix

and role assignment process are fixed. As a result, in a fully cooperative outcome future earning prospects are identical across individuals and rounds, and are completely unaffected by differences in past roles. Yet, the random assignment of roles provides an exogenous source of variation in cumulative earnings, and so it is likely to generate unequal economic results over the course of the supergame. As the supergame progressed, some participants could be *advantaged*, having been recipients more often than others (*disadvantaged*), thus having more chances of getting the higher payoff of 25 if cooperation occurred.

The duration of the supergame was uncertain because it was determined by a random continuation rule (Roth and Murnighan, 1978). A supergame began with 15 fixed rounds after which successive rounds occurred with probability $\beta = 0.75$. This continuation probability can be interpreted as the discount factor of a risk-neutral subject. *A priori*, the expected duration of a supergame was 18 rounds because from round 15, in each round the supergame is expected to last 3 more rounds. At the end of each round a computer drew an integer number between 1 and 100 with equal probability, which was then revealed to all subjects. A draw equal to or below 75 informed subjects that the supergame would continue (otherwise, it would end).³

At the end of each round, subjects observed whether or not the outcomes were identical in both pairs of their group. This form of anonymous public monitoring allows public detection of deviations from a social norm, and it could also simplify coordination tasks, but it does not allow agents to identify opponents (see screenshots from instructions in Appendix C). Hence, because individual histories remained private, subjects could neither build a reputation nor engage in relational contracting. Public monitoring ensured that the minimum discount factor supporting full cooperation in equilibrium was invariant across treatments (see next section).⁴

Supergames terminated simultaneously for all concurrent groups. After

³This number could also serve as a public coordination device, at the group level.

⁴Subjects had access to a record of their own past outcomes, and also had pen and paper.

each of the first four supergames, subjects were placed into new four person matching groups and began playing another supergame. Matching groups were constructed so that no one was ever in a group with someone else more than once, a perfect strangers design. Subjects were aware of this fact and, as a result, we have twenty unique groups per session. At the conclusion of a session, one supergame was selected randomly (Sherstyuk et al., 2013) and subjects were paid based upon their earnings in that supergame at the rate of \$0.20 per point. The specifics of the procedural details are discussed in Appendix A. Here we note that the experiment involved 16 sessions and 256 subjects with average salient earnings of \$27.00 per subject; see Table 3.

Other treatments: We adopt a 2x2 design where one factor is the donors’ knowledge of relative past roles and the other is knowledge of relative past earnings. In all treatments the socially optimal outcome is an equilibrium, subjects cannot build reputations or engage in reciprocity, and past actions cannot affect future opportunities or feasible outcomes; see Table 3.

Table 3: Overview of the four treatments and the sixteen sessions

Variable	Treatment			
	<i>Baseline</i>	<i>Roles</i>	<i>Wealth</i>	<i>History</i>
Blue Index	No	Yes	No	Yes
Earnings Index	No	No	Yes	Yes
Subjects/Sessions	64/4	64/4	64/4	64/4
Superg./Avg. rounds	80/18.5	80/18.4	80/18.5	80/19.6
Salient \$ Earnings				
average	26.38	25.94	25.00	30.08
min, max	8.75, 54.00	6.50, 55.50	9.25, 54.00	10.00, 54.00

In the *Roles* treatment, donors observed explicit information about past roles before making their choice. At the start of every round after the first, we calculated the proportion of past rounds in which each subject was a recipient. Before making a choice, donors observed the normalized recipient rate for each group member. This information was called the “blue index” as it measures how often players had been in the blue role. Index values masked identities

and preserved anonymity because they were not associated with individual identifiers, and were unobservable to recipients.⁵ Unequal blue indices give rise to inequality in past earnings, especially in cooperative groups, given the large point spread. Hence, as a robustness check we ran two additional treatments. In *Wealth*, donors observed the distribution of the subjects’ running total of points earned (wealth). Since the mean value varied from round to round, this information was presented in relative form, with the mean index normalized to 100 (“earnings index”). A donor observed her own relative wealth, that of the matched recipient, and of the other two group members. In *History*, donors saw both the “blue” and the “earnings” indices.

Remarks: In all treatments there is “equal opportunity.” Players’ ex-ante earnings potential is governed by a payoff matrix and a role assignment process that are fixed, identical across players, and independent of past roles and actions. Ex-post inequality in roles is thus payoff-irrelevant. Adding the indices expands the strategy set—as donors can condition their choice on the provided information in earlier rounds—but neither expands the action set relative to the no information *Baseline*, nor affects payoffs in the stage game.

With high levels of inequality in an index, a donor might be able to determine if they previously met that person only as a recipient-counterpart, but not as a donor-counterpart—thus preventing direct reciprocity. Moreover, since indices do not disclose information about actions undertaken by that person, they make past conduct opaque and hinder reputation-building. The “blue index” prevents reputation-building because role assignment is independent of past conduct and the donors’ index is hidden from recipients. The “earnings index” might offer a noisy reputation signal, as it is correlated with past conduct. If so this allows targeting punishment to specific individuals; e.g., high-wealth players, if high-wealth is considered statistical evidence of

⁵The donor observed her blue index, the paired recipient’s index, and the index of the two others in a random order. To facilitate comparisons, the average relative frequency of 0.50 was normalized to 100, so a value of $100 + x$ indicated a $x\%$ departure from the average.

free-riding behavior. It is thus conceivable that the information provided in *Wealth* and *History* facilitates coordination on efficient play as compared to the other two treatments, by allowing identification and sanction of free-riders without requiring coordination on the grim, community punishment.

Full cooperation supports income inequality ex-post because the realized sequences of donor and recipient roles inherently vary across subjects, over time. These disparities are an exogenous source of earnings variation in the experiment, which is uncontrollable and does not alter the structure of economic incentives because it does not affect continuation payoffs in the efficient equilibrium. In a cooperative group it is the only source of earnings variation. Yet, ex-post income or roles inequality cannot alter the power structure in the game as high-income players have no greater control over the earnings of others than low-income players.

4 Theoretical considerations

In our setup, full cooperation is efficient and benefits everyone. To demonstrate that groups can attain this outcome, we employ standard arguments from the theory of supergames, showing that players can exploit anonymous public monitoring to build an incentive-compatible rule of cooperation supported by a punishment convention that is triggered if the rule gets broken.

A theoretical premise is that players are risk-neutral and payoff-maximizing, where payoffs correspond to players' *expected* earnings. Payoffs can be calculated at the start of any round, as the present value of the anticipated stream of earnings in the continuation game, using the continuation probability β as the discount factor. Payoffs are thus ex-ante measures of performance that only include future anticipated earnings, and should not be confused with ex-post measures including performance in past rounds (e.g., total earnings in the game). Inequalities in past roles or incomes are thus theoretically *payoff-irrelevant*: they can neither affect the earnings matrix in Table 2, nor the

matching or role assignment processes in future rounds. It follows that though variation in past roles or incomes affects total earnings in the game, it does not affect the theoretical structure of incentives of payoff-maximizing players.

To begin the analysis, start by observing that defection is the unique Nash equilibrium in a one-shot interaction because cooperation is costly to a donor. It follows that full defection is a sequential Nash equilibrium because it consists of an indefinite repetition of the one-shot Nash equilibrium. However, payoffs are minimized under full defection. Instead, full cooperation maximizes payoffs. Full cooperation can be supported as a (sequential) equilibrium if a subject cooperates from the beginning of the game whenever she is a donor, but switches to defection forever after someone defects (Kandori, 1992, Proposition 1). When everyone adopts this strategy, then we say that cooperation is a *social norm*. Here, any defection is punished with permanent defection by the entire group.

Proposition 1. *In the experiment, the equilibrium set includes full cooperation & full defection.*

Cooperation is an equilibrium when two conditions apply: in equilibrium, every donor prefers to choose C; out of equilibrium *no* donor prefers to chose C. The latter condition is immediately verified: once someone moves off equilibrium, that deviation is publicly observed. Hence, every donor defects thereafter and there is no longer an incentive to cooperate. The first condition requires checking that a donor cannot improve her payoff by moving off equilibrium (unimprovability criterion). In Appendix A we prove that this is the case as long as players are sufficiently patient, if $\beta \geq \beta^* := \frac{2d}{g+l}$. In the experiment $\beta^* = 4/9$, and since the continuation probability of the game corresponds to $\beta = 0.75$ under risk neutrality, cooperation is an equilibrium in every treatment.⁶ The threshold value β^* is the ratio between the cost of cooperation for

⁶Full cooperation cannot be ruled out as an equilibrium under empirically reasonable risk aversion. The coefficient β^* depends on the assumption of linear preferences. One can show that with CRRA preferences of the type $u^{1-\gamma}/(1-\gamma)$ full cooperation remains an

a donor d and the surplus difference $\frac{g+l}{2}$ expected next round.

