

Competition between organizational groups: Its impact on altruistic and anti-social motivations*

Lorenz Goette[†] David Huffman[‡] Stephan Meier[§] Matthias Sutter[¶]

July 15, 2011

Abstract

Firms are often organized into groups. Group membership has been shown empirically to have positive effects, in the form of increased pro-social behavior towards in-group members. This includes an enhanced willingness to engage in altruistic punishment of inefficient defection. Our paper provides evidence of a dark side of group membership. In the presence of cues of competition between groups, a taste for harming the out-group emerges: punishment ceases to serve a norm enforcement function, and instead, out-group members are punished harder and regardless of whether they cooperate or defect. Our results point to a mechanism that might help explain previous mixed results on the social value of punishment, and also contribute to understanding the sources of conflict between groups. They also point to an important tradeoff for firms: introducing competition enhances within-group efficiency, but also generates costly between-group conflict.

*We thank seminar and conference participants at Harvard University, George Mason University, UC Berkeley, UC Santa Barbara, University of Pittsburgh, Temple University, Tel Aviv University, Wharton School, Tilburg University, ESA Tucson for very helpful comments.

[†]Department of Economics, University of Lausanne, Internef, 1015 Lausanne-Dorigny, Switzerland

[‡]Department of Economics, Swarthmore College, 500 College Avenue, Swarthmore, PA 19081, USA

[§]Graduate School of Business, Columbia University, 3022 Broadway, New York, NY 10027, USA; sm3087@columbia.edu

[¶]Department of Public Finance, University of Innsbruck, Universitaetsstrasse 15, 6020 Innsbruck, Austria; Department of Economics, University of Gothenburg, Vasagatan 1, 40530 Gothenburg, Sweden

1 Introduction

Many beneficial exchanges within societies and organizations require individuals to cooperate and trust each other, even though narrow self-interest may tempt them to act selfishly. The pervasive tendency for formal contracts to be incomplete enhances the importance of such pro-social behaviors for efficiency of organizations and firms (Williamson *et al.*, 1975; Knez and Simester, 2001; MacLeod, 2007). Choice experiments designed to isolate such altruistic motives have found evidence for altruistic cooperation, but show that on its own altruistic cooperation typically breaks down (e.g., Andreoni, 1995; Fischbacher and Gaechter, 2010). Rather, high levels of cooperation are often sustained by a willingness of subjects to expend resources to impose harm on others who act selfishly. This altruistic, and costly, informal sanctioning has been shown to deter defection (Ostrom, 1990; Güth *et al.*, 1982; Fehr and Gaechter, 2000; Fuster and Meier, 2010). Thus, mutual monitoring and sanctioning in terms of social pressure or other means, could be particularly important in contributing to an organization’s efficiency (Kandel and Lazear, 1992; Mas and Moretti, 2009). A key question, then, is what factors can enhance (or undermine) pro-social norm enforcement within organizations?

This paper focuses on one particular factor that has received increasing attention in the organizational literature: membership in a group (see, e.g., Akerlof and Kranton, 2000, 2005). Organizations and firms are typically organized into groups, and it has been shown using choice experiments that membership in a group can increase pro-social behaviors towards in-group members (e.g. Charness *et al.*, 2007; Chen and Li, 2009; Chen and Chen, 2011), including greater willingness to enforce norms within the group (Goette *et al.*, 2006). This enhanced norm enforcement and cooperation within groups then has positive consequences for overall efficiency of the organization. However, there is also a long-standing hypothesis that group membership might have a dark side, generating hostility, and anti-social actions, towards outsiders (e.g., Durlauf, 1999; Hewstone *et al.*, 2002). If there is costly conflict between groups within an organization, this would be harmful for organizational efficiency. So far, however, this type of negative effect of group membership has not been observed in experiments studying group effects, and the mechanisms that could cause it to arise are not fully understood.¹ The contribution of this paper is to show how a negative impact of group membership can arise.

We identify a mechanism – a cue of competition between groups – that can generate between-group conflict through a particular channel, i.e., punishment behavior. In a baseline condition, without competition (Goette *et al.*, 2006), we found that punishment is pro-social, in the sense that it is selectively used against defectors. Adding competition, punishment behavior is transformed from beneficial norm enforcement, to “anti-social punishment” of outsiders, which means

¹Note that, conceptually, in-group favoritism need not be the flip-side of out-group hostility. Individuals might favor the in-group, without going out of their way to harm the out-group (see, e.g., Bahry *et al.* (2005) for evidence of inter-ethnic trust in Russia).

outsiders are punished harder, and regardless of whether they cooperate or defect. We find little evidence that people cooperate less with outsiders as the result of competition, so conflict appears to work mainly through punishment behavior. There are, however, especially high levels of cooperation with in-group members when competition is present. We argue that previous studies have not found such effects because it is necessary to have both competition between groups, and punishment options. We discuss how a shift from pro-social to anti-social punishment, depending on the presence of competition, is consistent with the predictions of recent theories on the coevolution of human altruism and aggression (Choi and Bowles, 2007; Bowles, 2009). Our results also indicate that different preferences may be dormant in one individual, and that these preferences may be triggered by cues about the nature of the surrounding economic environment (Bowles, 1998).

Our experiments were conducted using officers in the Swiss Army, and involved four treatments that varied the group membership of players in the choice experiments. The group manipulation was generated by the Army: for the duration of its officer training program, the Army randomly forms platoons of the officer candidates, thus providing a strong, yet exogenous manipulation of group membership. The choice experiments were one-shot interactions, and were designed to test for willingness to cooperate altruistically, or to use a costly punishment option for either pro-social or anti-social purposes. As we explain in section 2, competition between groups was implemented such that it left the monetary payoffs from punishment unchanged. Thus, the appearance of anti-social punishment directed towards the out-group must be due solely to the creation of a more competitive atmosphere, not because of a change in which punishment can affect outcomes. The “army as laboratory” is an interesting setting, as all officers in the Swiss army serve as reservists and in their civilian lives are active in business life, typically in leadership positions.² Furthermore, the army setting provides a largely unique opportunity to observe random assignment to real social groups, in combination with choice experiments and exogenously implemented competition.

Of course, it has already been shown experimentally that competition between groups has behavioral implications for in-group and between-group interactions, specifically in terms of generating more cooperation within groups (Bornstein and Ben-Yossef, 1994; Bornstein *et al.*, 2002; Augenblick and Cunha, 2009). Similar to these earlier studies, we also find that under competition there is a particularly strong tendency to cooperate with the in-group (Halevy *et al.*, 2008). Our results are also similar, in that we do not observe hostility being expressed through the decision to cooperate or defect.

²One might think that this type of individual would be especially sensitive to cues of competition. On the other hand, it is remarkable that introducing a cue of competition between groups can generate conflict, within an organization that relies so crucially on cooperation between groups. Furthermore, the same individuals also have leadership positions in their civilian lives, thus underscoring the relevance to examine whether mere cues of competition can affect an individual’s preferences towards other groups.

One way that we differ from previous studies on competition between groups is that we also look at punishment behavior. This is important, because conflict appears to be expressed through this channel, and because we show how this can completely subvert what has been argued to be a crucial mechanism for sustaining cooperation. Another difference is our approach of using army officers. This keeps many aspects from the field while still leveraging the advantages of experimental methods. We build, in general, on Goette *et al.* (2006), where we investigated in-group and out-group effects, but did not examine how cooperation and punishment behavior interact with absence or presence of competition. The use of randomly assigned real groups, i.e., platoons, has significant advantages over most previous approaches. Studies based on real groups or existing friends (e.g., Fershtman and Gneezy, 2001; Bahry *et al.*, 2005; Bernhard *et al.*, 2006; Leider *et al.*, 2009) analyze groups with social content and social ties which is an important aspect of real groups. However, these groups are endogenously formed or differ in other dimensions than just their group membership (i.e. ethnicity), making inferences about the effects of groups *per se* difficult. A solution to these confounds is to randomly assign individuals to so-called “minimal” groups which are nothing more than a label (e.g., Tajfel *et al.*, 1971; Charness *et al.*, 2007; Chen and Li, 2009; Sutter, 2009). By design, however, these groups lack any social content and behavior might be very different from real groups, as shown in Goette *et al.* (2010). Combining the advantages of both group manipulation methods, our groups do have social content as groups in the former method, but are at the same time randomly assigned as groups in the minimal group-paradigm.³

Our results are relevant for an ongoing debate about the social value of punishment. While many studies have found that punishment serves a pro-social norm enforcement function (Ostrom, 1990; Güth *et al.*, 1982; Fehr and Gaechter, 2000; Fuster and Meier, 2010), other studies question whether punishment has such beneficial effects (Dreber *et al.*, 2008; Egas and Riedl, 2008; Herrmann *et al.*, 2008; Houser *et al.*, 2008; Abbink *et al.*, 2010). For example, Herrmann *et al.* (2008) find anti-social punishment directed at people who are relatively cooperative, as well as defectors, in certain societies. Our experiments provide some of the first evidence on a mechanism that determines whether the pro-social or anti-social form of punishment behavior will emerge.