If cooperation is a social norm, then strategies and therefore payoffs are independent of uncontrollable inequalities in past roles or incomes, both in- and off-equilibrium. Considering for simplicity rounds when the random stopping rule has already begun (for general rounds see Appendix A), for any realization of past roles, the equilibrium payoffs to donor and recipient are, respectively,

$$a + \frac{\beta(g+a)}{2(1-\beta)} \quad \text{and} \quad g + \frac{\beta(g+a)}{2(1-\beta)}.$$

Off-equilibrium, instead, there is full defection so for any realization of past roles donor and recipient's payoffs are $d + \frac{\beta(2d-l)}{2(1-\beta)}$ and $d-l + \frac{\beta(2d-l)}{2(1-\beta)}$.

Proposition 2. *Adding inequality indicators eliminates none of the Baseline equilibria, and does not expand the set of equilibrium payoffs. Conditioning on indicators is neither necessary nor sufficient to attain efficiency.*

The possibility to condition behavior on inequality indices increases the set of available strategies compared to *Baseline*, hence can alter the equilibrium set. However, adding inequality indices cannot eliminate any of the equilibria possible in *Baseline* since players can always rely on strategies that ignore inequality indices. Moreover, adding inequality indices does not expand the set of equilibrium payoffs because the efficient outcome is an equilibrium in *Baseline*. The use of conditional strategies is not necessary to sustain full cooperation because defections are publicly revealed, so the efficient outcome can be attained in all treatments by exploiting anonymous public monitoring. The use of conditional strategies is not sufficient for efficient play because the indices mask the identity of counterparts, cannot be used to signal a cooperative intention, and do not reveal individual past conduct.

equilibrium if $\gamma \leq 0.37$. Estimates of CRRA coefficients vary widely depending on many factors; in experiments with a fixed recruitment fee, such as ours, we find a coefficient of 0.34 (Harrison et al., 2009), while in experiments with low stakes 60% of subjects lay below 0.41 (Holt and Laury, 2002).

Summing up, players can maximize their prospective payoffs by adopting a strategy that threatens a switch to uncooperative behavior in response to defections. The structure of incentives remains unaltered as we add the payoff-irrelevant information offered by the blue and earnings indices. Hence, there is no obvious reason to expect different behavior across treatments. Conditioning actions on information about past roles or earnings does not enhance prospective outcomes. Payoff maximizing players can easily coordinate on efficient play in all treatment by relying on public monitoring of defections in their four-person group. Conditioning actions on a (privately observed) statistic about past roles or earnings does not go in the direction of facilitating coordination on cooperation, and in fact may only contribute to increase coordination complexity. We thus put forward three testable hypotheses:

H 1. *Donors should condition actions on past play.*

H 2. *Donors should not condition actions on own past roles/incomes.*

H 3. *Donors should not condition actions on others' past roles/incomes.*

If, instead, inequality concerns matter, then individuals may condition actions on past roles/incomes.

The requirement $\beta \geq \beta^*$ is necessary and sufficient for the existence of a cooperative equilibrium, but there is no guarantee that full cooperation will emerge instead of a lower-efficiency equilibrium. Equilibrium selection is an open question because many other equilibria exist with efficiency degrees below 100%. So why should we pay particular attention to full cooperation? First, it is Pareto dominant and so it is a natural equilibrium to coordinate on, for income-maximizing players. Second, the grim strategy is risk dominant in the experiment (see proof in Appendix A), so strategic uncertainty has less of a bite. Third, experiments on indefinitely repeated social dilemmas reveal that subjects tend to select play that aims for efficiency not only in fixed-pair designs (Camera et al., 2013), but also under random-matching design with

non-anonymous public monitoring (Camera and Casari, 2009); there is no reason to believe this tendency should shift under anonymous public monitoring. Fourth, previous experiments on helping games reveal that subjects who have sufficient exposure to the game, and access to institutions that minimize exploitation risks, tend to coordinate on efficient play even in much larger groups (Bigoni et al., forth.).

5 Results

This section starts by documenting behavior when inequalities remained hidden (*Baseline* treatment), proceeds by studying behavior when past roles were revealed (*Roles* treatment), and finally discusses the robustness of these findings to revealing wealth inequalities (*Wealth* and *History* treatments). Throughout this section we will refer to Table 4 below, which provides an aggregate view of average cooperation in the groups across treatments and supergames. By design, realized efficiency in a group corresponds to the mean cooperation rate of in the group.

Table 4: Average Cooperation and Coordination on Cooperation

Treatment	Supergame					Overall	t=1	Coordination on C		
	1	2	3	4	5			≤20%	≥80%	100%
Baseline	0.44	0.51	0.62	0.67	0.62	0.57	0.63	9	21	10
Roles	0.31	0.49	0.51	0.53	0.54	0.48	0.59	14	11	2
Wealth	0.38	0.51	0.49	0.57	0.57	0.50	0.53	12	12	4
History	0.42	0.52	0.58	0.55	0.51	0.51	0.62	10	13	1

Notes: 1 obs. = four-person group in a supergame (N=80 per treatment). *Supergame* columns: average proportion of cooperative choices in a supergame (the standard errors, which are not reported, vary between 0.04 and 0.08). *t=1* column: average cooperation in round 1 of all supergames. *Coordination on C*: number of groups that attained a given cooperation level. Only one group achieved 0% cooperation (supergame 1, *Wealth* treatment).

Consider the *Baseline* sessions in which inequalities remained hidden.

Result 1. *In Baseline, average cooperation increased with experience. No group coordinated on the inefficient equilibrium, and 10 in 80 groups coordinated on the efficient equilibrium.*

Evidence is provided in Tables 4 and 5. In Baseline, the average cooperation rate in a supergame lies between 44 and 67 percent. Cooperation significantly increased as subjects gained experienced with the task. Table 5, reports marginal effects on the cooperation rate in the average group from a regression on Baseline data.⁷

Table 5: Cooperation in *Baseline*: Marginal Effects

Dep. variable:	Coefficient	S.D.
Cooperation rate		
Supergame	0.058***	(0.014)
<i>Controls</i>		
Male	-0.053	(0.063)
Duration	-0.047***	(0.016)
Previous duration	0.002	(0.012)
Response Time	-0.059**	(0.024)
Incorrect Answers	-0.006	(0.018)
Risk attitude	0.185***	(0.072)
N	80	

Notes: GLM Regression: the dependent variable is the relative frequency of cooperation in a group in a supergame (N=80 per treatment). *Controls* include standardized values of supergame duration, current and previous (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), and self reported measures of sex and attitudes toward risk. Marginal effects are computed at the mean value of regressors of continuous variables. Robust standard errors (S.E.) adjusted for clustering at the session level. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

The regression includes a standard set of individual and other controls (e.g., subject’s self-reported sex and duration of the supergame, see notes to Table). It also includes a supergame regressor to trace how experience with

⁷For a continuous variable, the marginal effect measures the change in the likelihood to cooperate for an infinitesimal change of the independent variable. For a dummy variable, the marginal effect measures the change in the likelihood to cooperate for a discrete change of the dummy variable from its base level (0).

the task affects cooperation; its coefficient is positive and significant.⁸ This evidence suggests that subjects do not frequently attain high efficiency levels unless they are repeatedly exposed to the game; it is in accordance with data from previous indefinitely repeated helping games (Camera et al., 2013) as well as PD games in fixed pairs (Dal Bó and Fréchette, 2018). It is in contrast with the dynamics of cooperation observed under deterministic horizons, as in that case cooperation tends to fall as subjects gain experience with the game (Dal Bó, 2005; Palfrey et al., 1994).

As our design admits multiple Pareto-ranked equilibria, it is possible that the average cooperation rate observed is the result of different groups coordinating on different equilibria. As an example, if 46 groups coordinate on full cooperation and the rest on full defection, then we obtain 56 percent average cooperation. The data do not support this conjecture. About one in four groups attained a cooperation rate of 80% or more, half of which fully cooperated. By contrast, less than one in ten groups attained a cooperation rate of 20% or less, and no group fully defected. A similar picture emerges when we measure subjects' cooperation rates, i.e., the proportion of cooperative choices taken as a donor in a supergame. One in three subjects cooperated every single time they were a donor, one in seven never cooperated, and the majority of subjects falls somewhere in between (N=320). This suggests a general tendency to seek high-payoff outcomes, but also an inherent difficulty to coordinate on efficient play.

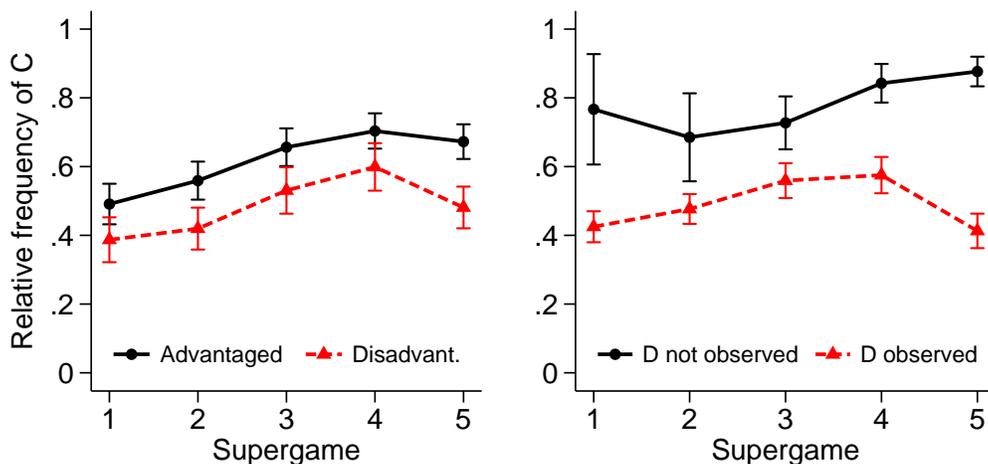
The data reveal that the payoff-irrelevant variation in past roles is partly responsible for these coordination issues, because disadvantaged players were less likely to cooperate. This impaired coordination on efficient play because defections triggered a long-lasting punishment response.

⁸Cooperation increased over the first three supergames, and then stabilized. In a regression with a factor variable tracing the impact of each supergame we can reject the hypothesis that the coefficient on supergame 2 is statistically similar to those on supergames 4 and 5. All other pairwise coefficient comparisons indicate statistical similarity in coefficients (p-value ranges from 0.541 to 0.829).

Result 2. *In Baseline, donors conditioned choices on past play, but also on their own past roles.*

The first part of Result 2 is consistent with the theoretical notion that defections trigger a long-lasting sanction, and is in line with H1. The second part is not, and leads us to reject H2. Figure 1 and Table 6 proved evidence.

Figure 1: Cooperation conditional on actions (right) and own past roles (left)



Notes: One obs.=one donor in a round > 1 , *Baseline* treatment. Each point is the average proportion of cooperative choices in the supergame, with 95% confidence intervals.

We trace a subject's roles history in round $t > 1$ of a supergame using the *blue index*, which corresponds to the relative frequency of the subject's past recipient roles in that supergame. This frequency is the exogenous source of variation in the experiment and it can be calculated in any treatment, even when hidden from the subjects (as in *Baseline*). The average blue index is set to 100 in each round. For $t > 1$, we classify a subject as *disadvantaged* if she has a blue index strictly below the 100 average; otherwise, the subject is *advantaged*. Based on this classification, we have four possible types of meetings, depending on the classification of donor and recipient.⁹

⁹The inclusion of the cutoff point 100 in the advantaged definition implies we have less than

The left panel in Fig. 1 reports the average cooperation rate of advantaged and disadvantaged donors in *Baseline* supergames (initial round excluded). Overall advantaged donors cooperated more than disadvantaged donors.

We ran a panel logit regression with random effects to determine if these differences are statistically significant. The dependent variable takes value 1 if the subject cooperated as a donor in a meeting, and is 0 otherwise. The panel variable is a subject in a session and we consider all rounds in a supergame after the initial one. In *Baseline*, donors could not observe individual histories but could see outcomes in their group. To determine how players responded to defections, we include dummy variables that trace the impact of observing D in the group, on the subsequent probability of choosing C, as done in Camera and Casari (2014). The *Grim trigger* regressor takes value 1 starting the round after the subject directly experiences or observes a defection for the first time (and is 0 otherwise). As subjects might delay their punishment response, we also include four *choice n* dummy variables. The *choice n* dummy takes the value 1 when the subject is a donor for the n^{th} time after experiencing or seeing the initial D in the group (0 otherwise).¹⁰

To determine if a donor conditioned actions on her own role history, we include the factor variable *Donor-Recipient Meeting* to classify meetings based on the advantage/disadvantage of donor and recipient (the exogenous source of

50% disadvantaged subjects in the data. Overall, 41% of subjects can be classified as disadvantaged and 59% as advantaged. For example, in round 9 of *Baseline*, the average subject has been a recipient in 4 of the previous 8 rounds; a subject who was a recipient in less than 4 rounds would have a blue index below 100 and would be classified as disadvantaged. Table B2 in Supplementary Information reports the distribution of the four possible types of meetings for all treatments pooled together.

¹⁰The sum of the coefficients on *Grim trigger* and each *choice* dummy identifies the average donor's reaction to a defection in her group, on the first and second occasion she had to react. The *Grim trigger* coefficient captures the long-run response. Subjects made choices at random points in time so these regressors trace an individual's behavior on the first two occasions in which she can make a choice, after suffering or observing an initial defection. As subjects were donors on average every two rounds, this traces the subjects' response between two and four rounds after the defection. Empirically, the first opportunity to react to an observed defection occurs, on average, in round 4. For a detailed discussion on this econometric technique see Camera and Casari (2014).

variation in the experiment). This can be done even if the information about past roles was hidden from the subjects in this treatment. Based on this classification, we have four possible types of meetings. The *Donor-Recipient Meeting* regressor takes the value 0 if both donor and recipient were disadvantaged (meeting DD), 1 if the donor was disadvantaged and the recipient was not (meeting DA), 2 in the reverse scenario (meeting AD), and 3 if both were advantaged (meeting AA) which we take as the base in the regression. Finally, we use a series of supergame dummies (supergame 1 is the base), a series of dummies to control for round-fixed effects, and we include the standard set of individual controls. Table 6 reports the marginal effects on the donor’s probability of cooperating.

Table 6: Past roles are hidden: marginal effects

Dep. var. =1 if C chosen	Coeff.	S.E.
<i>Donor-Recipient Meeting</i>		
DD	-0.063**	(0.026)
DA	-0.091***	(0.025)
AD	-0.002	(0.018)
<i>Punishment regressors</i>		
Grim trigger	-0.303***	(0.046)
Choice 1	0.140***	(0.047)
Choice 2	0.093***	(0.036)
Choice 3	0.061**	(0.031)
Choice 4	0.022	(0.028)
Supergame	0.023**	(0.011)
N	2672	

Notes: Logit panel regression with random effects at the individual level and robust standard errors (S.E.) adjusted for clustering at the session level (*Baseline* data only). Dependent variable = 1 if C chosen, 0 otherwise. One observation = choice of a donor in a round > 1. Base case = donor and recipient are both advantaged (AA meeting). *Controls* include round fixed effects through a series of dummy variables (a single dummy variable for rounds 19 and above), duration of previous supergame (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), and a self-reported measure of sex and risk attitudes. Marginal effects are computed at the mean value of regressors of continuous variables. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

The average donor conditioned her choice on past play in her group. Suf-

fering or observing a defection led to a permanent decline in cooperation. The *Grim trigger* coefficient is negative and highly significant, and the *choice n* coefficients decline in n , indicating an increasingly negative response to defections. The partial sums of *Grim trigger* and each *Choice n* coefficient is negative and significant (Wald tests, not reported). Hence, we cannot reject H1. However, in contrast with the arguments used to construct optimal strategies of payoff-maximizing players, donors based their choices also on *their own* role history. The coefficients on the DD and DA regressors are both negative and highly significant; they are also statistically similar (two-sided Wald test, p-value=0.142) and statistically different than the AD coefficient (two-sided Wald tests, p-values < 0.001 each), which is insignificant. The similarity of DD and DA confirms what we expected: in *Baseline* subjects cannot distinguish between advantaged and disadvantaged recipient, because past roles are private information, so there should not be differences in behavior. This same reason explains why AD is insignificant (AA is the base in the regression).