The findings of our experiment are also relevant for understanding why social groups are sometimes in conflict, and other times not. Research studies using field data have often conjectured that competition generates conflict. For example, the extent of competition for political power has been argued to explain why the Chewas and the Tubukas are enemies in Malawi, but are

³The advantage of the “minimal group” paradigm is that it allows investigating the pure “labeling effect” of groups versus the effect of actually identifying with the group. Our methodology does not differentiate between these two steps of group identity formation, categorization and identification, as they were completed before the experiment was conducted. Our results are thus not directly comparable to those of minimal group studies; they are, however, applicable to “real groups”, where both categorization and identification have occurred.

friends right across the border in Zambia (Posner, 2004), and conflict between natives and immigrants has been linked to the extent of competition in the job market (Esses *et al.*, 1998). Within organizations, it can happen that different groups/departments work well together or they fight each other leading to huge burdens for the organization. For example, Auletta (1985) describes how the fight between the “bankers” and the “traders” almost lead to the end of Lehman Brothers. The main challenge for this type of study, however, is the presence of many factors that confound a clean identification of group effects *per se*, and of the effect of the economic environment. For example, groups typically differ according to many characteristics, and these differences could drive behavior rather than group membership. The extent of competition is also often not randomly assigned, raising the concern that more hostile types self-select into competitive situations. Also, behavioral measures have typically not allowed disentangling strategic motives from non-strategic motives that lead to conflict between groups.⁴

Finally, our findings indicate that organizations and firms may face a tradeoff when it comes to trying to influence whether the “corporate culture” encourages internal competition between groups. Our results suggest that, while a sense of competition may enhance efficiency within groups, it may also lead to costly conflict between groups. If between-group conflict is sufficiently costly, the net effect of competition on efficiency may even be negative. The trade-off is also relevant for the choice of incentive schemes within organizations. Explicit incentives to compete in the form of, e.g., team tournaments, would presumably generate even more salient competition, and even stronger hostility effects.

The outline of the paper is as follows. In section 2 we introduce the experimental design. Section 3 presents behavioral hypotheses, and section 4 reports the experimental results. Section 5 concludes the paper.

2 Experimental Design

We conducted experiments in the Swiss Army, which allow us to exploit the random assignment of individuals into platoons as our group manipulation. Platoon assignment generated different treatments in terms of four different combinations of group membership for players in our experiments. We also used two types of experimental games, and two conditions in each game, to investigate the effect of competition between groups on cooperation and particularly on norm enforcement.

⁴The seminal study in psychology involved young boys being randomly assigned to different groups at a summer camp, and being observed as they first played competitive games, and subsequently engaged in cooperative activities (Sherif *et al.*, 1961). Our study is different because we use controlled choice experiments where anonymity allows disentangling strategic and non-strategic motives. Other differences include randomization of competition, rather than a fixed order, and using adults ($N = 525$) rather than young boys ($N = 24$).

2.1 Subject Pool and Random Group Assignment

All Swiss males are required to perform at least 300 days of military service, beginning with twenty-one weeks of basic training. In week seven, about one fourth are selected to go through ten weeks of officer-candidate training. Of these, one fourth are promoted to officers. Since the Swiss Army is a reserve army up to the highest officer ranks, all officers have a civilian life and typically hold leadership positions in organizations. In part, the Army selects officers based on their education. Among our officer candidates, 41 percent have completed an apprenticeship, while the remaining 59 percent are on a university track. The comparable age group of 19 to 26 year-olds in Switzerland have significantly lower educational attainments. According to the Swiss Labor Force Survey 2008, 23 percent only completed compulsory schooling, 48 percent have completed an apprenticeship and only 30 are on the university track. Jann (2003) documents that officers in the Swiss Army earn a wage premium of 25 percent compared to other Swiss males.⁵ Thus, officers in the Swiss Army are a select group that is particularly likely to be in management positions in civilian life. From this viewpoint, they also an interesting set of subjects.

As part of the training program, all officer candidate participate in the Joint Officer Training Program (JOTP). Whereas officer-candidate training is specific for each branch of service, and occurs in separate locations, JOTP brings new officers from all branches of service together, to the same location, for four weeks. Officers are randomly assigned to a platoon at the beginning of JOTP, and spend virtually all time during the day with their platoon. Training involves mainly coursework on principles of security, combat in large military units, logistics, and leadership. At the end of JOTP, the platoons are dissolved and officers are once again sent to separate locations, for further, advanced training specific to each branch of service.

We use the random assignment of candidates to platoons in JOTP as our manipulation of social groups. Each platoon is identified by a different number. Assignment to platoons is random, and stratified according to the different branches of service. The Army assigns platoons orthogonally to any previous social ties among officers with the aim of promoting exchanges of perspectives among different individuals and branches of service.

The assignment mechanism is ideal, in several ways, for investigating the impact of group membership on behavior. First, trainees know that platoon composition is designed to be identical and that nobody could choose which platoon to join. Indeed, statistical tests reveal no significant differences in platoon composition, by branch of service, education, or age. Second, there is no competition between the groups (or trainees) for evaluations or other resources. Relative performance evaluations were completed previously, in candidate training. Thus, there is no function of the group assignment, other than to affect the circle of individuals with whom an

⁵This is the unconditional wage premium, which is the relevant one for our comparisons.

officer interacts most frequently. Third, social interactions within a platoon are intense. Platoon members spend the whole workday with their group, for the three weeks leading up to our experiments. Despite the fact that platoons are assigned orthogonally to previous social ties, social interactions and ties also arise endogenously within platoons in after-work time: In a questionnaire, officers in our study report to a question on “*How often do you spend off-duty time with members of a) your own platoon or b) the other platoons?*” that they spend significantly more time off-duty with members of their own platoon. This is notable given that 79.8 percent of the trainees know people in other platoons, mostly from earlier stages of their training. Yet, as illustrated in Table 1, they choose to spend most of their off-duty time with members of their newly assigned group.

[Table 1 about here.]

By using randomly assigned real groups we incorporate the social ties aspect of real groups that is abstracted from in arbitrary “minimal groups”. The groups are also attractive because they are not endogenously formed (as in, e.g., Leider *et al.*, 2009) nor do they differ in other dimensions than just membership to different groups, e.g. nationality or ethnicity (as in, e.g., Bernhard *et al.*, 2006; Habyarimana *et al.*, 2007; Bahry *et al.*, 2005). This allows us to make inference about the causal effect of real groups on behavior.

2.2 Experiments and Group Conditions

In this subsection we present the two types of choice experiments used in our study, and in the next subsection we introduce the two conditions with which we varied the economic environment as being competitive or non-competitive. Experiments were one-shot and were always conducted in the third week of the four-week training period.

Experiment 1: Cooperation. The game was a simultaneous prisoners’ dilemma (PD). The players, labeled A1 and A2, were each endowed with 20 points worth real money (4 points = 1 CHF). They simultaneously decided whether to keep the points or pass all of them to the other player. Passed points were doubled. Thus, if both players passed their points (cooperation), they each got 40 points. However, a selfish player could always do better by keeping the points (defecting), regardless of the other player’s decision: Defecting when the other defected would yield 20, whereas cooperating would sacrifice the endowment and yield nothing in return; defecting when the other cooperated would yield 60, the maximum possible individual payoff in the game (while leaving the cooperator with 0). Cooperation thus entails lowering one’s own payoff, and improving the payoff of the other player, and is an indicator of non-selfish motives. We use the game as our workhorse for studying how group boundaries, and economic environment, affect non-selfish motives for cooperation.

Experiment 1 involved two conditions in a between-subject design; individuals did not know about treatments besides the one in which they participated. In all conditions, a subject never learned the individual identity of their partner. In the *in-group* condition, subjects interacted anonymously, except for being informed that the other player was a member of their own platoon. The *out-group* condition was the same, except subjects were informed that the other player was a member of another platoon. Group affiliation was clearly marked on the decision sheets. These conditions allow us to examine how group assignment affects cooperation. For a selfish individual, the group affiliation will not change the prediction that he will always defect.