This is evidence that donors with infrequent past opportunities to receive a benefit, cooperated significantly less than the rest. Hence, H2 is rejected. An interpretation is that subjects acted to reduce their own exposure to unfavorable earning opportunities due to past role assignments. They cooperated less when they had few past chances to benefit from the cooperation of others. This behavior is inconsistent with expected payoff maximization because, as we have seen, subjects punish defections by cooperating less in subsequent rounds. Hence, conditioning on own past roles can only reduce the future chances to attain high-payoffs. The open question is what happens when donors can compare their role history to that of others, before making a choice.

5.1 Past roles are observable

In *Roles*, donors saw the counterparts' relative frequency of past roles (their *blue index*) before making a choice. These disclosures made salient inequalities in past role assignments, allowing easy interpersonal comparisons of relative

positions (differences from a mean of 100). Yet, these disclosures do not alter the incentives' structure compared to *Baseline* as they neither reveal the counterparts' past conduct, nor their future intentions. If individuals are seeking to coordinate on efficient play then, given standard theoretical arguments, these disclosures should neither affect behavior nor outcomes. Payoff-maximizing players should not condition their actions on the past roles of others (H3).

Result 3. *In Roles, efficiency and coordination on full cooperation declined relative to Baseline.*

Support comes from Tables 4 and 7, concerning group-level analyses. Average cooperation is lower in *Roles* relative to *Baseline* in each supergame (see Table 4). These cooperation differences are significant according to a regression similar to that in Table 5, which pools data from both treatments. The marginal effects are reported in Model 1 of Table 7.

Table 7: Cooperation: marginal effects

	Model 1	Model 2
Dep. var. =	Coop.	Full Coop.
Roles dummy	-0.110*	-0.103***
	(0.059)	(0.035)
Supergame	0.059***	0.071***
	(0.012)	(0.008)
Controls	Yes	Yes
N	160	160

Notes: One observation is a group in a supergame (N=80 per treatment). Model 1: GLM Regression; the dependent variable is the relative frequency of cooperation in a group. Model 2: Logit regression; the dependent variable = 1 if group attained 100% cooperation, 0 otherwise. The regressions include interaction terms between treatment and supergame; *Controls* include supergame duration, current and previous (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), and a self-reported measure of sex and of risk attitudes. Robust standard errors (S.E.) adjusted for clustering at the session level. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

The regression includes a treatment dummy (*Baseline* serves as the base), and a continuous supergame regressor to trace the effect of experience that

is interacted with the treatment dummy. The *Roles* dummy is negative and significantly different from zero (p-value=0.062). Full coordination on efficient play is also less frequent in *Roles* than *Baseline* (2 vs. 10 groups). This difference is statistically significant according to a logit regression. Model 2 in Table 7 presents the marginal effects of a logit regression where the dependent variable takes value 1 if a group fully cooperated (and 0 otherwise). The coefficient on the *Roles* dummy is negative and significant (p-value = 0.004).¹¹

What lies behind this cooperation decline? Earning opportunities (roles) were randomly assigned by a computer program. Though this procedure guarantees ex-ante balance, there could still be ex-post imbalance in realizations.¹² A conjecture is that donors acted more uncooperatively in *Roles* partly out of a desire to reduce income inequality in their group. We find no support for this conjecture. Coordinating on full defection, which minimizes per-capita income, would have eliminated almost all income inequality because recipients and donors' round earnings are similarly low (4 vs. 6). Yet, no group did so in the *Roles* treatment. Moreover, income inequality—as measured by the Gini coefficient—was not lower in *Roles* than in *Baseline* (0.143 vs. 0.121).¹³

A second conjecture is that disclosing extant inequalities in past roles influenced behavior by making more salient any differences in past economic opportunities. If so, then donors might have conditioned their choices not only on their own past roles (as it happened in *Baseline*), but also on others' past roles, pushing players further away from coordination on efficient play.

¹¹The regressions account for the lack of independence within a session. Statistical analysis based on aggregating data at the session level gives us only four independent observations per treatment and is unsurprisingly inconclusive (data not reported).

¹²Table B1 in Supplementary Information reports statistics on inequality in opportunities experienced by donors, and its evolution over the supergame.

¹³One observation is the Gini measure for one group in a supergame (N=80 per treatment); income is measured as the average payoff of a subject in a supergame. Income exhibits a higher degree of inequality than in counterfactual simulations were roles alternate as in the experiment but choices are imposed. The average Gini for income is around 0.02 in the counterfactual full-defection outcome, and around 0.11 in the counterfactual full-cooperation outcome.

We find some support for this view.

Result 4. *In Roles, disadvantaged donors discriminated against advantaged recipients.*

Evidence comes from Fig. 2 and a logit panel regression similar to the one in Table 6. The dependent variable equals 1 if a donor cooperated (0 otherwise) in a round > 1 . The regression allows us to study the effect of this observable information on choices, because each donor-recipient pair is categorized based on what subjects observed in the meeting, using the *blue index*, as described before. The marginal effects are reported in Table 8.

Table 8: Past roles are visible: marginal effects

Dep. var. =1 if C chosen	Coeff.	S.E.
<i>Donor-Recipient Meeting</i>		
DD	-0.018**	(0.009)
DA	-0.064***	(0.024)
AD	0.002	(0.018)
<i>Punishment regressors</i>		
Grim trigger	-0.182***	(0.048)
Choice 1	0.072***	(0.027)
Choice 2	0.033	(0.021)
Choice 3	0.042**	(0.020)
Choice 4	0.025	(0.028)
Supergame	0.037***	(0.006)
N	2680	

Notes: Logit panel regression with random effects at the individual level and robust standard errors (S.E.) adjusted for clustering at the session level (*Roles* data only). Dependent variable = 1 if C chosen, 0 otherwise. One observation = choice of a donor in a round > 1 . Base case = donor and recipient are both advantaged (AA meeting). *Controls* include round fixed effects through a series of dummy variables (a single dummy variable for rounds 19 and above), duration of previous supergame (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), and a self-reported measure of sex and risk attitudes. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

As in *Baseline*, donors conditioned actions on their own role history. The coefficients on the DA and DD regressors are both negative and significant. Unlike *Baseline*, donors conditioned their cooperation also on the recipient's

visible history of past economic opportunities (something that could not be done in *Baseline*). This can be ascertained through a series of two-sided Wald tests. The DA coefficient is smaller than the DD coefficient (p-value=0.096), suggesting discrimination against advantaged recipients.¹⁴ Thus, H3 is rejected. Note that providing information on past roles could have led advantaged donors to act more cooperatively with recipients known to be at a disadvantage, and less with those advantaged. However, this is not so: the AA and AD coefficients are similar (p-value = 0.897), suggesting that advantaged donors did not discriminate against advantaged recipients.¹⁵

Discriminating against advantaged recipients harmed cooperation because donors switched to punishing after suffering or observing D; see the right panel in Fig. 2. The coefficient on the *Grim trigger* regressor is negative and highly significant, the sums of this coefficient and each of the *Choice n* coefficients are also negative and significant (p-values 0.025, 0.012, < 0.001, and 0.008), meaning that the cooperation decline was long-lasting, as seen in *Baseline*.