We also elicited individual's beliefs about in-group and out-group cooperation. Independent of the condition they were in, we asked participants to state both their belief for in- and out-group cooperation. We asked them to predict the percentage of the in- and out-group that would send all of the points (cooperate). They were given an incentive to make their best guess: they knew that their prediction would be compared to the percentage actually observed. If the deviation was less than 10 percentage points, then they would get one extra point.

At the very end of the experimental sessions, we conducted a short survey in which we asked participants whether they agreed or disagreed with three statements about trust: 1) "*In general, people can be trusted.*", 2) "*Nowadays, you can't rely on anybody.*", and 3) "*Dealing with strangers, it is better to be cautious before trusting them.*". Participants answered on a 4-point scale (1 "Agree Strongly", 2 "Agree Slightly", 3 "Disagree Slightly", and 4 "Disagree Strongly"). We created an individual variable, *Trust*, by adding the answers to the three questions and assigning a 1 for the least amount of trust and 4 the highest amount of trust per question (answers to question 1 are reversed coded). This is used to help capture individual differences in beliefs about trustworthiness in our statistical analysis.

Experiment 2: Punishment. In Experiment 2, two players A1 and A2 played a PD as in Experiment 1, but we added two additional players, B1 and B2. Each B-player was endowed with 70 points. B1 could assign up to 10 deduction points to A1, and B2 could do the same to A2. Each deduction point subtracted three points from the A-player, and cost the B-player one point of his endowment. The B-players could condition their choices on the actions of A1 and A2. Thus Experiment 2 incorporated the possibility of third-party punishment (Fehr and Fischbacher, 2004), and is suited for examining determinants of whether punishment takes the form of norm enforcement (selectively punishing defection) or anti-social punishment (punishing both cooperation and defection). As punishment is costly, a selfish B-player would never punish.

To examine the impact of group membership on norm enforcement, we varied the composition of players in each game in a between-subject design. For the remainder of the paper, we refer to the group composition in Experiment 2 from B1's perspective. Thus, A1 always refers to the player that the B-player can punish, while we refer to the other A-player as A2. The four

different group compositions we implemented are shown in Figure 1.

[Figure 1 about here.]

Varying the group membership of A1 (while keeping constant the group membership of A2) allows us to investigate how the group identity of the person being punished (A1) matters. We also study how punishment varies with the group affiliation of A2, the person affected by A1's actions. Appendix C provides a translation of the instructions for one group composition in the *Neutral Group Environment* treatment.

2.3 Economic Environment Treatments

We compare two conditions, which differ in terms of the absence or presence of cues of competition between groups.

Neutral Group Environment (NG): In this baseline condition, we used the randomly assigned groups as our group manipulation and varied the group composition as described above. There was no economic competition between the platoons. Preliminary results from baseline were previously presented in Goette *et al.* (2006).

Competitive Group Environment (CG): We added competition to the 'Neutral Group Environment' treatment by offering a bonus to the platoon that got the highest payoff in the PD stage. The bonus was 20 points for each member if the platoon got the highest average payoff in the PD. In case of a tie between two platoons, the winner was randomly determined. Because the bonus was based on average payoffs for pairs playing the PD, and cooperation maximized payoffs for the pair, cooperation facilitated winning the bonus for the platoon. Importantly, however, the bonus did not change the incentives for a selfish individual: the best strategy for a selfish A-player was still to defect (for the intuition and a formal test, see section 3 on the behavioral hypothesis). Furthermore, in Experiment 2 the bonus was calculated based on the A-player average payoffs *before* deducting any punishment points imposed by the B-players. B-players (and A-players) knew this. This design feature makes punishment different from sabotage, in which 'hurting' the other group can affect the probability of winning the prize in a tournament (see Harbring and Irlenbusch, 2011, for a paper on sabotage). Thus, the bonus was irrelevant for the choices of the B-players, regardless of whether they were selfish or altruistic. The rules of the game, in terms of the size of the bonus, and the irrelevance of punishment points for influencing the competition outcome, were made clear in the instructions. Furthermore, we only began the experiment after control questions verified that all participants understood these specific features.

2.4 Experimental Procedures

The experiment was conducted with paper-and-pencil in a large auditorium and lasted 45 minutes. Subjects did not know of the experiment in advance. The different treatments were conducted in different sessions.

Special care was taken to ensure particularly strong anonymity conditions. First, subjects were never told the identity of their partner(s). Second, they knew that payoffs would be mailed to home addresses ten days after the experiment, so that all participants would only learn the outcome of the experiment (including the outcome of the competition) *after* JOTP was over and they were no longer with their platoon. These conditions ensured that the experiment was truly one-shot, and that defection was the optimal choice for a selfish individual. For example, subjects did not need to fear reprisal after the experiment if they chose to defect; no-one even knew if someone had or had not defected, until after training was over, and on top of that, identities were anonymous even when that information was revealed. Points earned were converted into Swiss Francs (one point = 0.25 CHF) and the subjects earned on average CHF 14.4 (approximately \$14). There was no show-up fee.

Overall, 525 subjects participated in the experiments: 228 in the ‘Neutral Group Environment’ treatment and 297 in the ‘Competitive Group Environment’ treatment. The different treatments were conducted in different sessions. 281 were assigned the role of A-players and participated in Experiment 1. Half were assigned to the in-group treatment, and half to the out-group treatment. In the few cases in which the groups had an uneven number of A-players, we randomly used the action of some A-players twice to calculate payoffs. The payoff of these players was determined by the decisions associated with the first match. After participating in Experiment 1, these same subjects participated as A-players in Experiment 2. This procedure introduces a possible order effect for the A-players when looking at choices in the second experiment (A-players did not know about the second experiment, however, when making choices in the first experiment). Choices of the A-players in Experiment 2 are not of interest for our purposes, however, as we analyze cooperation of A-players in Experiment 1, and norm enforcement of B-players in Experiment 2. 244 subjects were assigned the role of B-players. They participated only in Experiment 2, and were assigned to one of four conditions (see Figure 1). We elicited B-players’ deduction points using the strategy method, i.e., they specified how many points to deduct from their associated A-player for each possible combination of actions by A1 and A2.

3 Behavioral Hypotheses

This section develops behavioral hypotheses on how the competitive environment might affect cooperation and punishment behavior. If individuals do not have (group-specific) pro-social

preferences, individuals will always defect in the PD game, since this is a dominant strategy. Similarly for punishment, a selfish individual would never punish another player as punishment is costly and there is no benefit of punishment in this one-shot interaction.

The competitive environment, i.e. the small bonus in the *CG* treatment, does not change the predictions for a selfish player. The intuition is straightforward: Cooperation never leads to an increased payoff, because it costs 20 points, and the bonus is only 20 points. In fact, our rules for tie-breaking in case two groups have the same number of points imply that individuals always must expect to lose money when cooperating, because the bonus is only 10 in expected terms. Thus, adding competition cannot generate an increase in cooperation rates through selfish incentives; an increase in cooperation under competition must reflect an effect working through non-selfish motives (see Appendix A for a proof). The competitive environment also has zero impact on punishment choices of a selfish B-player, by construction. The rules of the game are such that the competition is determined without taking into account punishment. Thus, punishment can have no influence on the likelihood of winning the bonus. Hence, our null-hypothesis can be summarized as follows:

H₀ : With selfish players, defection by A-players and no punishment by B-players will be the dominant strategies - both in NG and in CG.

Of course, past research has shown that people are not only willing to cooperate and to punish (for surveys, see Fehr and Schmidt, 2003; Meier, 2007) but that they have group specific social preferences (for evidence with minimal groups, see, e.g., Chen and Li, 2009). With group-specific social preferences, a competitive environment can change individuals' behavior. It has been shown that inter-group competition increases intra-group cooperation and coordination within "minimal" groups (Bornstein and Ben-Yossef, 1994; Bornstein *et al.*, 2002) and real, self-selected, groups (Augenblick and Cunha, 2009).

Recent evolutionary models provide an explanation for how such group-specific social preferences can survive. In particular, the idea of (cultural) group selection implies that a pattern of altruistic cooperation, and altruistic punishment of defectors, can emerge within groups. These altruistic behaviors can survive because they enhance group fitness, and make groups composed of altruists more likely to survive environmental shocks (Henrich, 2004; Boyd *et al.*, 2003). Crucially, altruism must be parochial, or preferentially directed towards own group members, otherwise altruistic groups lose their relative fitness advantage.