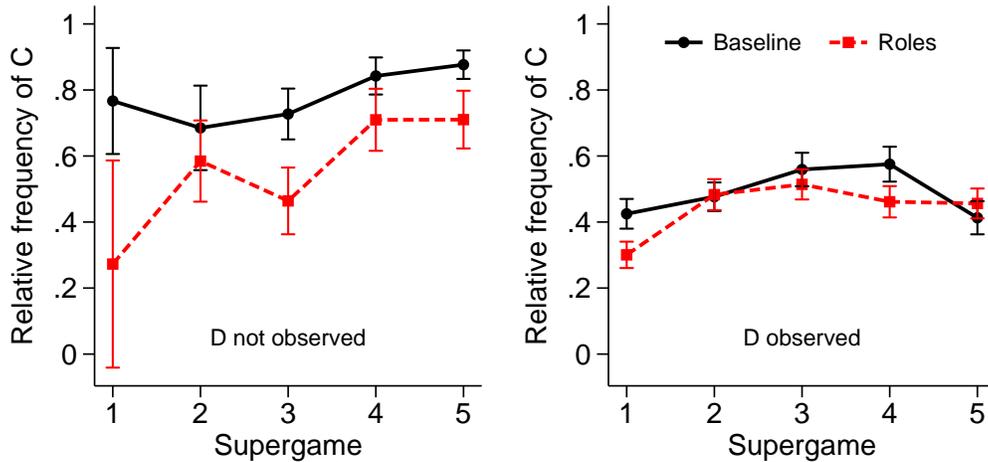
One can imagine that donors react less at the beginning of a game either because a defection has not yet been observed, or because they anticipate that an initial disadvantage can be overturned in the continuation game, or simply because “bad luck” must be persistent to become salient to a subject in the experiment. It is thus helpful to study the actions of donors in the initial rounds of the game, when defections have not yet occurred. Contrasting the *Baseline* to the *Roles* treatment provides an additional piece of evidence in support of the view that revealing inequalities in past opportunities drove

¹⁴We also see that the coefficient on the AD regressor is statistically zero, which means that donors behaved similarly in AA and AD meetings. One may think that advantaged donors would have *lowered* cooperation with disadvantaged players, as a way to punish for discriminating them. Our design hinders this kind of negative reciprocity. While donors saw the recipient’s index, the converse is not true, which is probably a reason why advantaged donors did not discriminate against the disadvantaged.

¹⁵In *Roles*, subjects also saw the past roles of the other two group members, but did not condition on this information. When we add a covariate that controls for the donor’s ranking in the distribution of roles (top, bottom or neutral), we find an insignificant impact on the cooperation probability; see col. 2 in Table B3 in Supplementary Information.

subjects away from efficient play.

Figure 2: Cooperation conditional on past actions



Notes: One obs.=one donor in a round > 1 , *Baseline* and *Roles* treatments. Each point is the average proportion of cooperative choices in the supergame, with 95% confidence bands. Left: no D experienced or seen in the group; Right: D experienced or seen in the group.

Result 5. *Donors who neither suffered nor observed a defection cooperated less in Roles than in Baseline.*

Evidence comes from Fig. 2 and Table 9, which deal with individual-level analyses. The left panel in Fig. 2 reports average cooperation rates for donors who have neither suffered nor seen a defection in the group (other than their own, possibly), in rounds after the first. The average cooperation rate is 0.78 in *Baseline* and it falls to 0.55 in *Roles*. This difference is highly significant according to a panel logit regression (p -value <0.001 , see Table 9), suggesting that in *Roles* cooperation is more fragile: there is a greater reluctance to conform to a norm of full cooperation, even if no one yet defected, as compared to *Baseline*. By contrast, we do not see a greater tolerance to defections. Defections led to a permanent decline in cooperation much it happened in *Baseline*; see right panel in Fig. 2.

Table 9: Choice when no defection observed or experienced: marginal effects

Dep. var.=1 if C chosen:	Coeff.	S.E.
Roles dummy	-0.168***	(0.044)
Supergame	0.033***	(0.005)
Controls	Yes	
N	981	

Notes: Logit panel regression with random effects at the individual level and robust standard errors (S.E.) adjusted for clustering at the session level. Dependent variable = 1 if C chosen, 0 otherwise. One obs. = choice of a donor in a round > 1, *Baseline* and *Roles* data only, in meetings where the donor has not previously suffered or seen a defection in the group (other than their own, possibly). *Controls* include round fixed effects through a sequence of dummy variables (one per each round 1-18 and one for rounds 19 and above), duration of previous supergame (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), a self-reported measure of sex and of risk attitudes. The supergame regressor is interacted with the treatment dummy. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

Summing up, we presumed that people should act with the intent to maximize payoffs and coordinate on efficient equilibria. In fact, we documented behavior that is not predicted by that theory because in our experiment individuals took actions that are inconsistent with efficient play. Donors with an unlucky streak of past economic opportunities were less cooperative than others. Second, revealing inequalities in past opportunities gave rise to discriminatory strategies, with unlucky donors withholding cooperation from more fortunate recipients. Third, revealing inequalities in past opportunities made cooperation more fragile in the initial phase of the game. It is possible that inequality or fairness concerns might have influenced choices: subjects might have sometimes acted uncooperatively to reduce their exposure to unfavorable realizations of past opportunities. Choosing D grants an immediate redistribution of earnings to a disadvantaged donor (6 points instead of 0 for the donor, 4 points instead of 25 for the recipient), but ultimately backfires because it triggered a long-lasting cooperation decline that damaged the subject’s earning prospects in the continuation game. Basing cooperation on past roles made coordination on high expected-payoff outcomes more challenging, and led to a

decline in realized efficiency.

5.2 Robustness: inequalities in earnings are observable

Disparities in past roles are a proxy for average earnings in the supergame (or, wealth), as they strongly correlate with wealth in the data; the correlation between wealth and the recipient frequency in the supergame is 0.37 in both *Baseline* and in *Roles* (one observation is one subject in a supergame, N=320 per treatment). Here, we discuss two additional treatments where we examine if the conduct observed in *Roles* is robust to making the distribution of past earnings in the group explicit.

In *Wealth*, donors observe relative past earnings of everyone in the group (earnings index). The *History* treatment combines this information with that about relative positions in past roles (blue index). Providing this information again interfered with establishing norms of long-run cooperation, although there are mixed results about the overall impact.

Result 6. *In Wealth & History coordination on the efficient outcome declined relative to Baseline, but average cooperation was similar.*

Supporting evidence is in Table 4 about group-level data, Fig. 3 about individual-level data, and Table 10, which gives an overview of all treatment effects. Table 4 reveals a decline in average cooperation and also in coordination on full cooperation, relative to *Baseline*. Regression analysis provides evidence that the decline in full cooperation is generally significant, but not the decline in average cooperation.¹⁶

¹⁶A regression based on Model 1 in Table 7, where the unit of observation is one group in a supergame (N=80 per treatment), confirms this lack of significance; see column 1 in Table B4 in the Supplementary Information. A logit regression based on Model 2 in Table 7 show a significant decline of full cooperation in the *History* treatment, but not in the *Wealth* treatment; see column 2 in Table B4 in Supplementary Information.

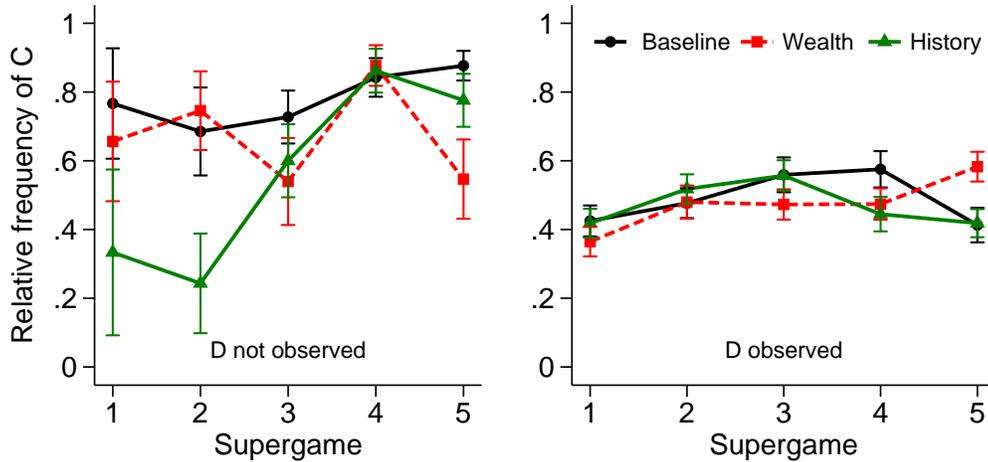
Table 10: Treatment effects: p-values

	Roles	Wealth	History
<i>Group-level regression analysis</i>			
(a) Average Cooperation	0.062	0.461	0.337
(b) Full Cooperation	0.036	0.208	0.622
<i>Individual-level regression analysis</i>			
(c) Cooperation when no D seen/experienced	< 0.001	0.079	0.043

Notes: (a),(b) Data from regressions in Tables 7 and B4 (1 obs. = one group); (c) Data from regressions in Tables 9 and B5 (1 obs. = one individual in a round).