Conflict between groups can emerge, however, with the introduction of competition for resources between groups. In this case the seemingly benign trait of altruism can play a surprising role, because enhancing own-group fitness is not the only way to win: damaging competitor groups is also a viable strategy. In addition to being even more cooperative within their group, altruists might become hostile towards other groups, and use anti-social punishment as a way to damage

outsiders. This taste for harming the out-group could survive because it reinforces the relative fitness advantage of groups with altruists (Choi and Bowles, 2007; Bowles, 2009). These theories thus predict that cooperation should be reinforced by competition:

H₁ : A-players with group-specific social preferences cooperate more often with in-group members in CG than in NG.

These theories also predict that punishment of the out-group should become stronger under competition, because hurting the other group generates an advantage for one’s own group. This should reflect a taste for harming the out-group, because such non-strategic motives help solve the free riding problem associated with the “public good” of punishing the out-group. In our experiment, competition did not affect the material consequences of punishment, and thus using punishment to harm the out-group must reflect a taste for hostility rather than strategic motives. While the deeper evolutionary rationale for such hostility, at the group level, depends on this helping win the competition between groups, it is likely that the cue of competition is already sufficient to trigger an individual’s taste for harming the out-group. Previous evidence suggests that a mere cue of competition can create aggression that spills over into relations between groups that are completely unrelated to the competition itself.⁶ These results are corroborated by recent studies showing that a threat to one’s group identity can affect behavior in domains that are unrelated to the threat (Inzlicht and Kang, 2010), and that emotions of anger can create cognitive bias against an outgroup (Dasgupta *et al.*, 2009). A cue of competition is likely interpreted as a threat to one’s group, and likely also increases anger towards the other group. Thus, even though the competition introduced in *CG* does not affect the material consequences of punishment, and has no strategic value to the individual, it may still spill over into the social preferences underlying punishment, as it provides a direct means of hurting the other group.⁷ We thus put forward:

H₂ : B-players with group-specific social preferences punish out-group members more often in CG than in NG.

⁶For example, Sherif *et al.* (1961) report that the equivalent of war broke out between the two groups of boys in the camp, even though competition was confined to a tug-of-war competition. The two groups stole household items from each other, which clearly do not present an advantage in a tug-of-war.

⁷There are many examples where emotional reactions are triggered by cues in the environment, before or in spite of efforts at a more cognitive evaluation that might lead to a different characterization of the situation than suggested by the cue alone. This may reflect the evolutionarily old origins of such emotional systems, which were programs designed to respond to cues that were reliable indicators in the adaptive environment (Cosmides and Tooby, 1994, 2000).

4 Results

We present the results in two steps: first we analyze the impact of group membership and the presence of competition on cooperative behavior. Second, we show how group boundaries and a competitive environment affect punishment behavior.

4.1 Cooperation and Beliefs About Cooperation

Panel A in Figure 2 shows the fraction of individuals cooperating as a function of the group composition and the presence or absence of competition. In general, the figure shows that individuals are willing to cooperate in the PD and that they exhibit in-group favoritism. In the *NG* baseline, we replicate the usual finding in the literature on group effects: there is a significant and large increase in cooperation if individuals are paired with someone from their own platoon rather than another platoon. In fact, cooperation rates are 18 percentage points higher for within-group interactions than between-group. Notably, the lower cooperation rates with out-group members need not indicate hostility, but might simply indicate less willingness to deviate from the dominant selfish strategy when paired with an out-group member. Turning to the *CG* treatment, we see that favoritism towards the in-group is even more extreme when competition is present, with cooperation rates being 36 percentage points higher in within-group than between-group interactions. Thus, in the presence of competition, the difference between out-group and in-group cooperation rates is especially pronounced. However, our results also show that out-group cooperation is not decreasing in *CG*. Thus, competition does not lead individuals to express hostility toward the out-group by defecting more often. This could reflect the limited “expressive value” of defection; individuals might not see defection as a way to unambiguously express hostility, given that it also coincides with the dominant selfish strategy. As we will see below, a design that includes punishment behavior does reveal evidence of a taste for harming the out-group.

Panel B in Figure 2 shows that the results on cooperation behavior are fully reflected in the individuals’ beliefs; people report that they expect in-group favoritism in *NG*, and significantly greater favoritism in *CG*.

[Figure 2 and Table 2 about here.]

The results in Figure 2 are also confirmed in logit models of the following form

$$coop_i = \alpha + \gamma_0 IG_i + \gamma_1 IG_i \times CG_i + \delta CG_i + x_i \beta + e_i \quad (1)$$

where *coop* is an indicator variable equal to 1 if individual *i* cooperates, and zero otherwise. *IG*

is an indicator variable equal to 1 if the individual is paired with another individual from his platoon (*in-group*) and zero if the other player is from another platoon. The indicator variable CG_i is equal to 1 for the ‘Competitive Group Environment’ treatment, and zero otherwise. In some specifications, we also add control variables x , an index of a person’s self-reported trust (explained in Section 2.1), to increase the precision of the estimates. For ease of interpretation, we report marginal effects.

Results in column (1) of Table 2 show that there is a significant overall in-group effect of almost 30 percentage points. Columns (2) and (3) separate the effects of group membership in the two treatments. In the *NG* treatment, cooperation is about 20 percentage points higher if the interaction is in-group ($p = 0.03$ in column (2), and $p = 0.05$ in column (3)). The strength of the in-group effect depends on the presence of competition. The interaction term between *IG* and *CG* shows that the cooperation differential in in-group interactions is about 20 percentage points larger when there is competition ($p = 0.07$ in column (2), and $p = 0.021$ in column (3), where we include an index of trust questions).

We can also estimate OLS models⁸ similar to those above with the dependent variable being bel_{ik} , which is individual i ’s belief about the percentage of individuals cooperating in the two group configurations, k . Because we use two observations per individual, we adjust the error terms by clustering on individuals for possible correlations in e_{ik} within individuals. The results, displayed in columns (4) to (6) of Table 2, show that there is a strong overall in-group effect in beliefs, of almost the identical magnitude as observed in behavior ($p < 0.01$; column (4)). We then separate the in-group effect in the two environments. Beliefs about cooperation are significantly higher for in-group pairings in *NG*, about the same magnitude as we find for behavior. There is a significant interaction with the economic environment: The in-group differential is 13 percentage points larger in *CG* than in *NG* ($p < 0.01$ in both specifications). All in-group differentials in beliefs are within a standard deviation of the in-group differentials in cooperation, showing that the individuals had well-calibrated beliefs.

In sum, group membership *per se* creates in-group favoritism, i.e. individuals cooperate more with in-group members than with out-group members. This effect is also reflected in people’s beliefs. Randomly adding competition between the groups increases the in-group favoritism even though it does not change the predictions for individuals under the assumption of selfish preferences. This indicates that a competitive environment has an impact of group-specific social preferences. Importantly, competition increases in-group cooperation without reducing out-group cooperation. Thus, just looking at cooperation, one would conclude that competition between groups increases social efficiency.

⁸Estimating the same specification with tobit models does not change the results. Results are available from the authors upon request.

4.2 Punishment

We now turn to the analysis of B-players' punishment behavior. Figure 3 displays the results for punishment in situations in which A2 cooperated. The figure allows us to highlight two distinct motives related to the group membership. By varying the identity of A1, the person who can be punished, we can see if punishment depends on whether A1 was a member of the punisher's own group (dark lines) or another group (grey lines). By varying the identity of A2, the player who is the potential victim of defection, we can examine if punishment of A1 depends on whether the victim of defection was from the punisher's group (solid lines) or some other group (dashed lines). The figure also distinguishes between whether A1 cooperated or defected.

Panel A displays the results for the *NG* baseline. There is a clear pattern of norm enforcement in the data: A1 is punished more strongly for defection than cooperation. Punishment of A1 also depends on the identity of A2. If A1 defects, the solid lines (A2 from the punisher's group) are always above the dashed lines (A2 from another group). Thus, individuals are especially prone to punish defection if the "victim" of defection is from the in-group. These results are consistent with the prediction that punishers engage in altruistic punishment in a way that enforces a norm of cooperation toward members of their own group. They also mirror the in-group favoritism observed for cooperation behavior. It is also evident from the figure that the identity of A2 does not matter if A1 cooperates. This indicates a lack of hostility in this treatment. Hostility would imply stronger punishment of an A1 that belongs to another group, regardless of what A1 does. As can be seen in the graph, there is essentially no difference as a function of A1's group affiliation. In sum, group boundaries *per se* do not create hostility in punishment.

[Figure 3 about here.]