Defections triggered a decline in cooperation, much as it happened in *Baseline*. Fig. 3 shows average cooperation for donors that did and did not suffer or observed a defection in the group.

Figure 3: Cooperation conditional on past actions in *Wealth* and *History*



Notes: One obs.=one donor in a round > 1, *Baseline* (included for comparison), *Wealth* and *History* treatments. Each point is the average proportion of cooperative choices in the supergame, with 95% confidence intervals. Left: no D experience or observed in the group; Right: D has been experienced or observed in the group.

A regression reveals that there is a statistically significant switch to a per-

manent punishment mode. The econometric model is the same as the one used earlier because we use the blue index, which is a source of exogenous wealth variation, as an instrument for the earning index (which is correlated with the donor’s past actions); see Table 11.¹⁷ The regression also reveals that donors conditioned their actions either on their own role history or that of others, albeit with some differences across treatments.

Result 7. *In Wealth and History, disadvantaged donors discriminated against advantaged recipients; donors who neither suffered nor observed a defection cooperated less than in Baseline.*

Table 11 provides evidence of differences in cooperativeness between disadvantaged and advantaged donors. In both columns, the coefficients on the DD and DA covariates are negative, and statistically different (two-sided Wald tests, p-values= 0.051 and < 0.001 for columns 1 and 2). Since $DA < DD$, this is evidence that disadvantaged donors discriminated against the advantaged.

Fig. 3 also suggests that revealing inequalities in past earnings made cooperation more fragile than in *Baseline*. Table B5 in Supplementary Information provides evidence that donors who had not seen a defection in their group are more likely to defect in *Wealth* and in *History*, as compared to *Baseline*. Overall, this is evidence that explicitly revealing the distribution of wealth in the group impacted behavior and led to discrimination against the more fortunate players, in contrast with H2 and H3.

Yet, with wealth explicitly revealed, we also see an increase in pro-social behavior as compared to *Roles*. In *History* donors do not appear to condition actions on their own past roles, if the recipient is known to be at a disadvantage; the coefficients on DD and AD are statistically similar and indistinguishable from zero (two-sided Wald test, p-value= 0.189). Second, in *Wealth* advantaged donors are more likely to cooperate with a recipient known

¹⁷In Table 11, the coefficient on the *Grim trigger* regressor is negative and highly significant, the sums of this coefficient and each of the *Choice n* coefficients are also negative and significant, meaning that the decline in cooperation was long-lasting.

to be disadvantaged (the AD coefficient is positive and significant). An interpretation is that making earnings inequalities explicit might have increased the desire to help the less fortunate, which might explain why there no longer is a significant decline in average cooperation as compared to *Baseline*.

Table 11: Past roles & cooperation: marginal effects (other treatments)

Dep. variable: =1 if C chosen	Wealth		History	
	Coeff.	S.E.	Coeff.	S.E.
<i>Donor, Recipient meeting</i>				
DD	-0.082**	(0.034)	-0.021	(0.022)
DA	-0.115***	(0.036)	-0.069***	(0.018)
AD	0.025**	(0.012)	0.024	(0.021)
<i>Punishment regressors</i>				
Grim trigger	-0.204***	(0.069)	-0.255***	(0.073)
Choice 1	0.102***	(0.038)	0.148***	(0.024)
Choice 2	0.065	(0.055)	0.126***	(0.026)
Choice 3	0.031	(0.047)	0.089***	(0.026)
Choice 4	0.030	(0.028)	0.050**	(0.025)
Supergame	0.041***	(0.010)	0.012	(0.016)
Controls	Yes		Yes	
N	2720		2856	

Notes: Logit panel regression with random effects at the individual level and robust standard errors (S.E.) adjusted for clustering at the session level Column 1: *Wealth* data only; Column 2: *History* data only. Dependent variable = 1 if donor chooses C, 0 otherwise. One observation = choice of a donor in a round > 1. Base case = donor and recipient are both advantaged (AA meeting). *Controls* include round fixed effects through a series of dummy variables (a single dummy variable for rounds 19 and above), duration of previous supergame (set to 18 rounds, in supergame 1), two measures of understanding of instructions (response time and wrong answers in the quiz), and a self-reported measure of sex and risk attitudes. The coefficient ranking is DA₁DD₁AD in column 1 and DA=DD₁DA in column 2: in both columns the DD and DA coefficients are different (Wald tests, p-values= 0.051 and < 0.001 for columns 1 and 2); the DD and AD coefficients are significantly different only in column 1 (two-sided Wald tests, p-values= 0.004 and 0.189 for columns 1 and 2). Two-sided Wald tests also allow us to establish that the partial sums of *Grim trigger* and each of the *Choice n* coefficients is negative and significant. Symbols ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

6 Discussion

Our experiment induced payoff-irrelevant inequality in groups of four individuals playing a game with an uncertain ending. Random shocks—over which players had no control—ensured equal *future* earning potential, while inducing variation in *past* earning opportunities, and thus realized earnings. Inequalities in past shocks could neither alter the power structure in the group, nor the future earning potential or the expected return from cooperation. Theoretically, income-maximizing players do not have to condition on these inequalities due to past shocks because they do not alter the payoffs that can be earned in the future (payoff-irrelevance).

Instead, the data lead to two broad observations. First, theoretically irrelevant shocks affected behavior in our laboratory supergames. This suggests that a design with exogenous variation in tasks may affect the incentives for cooperation in indefinite-horizon experiments, biasing the results away from efficient play. Second, disclosing inequalities eroded norms of mutual support, impairing coordination on efficient play. We observe this even when we added information about past earnings, which conceivably could have facilitated coordination on efficient play as high-wealth is correlated with free-riding behavior. This provides some support to the notion that fairness or inequality motivations present an obstacle to a society’s cohesion and prosperity as suggested in Putnam (2000). Field data offers ambiguous evidence on this point as many institutional and environmental factors co-vary with economic inequality. Reduction in social cohesion may stem from migration, not economic factors; inequality may stem from a mix of factors (choice, luck, power, ability) or may alter the return from cooperation. Our experiment controls for these kinds of confounding factors.

In the experiment, subjects became less cooperative as a result of repeated random assignment to the disadvantaged role; this conduct is inconsistent with maximization of ex-ante payoffs in the experimental group. Inequal-

ity disclosure gives rise to discriminatory behavior. This type of backward-looking non-strategic behavior has not been documented before in experiments on indefinite-horizon dilemmas, but it is consistent with results from finite-horizon experiments (Loch and Wu, 2008; Sonnemans et al., 1999) and yet quite puzzling because in our setup it crowded out Pareto-superior equilibria. In our setting there is no clear economic incentive to depart from efficient play by conditioning on past luck. Variation in past luck cannot alter the matrix of earnings in a round and future assignment of opportunities, hence cannot alter the equilibrium earning potential, which is always maximal under full cooperation. In fact, defecting in response to bad luck is likely to backfire, lowering the future earning potential if free-riding is sanctioned with future defections (which is what happened in the experiment). Off-equilibrium there is also no reason to condition on past opportunities because—presuming that individuals wish to coordinate on efficient play—standard theoretical arguments suggest that the threat of unconditional defection is the most effective deterrent against free-riding.

So why did subjects modulate their cooperation on payoff-irrelevant inequalities? An interpretation is that they acted with the apparent intent to counteract unfavorable past economic opportunities, attempting to equalize ex-post payoffs in the face of randomness. But this kind of play did not pay off in the experiment. Ex-post payoffs would have been much larger under full cooperation for about 90% of subjects, and only slightly smaller for the rest.¹⁸ This lack of long-run perspective is also consistent with the view that disadvantaged donors—those who had infrequent past opportunities to benefit from cooperative play—acted according to a “poverty mindset.” There is evidence

¹⁸We calculated counterfactual ex-post payoffs for each subject, under full cooperation and full defection, using the realized sequences of roles. Full cooperation would have benefited about 90% of subjects, roughly doubling their payoff (on average, 4.2 points/round gained vs. 8.6 realized in the experiment—all treatments combined). The rest would have lost 1.3 points (vs. 11.4 points realized). Full defection would have generated a slight ex-post payoff increment (0.8 points/round) for about 10% of players—those with especially long donor sequences, and created a 4.4 points average loss for everyone else.

that poverty causes stress and negative affective states, a mindset which leads to short-sighted and risk-averse decision making (Haushofer and Fehr, 2014). It is conceivable that disadvantaged donors might have internalized the feeling of being poor in their experimental group, which in turn re-aligned their incentives away from the long-term but risky gains made possible by cooperation, toward the short-term and sure gains granted by defection.

Another interpretation is that fairness concerns might have dominated efficiency concerns—as documented in Güth et al. (2003) for the case of one-shot dictator games. Subjects might have been unwilling to follow a norm of unwavering mutual support when the associated benefit did not reflect *their* contribution to the prosperity of others, thus acting out of a desire to balance their contribution to others’ prosperity and their own benefit. This behavior is consistent with a “luck egalitarian” norm, according to which subjects act to reduce a disadvantage due to uncontrollable factors (the random role assignment in our case), as observed in non-strategic settings (Konow, 2000; Mollerstrom et al., 2015). Finally, it is also possible that providing information about past roles or earnings might have itself affected cooperation levels either by causing an informational overload, or by affecting inference judgment.

In conclusion, we know from one-shot experiments that individuals may be driven by a mix of motives, including fairness or aversion to inequality in outcomes. Our experimental results suggest that these considerations apply to indefinitely repeated laboratory supergames as well. It is conceivable that equilibria exist in which defections are tolerated under certain circumstances (e.g., being disadvantaged) but trigger punishment in others. If a norm of this kind is indeed adopted, then this might explain the decline in coordination of efficient play when past roles and earnings were disclosed. We see this as an important direction for future work.

Provided that our findings scale up to larger economies, we offer two additional insights. The first one concerns economic policy at the national and

international level. A nation's fiscal policy is inherently redistributive. Our experiment suggests such policies should be evaluated not only in terms of their effectiveness at raising revenue while mitigating market distortions, but also their potential impact on society's cohesion. This consideration can be extended to the transnational domain of economic unions, where transfer policies are often justified by a mix of solidarity and efficiency-enhancing motives (e.g., the "cohesion policy" of the European Union). Our study provides an additional reason why, if countries have heterogeneous pre-existing economic conditions, redistributing resources can benefit the entire economic area. A second consideration concerns the social implications of inequality. Much of the economics literature has focused on inequality's potential to lower society's welfare by misshaping the structure of economic incentives, and hence performance. Our study suggests that inequality alters economic behavior even when it leaves the economic incentives unaltered (payoff-irrelevance), as it inhibits pro-sociality and fosters discrimination. The study did not uncover the ultimate causes of this influence on behavior (e.g., did envy play a role?) and suggests that a more systematic approach is needed to make progress.

References

- Abreu, D., D. Pearce and E. Stacchetti. 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58, 1041-1063.
- Aghion, P. and J. Williamson. 1998. *Growth, Inequality, and Globalization*. New York: Cambridge University Press.
- Anderson, L. R., J. M. Mellor, and J. Milyo. 2006. Induced heterogeneity in trust experiments. *Experimental Economics* 9, 223-235
- Andreoni, J. and H. Varian. 1999. Preplay contracting in the Prisoners' Dilemma. *Proceedings of the National Academy of Sciences of the United States of America* 96(19) 10933-10938.
- M. Bigoni, G. Camera, and M. Casari, 2018. Partners or Strangers? Coop-

- eration, monetary trade, and the choice of scale of interaction. *American Economic Journal: Microeconomics* (forthcoming).
- Blanco, M., D. Engelmann and H. T. Normann. 2011. A within-subject analysis of other-regarding preferences. *Games and Economic Behavior* 72(2), 321-338
- Bolton, G.E, Ockenfels, A., 2000. ERC: a theory of equity, reciprocity and competition. *American Economic Review* 90, 166-193.
- Camera, G. and M. Casari. 2009. Cooperation among strangers under the shadow of the future. *The American Economic Review*, 99(3), 979-1005.
- Camera, G. and M. Casari. 2014. The Coordination Value of Monetary Exchange: Experimental Evidence. *American Economic Journal: Microeconomics*, 6(1), 290-314.
- Camera, G., Casari, M., and Bigoni, M. 2013. Money and trust among strangers. *Proceedings of the National Academy of Sciences of the United States of America*, 110(37), 14889-14893.
- 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.
- Deck, C. 2001. A Test of Game Theoretic and Behavioral Models of Play in Exchange and Insurance Environments. *American Economic Review* 91(5), 1546-1555.
- Ellison, G. 1994. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61, 567-88
- Fehr, E. and K. M. Schmidt. 1999. A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* 114(3), 817-868
- Fischbacher, U. 2007. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2), 171-178.

- Gangadharan, L., N. Nikiforakis and M. C. Villeval. 2015. Equality Concerns and the Limits of Self-Governance in Heterogeneous Populations. Manuscript.
- Goeree, J. K., and C. A. Holt. 2000. Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *European Economic Review* 44(4), 1079-1089.
- Greiner, B., A. Ockenfels, and P. Werner. 2012. The Dynamic Interplay of Inequality and Trust—An Experimental Study. *Journal of Economic Behavior and Organization* 81, 355-365.
- Güth, W., H. Kliemt, and A. Ockenfels. 2003. Fairness versus efficiency: An experimental study of (mutual) gift giving *Journal of Economic Behavior & Organization* 50(4), 465-475
- Harrison, G. W., M. I. Laub, and E. E. Rutström. 2009. Risk attitudes, randomization to treatment, and self-selection into experiments. *Journal of Economic Behavior & Organization* 70(3), 498-507
- Harsanyi John C., Selten Reinhard (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.
- Haushofer, Johannes and Ernst Fehr, 2014. On the psychology of poverty. *Science* 344(6186), 862-867.
- Holt, C. A, and S. K. Laury. 2002. Risk Aversion and Incentive Effects. *American Economic Review* 92(5), 1644-1655
- Kagel, J. H., and K. Willey-Wolfe. 2001. Tests of fairness models based on equity considerations in a three-person ultimatum game. *Experimental Economics* 4(3), 203-219.
- Kandori, M. 1992. Social norms and community enforcement. *Review of Economic Studies*, 59, 63-80.
- Konow, J. 2000. Fair shares: accountability and cognitive dissonance in allocation decisions. *American Economic Review* 90, 1072-1091.
- Loch, C. H. and Y. Wu. 2008. Social Preferences and Supply Chain Performance: An Experimental Study. *Management Science* 54(11), 1835-1849,

- Mollerstrom J., B. Reme, and Erik Sørensen. 2015. Luck, choice and responsibility—An experimental study of fairness views. *Journal of Public Economics* 131, 33-40
- Nishi, A. H. Shirado, D. G. Rand, and N. A. Christakis. 2015. Inequality and visibility of wealth in experimental social networks. *Nature*.
- O'Reilly, Charles A. 1980. Individuals and Information Overload in Organizations: Is More Necessarily Better? *The Academy of Management Journal* 23(4), 684-696
- Palfrey, T.R. and Rosenthal, H., 1994. Repeated Play, Cooperation and Coordination: An Experimental Study. *The Review of Economic Studies*, 61(3), 545-565.
- Piketty, T. 2014. *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.
- Putnam, R. 2000. *Bowling Alone: The Collapse and Revival of American Community*.
- Rabin, M. 1993. Incorporating Fairness into Game Theory and Economics. *American Economic Review* 83, 1281-1302.
- Roth, A. E., and K. Murnighan. 1978. Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17, 189-98
- Sadrieh, A., and H.A. Verbon. 2006. Inequality, cooperation, and growth: an experimental study. *European Economic Review* 50, 1197-1222.
- Sherstyuk, K., N. Tarui, and T. Saijo. 2013. Payment schemes in infinite-horizon experimental games. *Experimental Economics* 16 (1), 125-153.
- Smith, V. 1994. Economics in the Laboratory. *Journal of Economic Perspectives* 8(1), 113-131.
- Sonnemans, J., A. Schram, T. Offerman. 1999. Strategic behavior in public good games: when partners drift apart. *Economics Letters* 62, 35-41
- Stiglitz, J. E. 2012. *The Price of Inequality: How Today's Divided Society Endangers Our Future*. W. W. Norton & Company, New York.