Turning to the competition treatment in panel B, we see that the punishment choices are starkly different. Most importantly, there is now a clear difference in punishment depending on whether A1 belongs to the punisher's own group or not. Grey lines (A1 is from another group) are clearly above the dark lines (A1 is from the punisher's group). Thus, out-group individuals are punished significantly harder than in-group members, and importantly, this is true no matter whether the individual cooperates or defects (grey lines are above dark lines in both cases). Thus, the introduction of competition leads to conflict between groups, in the form of anti-social punishment or hostility. Furthermore, there is no relationship between the identity of A2 and punishment in *CG*, so the tendency to preferentially punish defection against the in-group is no longer present.

The two different conditions (neutral and competitive) generated qualitatively different patterns of punishment, as is evident in the figure. A formal statistical test confirms this impression: We

estimate the following OLS regressions:⁹

$$PP_{ik} = \alpha + \gamma_1 I_i(\text{A1 out-group}) + \gamma_2 I_i(\text{A2 in-group}) + e_i \quad (2)$$

where PP are the punishment points that individual i assigns in case k . We include two indicator variables to capture the effect of the group composition on i 's punishment; $I(\text{A2 in-group})$ is equal to 1 if player A2 is from the same group as B1 and 0 otherwise and $I(\text{A1 out-group})$ is equal to 1 if player A1 is from another group as B1 and 0 otherwise.

We estimate equation (2) separately for the two cases where A1 cooperates and the two cases where A1 defects, and estimate these again separately for NG and CG . The coefficients across columns are compared in the bottom panel using two-sided χ^2 -tests (see Appendix B for a formal expression of the tests). Table 3 displays the result for the case in which A2 cooperates. The table shows that in the NG treatment, i.e. in the neutral environment, we find stronger punishment of defection against a member of one's own group, i.e. A2 is an in-group member ($p = 0.05$, column(1)), but no effect of the identity of A1 on punishment ($p = 0.85$, column (1)). Importantly, in the NG environment, there is no effect of the group composition on punishment of cooperation (column (3)). In contrast, in CG , we observe a different pattern in punishment. This can be seen in columns (2) and (4) of Table 3. The identity of A2 is no longer significant. However, A1 gets punished more heavily, whether he cooperates or defects, if he is from a different group than the punisher, i.e. A1 is out-group ($p < 0.01$ in columns (2) and (4)). Hence, there is substantial anti-social punishment.

The comparison of the coefficients across columns, i.e. the neutral vs. the competitive environment, in the lower panel of Table 3 shows that in treatment CG , A1 is punished more heavily than in NG if he is out-group ($p < 0.01$). This is true for whether A1 defects or cooperates. The different punishment pattern conditional on the identity of A2 is not statistically significant. In sum, we clearly reject the hypothesis that the effect of A1's and A2's group affiliation on punishment are the same across the two treatments ($p < 0.01$ for both, defection and cooperation of A1). And the results show substantial anti-social punishment in the CG treatment.

[Table 3 about here.]

The results are also robust (qualitatively the same, but slightly weaker) to adding the cases in which A2 defected. Obviously, in the case in which A1 cooperates and A2 defects, A1 has a payoff of zero and punishment can't reduce his payoff further. So due to censoring there is no punishment in this case – even in CG . Nevertheless, individuals exhibit hostility in punishment also in the case when A1 defects and A2 defects. For reasons of succinctness, the detailed results have been relegated to the Appendix (see Figure A1 and Table A1 there).

⁹The results are maintained in tobit regressions and can be obtained upon request.

Figure 3 additionally indicates that norm enforcement might be weaker in *CG* than in *NG*, i.e. that punishment seems to depend less on A1's behavior in *CG* than in *NG*: While in Panel A of Figure 3 punishment is clearly higher when A1 defects, regardless of the group composition, that relationship is almost completely muted in Panel B. In order to examine the differential in punishment between cooperation and defection, we estimate for each treatment the following equations:

$$PP_{ik} = \alpha + \gamma_1 I_i(\text{A1 out-group}) + \gamma_2 I_i(\text{A2 in-group}) + \gamma_3 I_k(\text{A1 defects}) + e_i \quad (3)$$

$$PP_{ik} = \alpha + \gamma_1 I_i(\text{A1 out-group}) + \gamma_2 I_i(\text{A2 in-group}) + \gamma_4 I_k(\text{A1 defects}) \times I_i(\text{A2 out-group}) + \gamma_5 I_k(\text{A1 defects}) \times I_i(\text{A2 in-group}) + e_i \quad (4)$$

in which *A1 defects* equals 1 if A1 defects and 0 otherwise. Regression 4 adds two interaction terms for *A1 defects* and whether A2 is an out-group member or an in-group member.

The results (displayed in Table 4) show that defection is more strongly punished than cooperation in *NG* ($p < 0.01$, column (1)). This effect is less strong in *CG* as seen in column (2). A formal test comparing the two coefficients, γ_3 , shows that the difference is significant ($p < 0.01$) and conditioning of punishment on actions of A1 is much weaker in *CG*, while it is still significant. As can be seen from comparing coefficient γ_4 across columns (3) and (4), this effect of competition in *CG* even prevails when A2 is from the punisher's own group ($p = 0.03$), the case when norm enforcement was strongest in *NG*. Finally, it is noteworthy from column (4) that the norm enforcement pattern of punishment if defection is directed against one's own group has completely disappeared in the competitive environment. The fact that individuals cease to use punishment to enforce cooperation among in-group members in a competitive environment might seem surprising, from the perspective of increasing group fitness. However, our results on cooperation showed that competition leads to very high within-group cooperation rates, even in the absence of punishment threat, mitigating this problem for the group. In summary, competition causes punishment to stop functioning as a tool for norm enforcement, and instead to take the form of anti-social punishment directed towards the out-group, consistent with predictions given group-specific social preferences.

[Table 4 about here.]

5 Conclusions

Recent research has shown that people are willing to engage in costly punishment. This can be highly beneficial when punishment is directed selectively at defectors, making it possible to sustain cooperation even in the face of strong free-riding incentives (e.g., Fehr and Gaechter, 2000). These findings have important implications for organizations, given that contracts are typically incomplete and potential free-riding problems abound; altruistic norm enforcement can thus help firms avoid inefficient outcomes (Gaechter *et al.*, 2008). Results from this paper and others (Chen and Li, 2009) also suggest that the division of organizations into groups can affect norm enforcement, and can even enhance the willingness to punish defection in the case that defection is against in-group members (Goette *et al.*, 2006).

The novel result of this paper is to show that a simple change in economic environment, interacting with group boundaries, can not only eliminate willingness to enforce norms, but in fact generate a taste for destructive punishment. In particular, the introduction of competition between organizational groups triggers strong anti-social punishment directed towards outsiders, such that they are harmed even when they cooperate. Because cooperation is punished as hard as defection, punishment ceases to deter defection, while at the same time destroying resources. This anti-social punishment reflects a taste for harming the out-group, generated by cues of competition alone; our design excludes any strategic motive for punishing outsiders, i.e., there is no competitive advantage from harming the out-group. Also, the willingness to engage in norm enforcement on behalf of one's own group is reduced by competition. Thus, our results provide an important caveat to the view that punishment behavior is beneficial, and identify a mechanism that can cause costly conflict within organizations.

While competition can be harmful in terms of leading to anti-social punishment between groups, we also show that competition can have a positive effect on within-group cooperation, which is also fully reflected in individuals' beliefs. Thus, group cohesion is improved when there is a sense of competition with other groups. The fact that competition is a catalyst for both harmful and efficiency-enhancing motives within organizations means that firms face a complicated decision when thinking about introducing incentives for competitions between teams.

Our results also provide additional evidence that social preferences are endogenous to the economic environment (Bowles, 1998; Burks *et al.*, 2009). This literature argues that changes in economic environments bring about changes in preferences. Yet, these changes are typically assumed to be slow, e.g., operating through slow-changing norms of cooperation (Herrmann *et al.*, 2008). It is noteworthy that we observe a particularly strong form of endogenous preferences: Our treatments are between-subject manipulations, and still we immediately observe starkly different punishment and cooperation strategies, conditional on the economic environment. Thus, this evidence suggests that different motives of social preferences may be dormant in humans,

and triggered by different economic environments.