Tavoni, A., A. Dannenberg, G. Kallis, and A. Löschel. 2011. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proceedings of the National Academy of Sciences of the United States of America*, 108 (29), 11825-11829.

Troutman, C. M., and Shanteau, J., 1977. Inferences based on nondiagnostic information. *Organizational Behavior and Human Performance* 19, 43-55.

Appendix A

Procedural Details

We recruited 256 subjects through announcements to the standing subject pool for the Behavioral Business Research Laboratory at the University of Arkansas. No subject had previous experience with this type of game in the lab.¹⁹ After giving informed consent, subjects were seated at private terminals. Neither communication nor eye contact was possible among subjects at any time during the session. The experimenter publicly read the paper instructions at the start of the experiment, a copy of which were then left on the subjects' desks. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). On average, a session lasted 94 rounds for a running time of about 120 minutes including instructions, a paid post-instruction comprehension quiz, and post-experiment payment. Average earnings were \$27.00 per subject (min = \$6.50, max = \$55.50) excluding a \$5 fixed participation payment and an average of \$2.10 (min = \$.75, max = \$2.50) from providing correct answers to the comprehension quiz (\$0.25 for each correct answer out of 10 questions). Only one randomly selected supergame from the session was paid and subjects knew this fact in advance.

Proof of Proposition 1

This analysis is based on the existence of equilibrium proof in Camera et al. (2013). In each round $t = 0, 1, 2 \dots$ individuals in the group are matched in pairs, with uniform probability of selection. In each pair, a computer randomly determines who is the donor and who is the recipient (with equal probability). If cooperation (=Help) is the outcome, then g is the payoff to the recipient and for generality let a denote the payoff to the donor. If defection (=Do nothing) is the outcome, then d is the payoff to the donor and $d - l$ to the recipient. Round payoffs are geometrically discounted at rate $\beta \in (0, 1)$ starting from round $n \geq 0$.

The equilibrium payoff (=expected lifetime utility) at $t = 0$ is

$$v(n) := (n + 1) \times \frac{g + a}{2} + \sum_{j=1}^{\infty} \beta^j \times \frac{g + a}{2} = \frac{g + a}{2} \times \left(n + \frac{1}{1 - \beta} \right).$$

¹⁹About 55% of subjects were males, and the rest female. The subject pool is composed of about 90% undergraduate students with the remainder being primarily graduate students although some faculty, staff, and non-university associated people are in the pool. We include sex differences as controls in the econometric analysis.

A player is a donor or a recipient with equal probability in each round, hence expects to earn $\frac{g+a}{2}$ in each round. The payoff $v(n)$ is increasing in n because payoffs are discounted by β in rounds $t \geq n$.

The equilibrium payoff in the continuation game starting on any date $t \geq 0$, before any uncertainty is resolved, corresponds to

$$V_t = \begin{cases} v(n-t) & \text{if } t < n \\ v^* := \frac{g+a}{2(1-\beta)} & \text{if } t \geq n. \end{cases}$$

The equilibrium payoff of a donor at the start of any date t is

$$V_{dt} = \begin{cases} a + v(n-t-1) & \text{if } t < n \\ a + \beta v^* & \text{if } t \geq n. \end{cases}$$

We must check that in equilibrium donors have no incentive to defect; out of equilibrium, donors have no incentive to cooperate.

Defection is the dominant action off-equilibrium; i.e., it is always individually optimal to punish after a defection from equilibrium play is made public. To see this suppose a donor deviates by cooperating off equilibrium. She would earn a instead of d but her continuation payoff would not improve since everyone else keeps defecting—as prescribed by the rule of punishment. Since $d > a$, it is optimal to punish off equilibrium.

In equilibrium, cooperation is a best response in every round $t = 0, 1, \dots$, if $V_{dt} \geq \hat{V}_{dt}$. The left-hand-side denotes the payoff to a donor who cooperates; the right-hand-side denotes the donor's payoff when she moves off equilibrium under a one-time deviation. Such deviation is publicly observed, hence—when everyone follows the cooperative strategy—every donor will always defect in the future. The payoff to the deviator is thus

$$\hat{V}_{dt} = \begin{cases} \hat{v}(n-t) := d + (n-t)\frac{2d-l}{2} + \beta\frac{2d-l}{2(1-\beta)} & \text{if } 1 \leq t < n \\ \hat{v}^* := d + \beta\frac{2d-l}{2(1-\beta)} & \text{if } t \geq n \end{cases}$$

Now define

$$\Delta_t = V_{dt} - \hat{V}_{dt} = a - d + \frac{g+a-2d+l}{2} \times \begin{cases} n-t + \frac{\beta}{1-\beta} & \text{if } t < n \\ \frac{\beta}{1-\beta} & \text{if } t \geq n \end{cases}$$

The minimum value of Δ_t is achieved for $t \geq n$. The implication is that cooperation is individually optimal in all rounds t whenever

$$\beta \geq \beta^* := \frac{2(d-a)}{g+l-a}.$$

The grim strategy is risk dominant

Our game is sequential, does not have a two-by-two structure, and admits more than two equilibria. Therefore, we consider a notion of risk dominance that departs from the standard formalization in Harsanyi and Selten (1988), but conforms to the heuristic argument that motivates it. A risk-dominant equilibrium maximizes the expected payoff when players have uniformly distributed second order beliefs on the best and worst equilibrium (full cooperation and full defection). We also add structure to beliefs about strategies play, by exploiting the dynamic structure of the game. In doing so, we employ the technique in Bigoni et al. (forth.), to which we refer the reader.

Given random role alternation, the ex-ante expected payoff under full defection equilibrium is denoted

$$\hat{v} := \frac{2d-l}{2(1-\beta)}.$$

The ex-ante expected payoff under full cooperation is

$$v := \frac{g}{2(1-\beta)}.$$

Consider uncertainty over two competing strategies: “grim” (G) and “always defect” (AD). Initial donors select a strategy in round 1 and maintain it for the rest of the supergame. Initial recipients take no action in round 1, so we set them free to select G or AD in round 2. Given public monitoring, conjecture that those who choose strategy in round 2 coordinate their choices by best responding to the strategy play observed in round 1. If so, then all uncertainty about future play is resolved at the end of round 1 as if initial donors selected the equilibrium. G dominates AD after any history of play, for those who choose their strategy in round 2 (weakly dominates AD , if someone defected in round 1). Hence, if no-one (someone) defected in round 1, then every donor will cooperate (defect) in every future meeting.

Consider round 1. There is strategic uncertainty because an initial pro-

ducer is not sure what strategy the other initial producers will select. Suppose that every initial producer believes that in round 1 there is probability p that C is the outcome in any given pair; D is the outcome with the complementary probability. A special case is $p = 1/2$, which may be motivated by the “principle of insufficient reason” for a player who is unsure about what the others will do. That is, the individual believes that the other initial producer plays G with probability p , and AD otherwise. At the end of round 1 either C will be the outcome in every future meeting, or D will be the outcome in every future meeting.

Fix an initial donor. Let denote V_G and V_{AD} the expected utilities from choosing strategy G and AD where

$$\begin{aligned} V_G &= 0 + p\beta v + (1 - p)\beta\hat{v}, \\ V_{AD} &= d + \beta\hat{v}. \end{aligned}$$

Consider V_{AD} : the initial donor defects so all future donors will defect whether or not they chose G or AD . Consider V_G : with probability p the other initial donor is also a grim player. In that case both cooperate and there is full cooperation forever so the continuation payoff is βv . With probability $1 - p$ the other initial donor plays AD in which case the continuation payoff is $\beta\hat{v}$ because full defection ensues. We say that G is risk dominant if $V_G \geq V_{AD}$. If $p = 0.5$, which is a standard consideration, this implies

$$\beta \geq \frac{4d}{g + 2d + l} \quad \Rightarrow \quad \beta \geq 0.62$$

Therefore grim trigger is a risk-dominant strategy in our experiment since the discount factor is 0.75.