References

- ABBINK, K., BRANDTS, J., HERRMANN, B. and ORZEN, H. (2010). Intergroup Conflict and Intra-Group Punishment in an Experimental Contest Game. *The American Economic Review*, **100** (1), 420–447.
- AKERLOF, G. and KRANTON, R. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives*, **19**, 9 – 32.
- AKERLOF, G. A. and KRANTON, R. E. (2000). Economics and identity. *Quarterly Journal of Economics*, **115** (3), 715–753.
- ANDREONI, J. (1995). Cooperation in public-goods experiments: kindness or confusion? *American Economic Review*, pp. 891–904.
- AUGENBLICK, N. and CUNHA, J. (2009). Using competition to elicit cooperation in a political public goods game: A field experiment. *Working Paper*.
- AULETTA, K. (1985). The fall of Lehman Brothers: Part 2. The men, the money, the merger. *The New York Times Magazine*, **24**, 36–42.
- BAHRY, D., KOSOLAPOV, M., KOZYREVA, P. and WILSON, R. (2005). Ethnicity and trust: Evidence from Russia. *American Political Science Review*, **99** (04), 521–532.
- BERNHARD, H., FEHR, E. and FISCHBACHER, U. (2006). Group affiliation and altruistic norm enforcement. *American Economic Review*, **96** (2), 217–221.
- BORNSTEIN, G. and BEN-YOSSEF, M. (1994). Cooperation in inter-group and single-group social dilemmas. *Journal of Experimental Social Psychology*, **30**, 52–52.
- , GNEEZY, U. and NAGEL, R. (2002). The effect of intergroup competition on group coordination: an experimental study. *Games and Economic Behavior*, **41** (1), 1–25.
- BOWLES, S. (1998). Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, **36** (1), 75–111.
- (2009). Did Warfare Among Ancestral Hunter-Gatherers Affect the Evolution of Human Social Behaviors? *Science*, **324** (5932), 1293–1298.
- BOYD, R., GINTIS, H., BOWLES, S. and RICHEYSON, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, **100** (6), 3531–3535.
- BURKS, S. V., CARPENTER, J. P. and GOETTE, L. (2009). Performance pay and worker cooperation: Evidence from an artefactual field experiment. *Journal of Economic Behavior and Organization*.
- CHARNESS, G., RIGOTTI, L. and RUSTICHINI, A. (2007). Individual behavior and group membership. *American Economic Review*, **97** (4), 1340–52.
- CHEN, R. and CHEN, Y. (2011). The Potential of Social Identity for Equilibrium Selection. *American Economic Review*, p. Forthcoming.
- CHEN, Y. and LI, S. X. (2009). Group identity and social preferences. *American Economic Review*, **99** (1).
- CHOI, J.-K. and BOWLES, S. (2007). The coevolution of parochial altruism and war. *Science*, **318** (5850), 636.

- COSMIDES, L. and TOOBY, J. (1994). Better than rational: Evolutionary psychology and the invisible hand. *The American Economic Review*, **84** (2), 327–332.
- and — (2000). Evolutionary psychology and the emotions. *Handbook of emotions*, **2**, 91–115.
- DASGUPTA, N., DESTENO, D., WILLIAMS, L. A. and HUNSINGER, M. (2009). Fanning the flames of prejudice: the influence of specific incidental emotions on implicit prejudice. *Emotion (Washington, D.C.)*, **9** (4), 585–91.
- DREBER, A., RAND, D. G., FUDENBERG, D. and NOWAK, M. A. (2008). Winners don’t punish. *Nature*, **452** (7185), 348–351.
- DURLAUF, S. (1999). The case ‘against’ social capital. *Focus*, **20** (3), 1–5.
- EGAS, M. and RIEDL, A. (2008). The economics of altruistic punishment, and the maintenance of cooperation. *Proceedings of the Royal Society*, **275**, 871 – 878.
- ESSES, V. M., JACKSON, L. M. and ARMSTRONG, T. L. (1998). Intergroup competition and attitudes toward immigrants and immigration: An instrumental model of group conflict. *Journal of Social Issues*, **54** (4), 699–724.
- FEHR, E. and FISCHBACHER, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior*, **25** (2), 63–87.
- and GAECHTER, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, **90** (4), 980–994.
- and SCHMIDT, K. (2003). Theories of fairness and reciprocity - evidence and economic application. In M. Dewatripont, L. P. Hansen and S. J. Turnovsky (eds.), *Advances in Economics and Econometrics - 8th World Congress, Econometric Society Monographs*, Cambridge: Cambridge University Press, pp. 208–57.
- FERSHTMAN, C. and GNEEZY, U. (2001). Discrimination in a segmented society: An experimental approach. *Quarterly Journal of Economics*, **116** (1), 351–377.
- FISCHBACHER, U. and GAECHTER, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, **100** (1), 541–556.
- FUSTER, A. and MEIER, S. (2010). Another hidden cost of incentives: The detrimental effect on norm enforcement. *Management Science*, **56** (1), 57–70.
- GAECHTER, S., RENNER, E. and SEFTON, M. (2008). The long-run benefits of punishment. *Science*, **322** (5907), 1510.
- GOETTE, L., HUFFMAN, D. and MEIER, S. (2006). The impact of group membership on cooperation and norm enforcement: Evidence using random assignment to real social groups. *American Economic Review (Papers & Proceedings)*, **96** (2), 212 – 216.
- , — and — (2010). The impact of social ties on group interactions: Evidence from minimal groups and randomly assigned real groups. *Working Paper*.
- GÜTH, W., SCHMITTBERGER, R. and SCHWARZE, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, **3** (4), 367–388.
- HABYARIMANA, J., HUMPHREYS, M., POSNER, D. and WEINSTEIN, J. (2007). Why does ethnic diversity undermine public goods provision? *American Political Science Review*, **101** (04), 709–725.

- HALEVY, N., BORNSTEIN, G. and SAGIV, L. (2008). ‘In-Group Love’ and ‘Out-Group Hate’ as Motives for Individual Participation in Intergroup Conflict: A New Game Paradigm. *Psychological Science*, **19** (4), 405–411.
- HARBRING, C. and IRLBUSCH, B. (2011). Sabotage in tournaments: Evidence from a laboratory experiment. *Management Science*, **57** (4), 611–627.
- HENRICH, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*, **53** (1), 3–35.
- HERRMANN, B., THOENI, C. and GAECHTER, S. (2008). Antisocial Punishment Across Societies. *Science*, **319** (5868), 1362–1367.
- HEWSTONE, M., RUBIN, M. and WILLIS, H. (2002). Intergroup bias. *Annual Review of Psychology*, **53**, 575 – 604.
- HOUSER, D., XIAO, E., MCCABE, K. and SMITH, V. (2008). When punishment fails: Research on sanctions, intentions and non-cooperation. *Games and Economic Behavior*, **62** (2), 509–532.
- INZLICHT, M. and KANG, S. K. (2010). Stereotype threat spillover: how coping with threats to social identity affects aggression, eating, decision making, and attention. *Journal of personality and social psychology*, **99** (3), 467–81.
- JANN, B. (2003). Old-boy network: Military service and professional success in civilian life in switzerland. *Zeitschrift fuer Soziologie*, **32** (2), 139–155.
- KANDEL, E. and LAZEAR, E. P. (1992). Peer Pressure and Partnerships. *Journal of Political Economy*, pp. 801–817.
- KNEZ, M. and SIMESTER, D. (2001). Firm-wide incentives and mutual monitoring at Continental Airlines. *Journal of Labor Economics*, **19** (4), 743–772.
- LEIDER, S., MÖBIUS, M., ROSENBLAT, T. and DO, Q.-A. (2009). Directed altruism and enforced reciprocity in social networks. *Quarterly Journal of Economics*, **124** (4), 1815–1851.
- MACLEOD, W. B. (2007). Can contract theory explain social preferences? *American Economic Review*, **97** (2), 187–192.
- MAS, A. and MORETTI, E. (2009). Peers at work. *American Economic Review*, p. Forthcoming.
- MEIER, S. (2007). A survey on economic theories and field evidence on pro-social behavior. In B. S. Frey and A. Stutzer (eds.), *Economics and Psychology: A Promising New Cross-Disciplinary Field*, Cambridge: MIT Press, pp. 51–88.
- OSTROM, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.
- POSNER, D. (2004). The political salience of cultural difference: Why Chewas and Tumbukas are allies in Zambia and adversaries in Malawi. *American Political Science Review*, **98** (4), 529–545.
- SHERIF, M., HARVEY, O. J., WHITE, J., HOOD, W. and SHERIF, C. (1961). *Intergroup conflict and cooperation: The robbers cave experiment*. Norman, OK: University of Oklahoma Book Exchange.
- SUTTER, M. (2009). Individual behavior and group membership: Comment. *American Economic Review*, **99** (5), 2247–57.

TAJFEL, H., BILLIG, M., BUNDY, R. and FLAMENT, C. (1971). Social categorization and inter-group behavior. *European Journal of Social Psychology*, **1**, 149–177.

WILLIAMSON, O. E., WACHTER, M. L. and HARRIS, J. E. (1975). Understanding the employment relation: The analysis of idiosyncratic exchange. *The Bell Journal of Economics*, **6** (1), 250–278.

Tables and Figures

Table 1: Off-Duty Time Spent per Week

	Own Platoon	Other Platoon
Less than once	4.5%	30.4%
Between 1 and 2 times	45.0%	44.0%
Twice or more	50.5%	25.7%
N	489	491

Note: The two distributions are significantly different (Wilcoxon Signed-Rank test, $p < 0.001$).

Table 2: Results for Cooperation Rates and Beliefs about Cooperation

Dependent Variable:	Cooperation (=1)			Beliefs: % cooperating		
	(1)	(2)	(3)	(4)	(5)	(6)
Ingroup (=1)	0.28*** (0.05)	0.18** (0.08)	0.17* (0.09)	0.26*** (0.02)	0.18*** (0.02)	0.18*** (0.02)
CG (=1)	0.08 (0.06)	0.00 (0.08)	0.01 (0.08)	0.01 (0.03)	-0.06* (0.03)	-0.06* (0.03)
Ingroup×CG		0.18* (0.10)	0.22** (0.10)		0.13*** (0.03)	0.13*** (0.04)
Trust			0.06*** (0.02)			0.03*** (0.01)
Constant				0.37*** (0.02)	0.41*** (0.02)	0.34*** (0.04)
(Pseudo)- R^2	0.07	0.08	0.10	0.20	0.21	0.25
# of observations	281	281	267	538	538	515
# of individuals	281	281	267	274	274	262

Notes: In column (1) to (3), marginal effects from logit models. In columns (3) and (4) coefficients from OLS models. The model in columns (4) to (6) uses two observations per individual (if available), therefore standard errors of the estimates in column (4) to (6) are adjusted for clustering on individuals.

Level of significance: * : $0.05 \leq p < 0.1$, ** : $0.01 \leq p < 0.05$, *** : $p < 0.01$.

Table 3: Punishment as a Function of Group Membership

Behavior of A1: Environment:	(1)	(2)	(3)	(4)
	A1 defects		A1 cooperates	
	Neutral	Comp	Neutral	Comp
A1 out-group (γ_1)	0.155 (0.853)	3.742*** (0.658)	-0.099 (0.693)	2.898*** (0.684)
A2 in-group (γ_2)	1.694** (0.840)	0.868 (0.676)	0.535 (0.697)	-0.544 (0.693)
Constant	4.487*** (0.705)	1.636*** (0.447)	2.307*** (0.578)	1.988*** (0.517)
R^2	0.039	0.203	0.005	0.125
# of observations/individuals	111	132	111	132
<i>Tests across equations (environments):</i>				
Test that γ_1 differs	$p < 0.01$		$p < 0.01$	
Test that γ_2 differs	$p = 0.44$		$p = 0.27$	
Test that γ_1 and γ_2 differ	$p < 0.01$		$p < 0.01$	

Note: Dependent variable: # of deduction points. OLS estimates for the cases in which A2 cooperates. Robust standard errors in parentheses. p -values in cross-equation tests are all two-sided.


Level of significance: * : $0.05 \leq p < 0.1$, ** : $0.01 \leq p < 0.05$, *** : $p < 0.01$.

Table 4: Norm Enforcement Across Environments

Environment:	Neutral (1)	Comp (2)	Neutral (3)	Comp (4)
A1 out-group (γ_1)	0.028 (0.624)	3.320*** (0.589)	0.028 (0.625)	3.320*** (0.590)
A2 in-group (γ_2)	1.114* (0.609)	0.162 (0.595)	0.513 (0.691)	-0.533 (0.694)
A1 defects (γ_3)	2.838*** (0.445)	0.682** (0.330)		
A1 defects \times A2 out-group (γ_4)			3.510*** (0.722)	1.482*** (0.560)
A1 defects \times A2 in-group (γ_5)			2.306*** (0.549)	0.092 (0.385)
Constant	1.978*** (0.565)	1.471*** (0.465)	2.244*** (0.578)	1.766*** (0.495)
R^2	0.130	0.161	0.136	0.167
# of observations	222	264	222	264
# of individuals	111	132	111	132
<i>Tests across equations (environments):</i>				
Test that γ_1 differs		$p < 0.01$		$p < 0.01$
Test that γ_2 differs		$p = 0.26$		$p = 0.29$
Test that γ_3 differs		$p < 0.01$		
Test that γ_4 differs				$p = 0.03$
Test that γ_5 differs				$p < 0.01$
Test that $\gamma_1, \gamma_2,$ and γ_3 differ		$p < 0.01$		
Test that $\gamma_1, \gamma_2, \gamma_4,$ and γ_5 differs				$p < 0.01$


Note: Dependent variable: # of deduction points. OLS estimates. Robust standard errors clustered on the individual in parentheses.

Level of significance: * : $0.05 \leq p < 0.1$, ** : $0.01 \leq p < 0.05$, *** : $p < 0.01$. p -values in cross-equation tests are all two-sided.



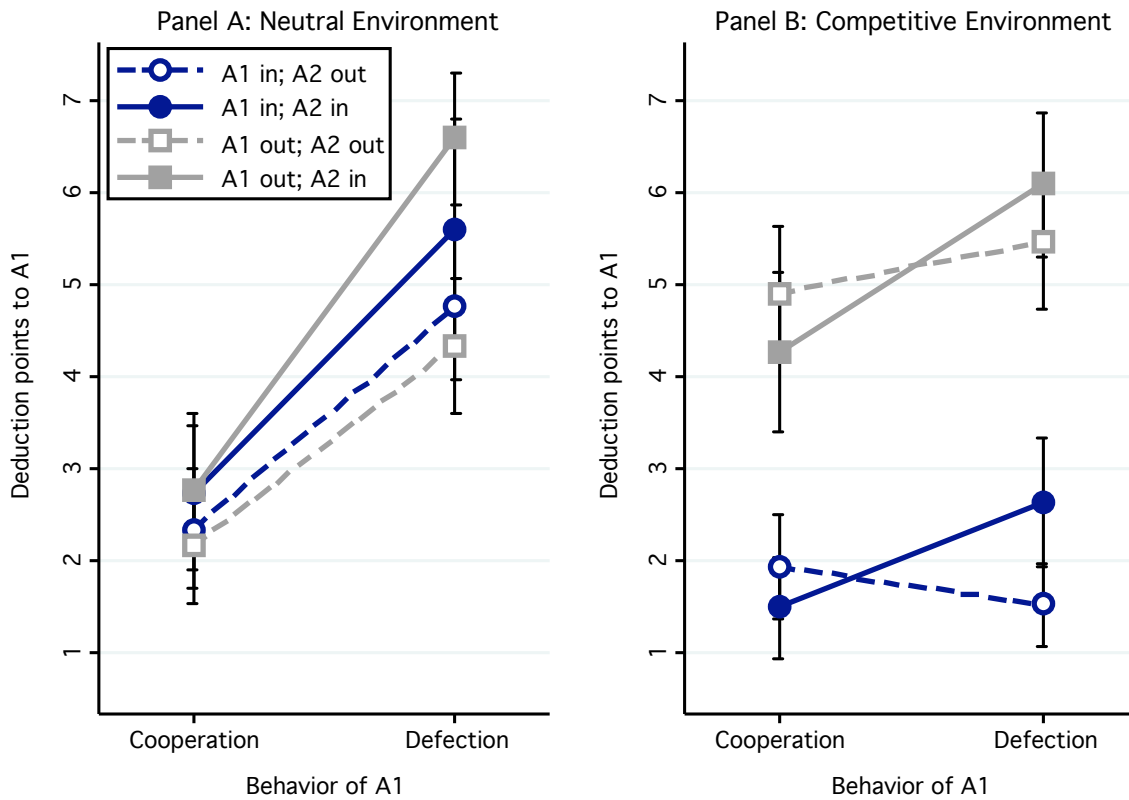
Note: The game allowed B1 to punish A1, and B2 to punish A2, conditional on the actions and A1 and A2 in a simultaneous prisoners' dilemma game. The dark shading indicates the four possible group combinations for B1, A1, and A2, which were implemented as different treatments (players with the same shading are from the same group). The design deliberately did not vary all possible combinations of B1 and B2 group roles, because of number of observations, so the effect of B2 group identity on B1 behavior is not studied. The pattern of B-player (and A-player) group compositions was identical across the Economic Environments, the *NG* and *CG* treatments.

Figure 1: Group Composition in Third-Party Punishment Game



Notes: The bars show standard errors of the mean. White indicates an out-group pairing while dark indicates an in-group pairing. Panel A shows the fraction of A-players passing their endowment to the other player in a simultaneous one-shot PD. Panel B shows beliefs about the fraction of A-players who pass their endowment in the PD.

Figure 2: Cooperation rates and beliefs about cooperation in *NG* and *CT*



Notes: B-players could deduct between 0 and 10 points. Each deduction point costs B-players 1 point and A1-players 3 points. Deduction points were made conditional on whether A1-players cooperated or defected using the strategy method. Error bars show standard errors of the mean.

Figure 3: Punishment in the case A2-player cooperated

Appendix

A Nash Equilibrium in the PD in *Competitive Group Environment*

In this section, we explain why it was optimal for selfish A-players to defect in the Competitive Group Environment. The intuition is straightforward: Cooperating never leads to an increased payoff, because cooperating costs 20 points, and the bonus is only 20 points. In fact, our rules for tie-breaking in case two groups have the same number of points imply that individuals always lose money when cooperating, because the bonus is only 10 in expected terms. Thus, adding competition cannot generate an increase in cooperation rates through selfish incentives; an increase in cooperation under competition must reflect an affect working through non-selfish motives. Below, we formalize the intuition that the Nash equilibrium doesn't change between *NG* and *CG*.

We show that the Nash equilibrium in the game involves all (selfish) A-players defecting.

(i) In our experiment, there are within-group and between-group pairings. Obviously, a selfish player never cooperates with a player from another group, since, on top of costing him 20 points, he may also be pivotal in losing the bonus. Therefore, what remains to be considered are within-group pairings. First consider the case of $K = 2$ groups, denoted X and Y . Now pick an arbitrary collection of strategies in which some individuals cooperate in within-group pairings. We ask whether this strategy can be a Nash equilibrium. Two possible cases can arise: Either one of the groups, say group Y , loses, or the two groups tie.

We first show that groups can never tie with some individuals cooperating.

- Pick an arbitrary member of group k who is cooperating. Since the groups are tied, he wins a bonus with probability 0.5. If he defects, his group will lose for sure. However, defecting saves 20 points, while costing only 10 points in expected bonus. Thus, when two groups are tied, cooperating players have an incentive to defect.

We now show that it is impossible to have a Nash equilibrium in which group Y loses for sure.

- If group Y loses, then it cannot be a Nash equilibrium for anyone in group Y to cooperate. Given the others' strategies, members of Y who cooperate can increase their payoff by 20 points if they defect.
- Given this result, it follows that in group X , at most one player will cooperate. If more than one player in group X cooperated, a player could switch to defection while still winning the bonus, holding the other players' strategies constant.
- However, if one player in X cooperates, the tying rule now implies that the player can defect, and save 20 points, but only lose 10 points in expected bonus (since the two groups now tie).

Thus, the only equilibrium for $K = 2$ groups involves both groups tying, and this equilibrium involves all players defecting.

(iii) The above arguments immediately generalize to $K > 2$ groups. The only difference is that the expected bonus in the case of a tie will be even smaller, $20/m$, where $m \leq K$ is the number of groups tying. Thus, the same reasoning applies.

B Comparisons across Equations

The bottom panels of Table 3 to 4 also display cross-equation tests. For single-coefficient tests, we calculate

$$z = \frac{\gamma_j - \tilde{\gamma}_j}{\sqrt{\Sigma_{jj} + \tilde{\Sigma}_{jj}}} \quad (5)$$

where γ_j and $\tilde{\gamma}_j$ are the two coefficients of interest from the two equations, and Σ_{jj} and $\tilde{\Sigma}_{jj}$ are the corresponding main diagonal elements in the covariance matrix (because the two coefficients come from two separate equations, their covariance, by construction, is zero). z has a standard normal distribution under the null of no difference. We report two-sided p -values to be conservative. In the case of coefficient vectors, we calculate the analogous test statistic

$$\chi = (\gamma - \tilde{\gamma})(\Sigma + \tilde{\Sigma})^{-1}(\gamma - \tilde{\gamma})' \quad (6)$$

which has a chi-square distribution with k degrees of freedom, where k is the number of variables in γ .

C Appendix with Experimental Instructions

C.1 Instructions Player A (Translation)

What is this about?

Two subjects participate in this decision situation. They will be called A1 and A2. Both, A1 and A2, will get an endowment of 20 points. Each participants has to decide between two options:

- **Keep:** The participant keeps his 20 points.
- **Transfer:** The participant transfers his 20 points to the other participants. The transfered points will be doubled.

Each participant has to decide whether to Keep or the Transfer without knowing how the other participant decided. So, the following payoffs can result:

		Payoffs in this case:
Case 1:	A1 keeps the points	A1: 20 points
	A2 keeps the points	A2: 20 points
		Payoffs in this case:
Case 2:	A1 transfers 20 points	A1: 0 points
	A2 keeps the points	A2: 60 points
		Payoffs in this case:
Case 3:	A2 keeps the points	A1: 60 points
	A2 transfers 20 points	A2: 0 points
		Payoffs in this case:
Case 4:	A1 transfers 20 points	A1: 40 points
	A2 transfers 20 points	A2: 40 points

How will you decide?

- You will be in the role of A1.
- Your assigned **participant A2 is from another platoon.**

None of the participants will ever find out to whom he was assigned. We guarantee total anonymity. When all the participants reached a decision, we will calculate the points and the resulting monetary payoffs in the following way:

4 points = CHF1

The amount will be delivered to you by mail.

Everything clear?

Before you decide, answer the following questions. The question make sure that all the participants understand the instructions.

If you have questions, please contact the staff.


1. A1 and A2 keep their points. Please calculate the resulting points for all participants. State all the steps in getting to the result.
2. A1 and A2 transfer their points. Please calculate the resulting points for all participants. State all the steps in getting to the result.
3. A1 keeps his points and A2 transfer his points. Please calculate the resulting points for all participants. State all the steps in getting to the result.

Please contact the staff when you are done with the questions or if you have questions.

Decision Sheet

- You were assigned the role of **A1**.
- Your assigned **participant A2 is from another platoon**.

In the following figure are the participants from the other platoons shaded.



Please decide which option to pick:

- Keep
- Transfer

Please let the staff know when you decided.

C.2 Instructions Player B (Translation)

What is this about?

Four subjects participate in this decision situation. They will be called A1, A2, B1 and B2. The decision situation will have two steps.

Step 1: A1 and A2 will get an endowment of 20 points. Each participants has to decide between two options:

- **Keep:** The participant keeps his 20 points.
- **Transfer:** The participant transfers his 20 points to the other participants. The transfered points will be doubled.

Each participant has to decide whether to Keep or the Transfer without knowing how the other participant decided. So, the following payoffs can result:

Step 2: B1 and B2 will get an endowment of 70 points each and A1 and A2 will get another 10 points each. In Step 2, B1 and B2 can assign deduction points. B1 can assigned deduction points to A1 and B2 can assign deduction points to A2. B1 and B2 can each assign a maximum of 10 deduction points.

		Payoffs in this case:
Case 1:	A1 keeps the points	A1: 20 points
	A2 keeps the points	A2: 20 points
		Payoffs in this case:
Case 2:	A1 transfers 20 points	A1: 0 points
	A2 keeps the points	A2: 60 points
		Payoffs in this case:
Case 3:	A2 keeps the points	A1: 60 points
	A2 transfers 20 points	A2: 0 points
		Payoffs in this case:
Case 4:	A1 transfers 20 points	A1: 40 points
	A2 transfers 20 points	A2: 40 points

Before explaining how B1 and B2 will make their decisions, we will describe how deduction points will change the payoffs. **Each deduction point will reduce the payoff of B by one point and the payoff of A by three points.** For example, if B1 assigns 3 deduction points, this will reduce A1's payoff by 9 points and B1's payoff by 3 points.

B1 and B2 will decide about the assignment of deduction points for each potential case in Step 1. That is, they will decide about assigning deduction points for the following four potential cases in Step 1:

- **Case 1:** A1 and A2 keep their points.
- **Case 2:** A1 transfers his points and A2 keeps his points.
- **Case 3:** A1 keeps his points and A2 transfers his points.
- **Case 4:** A1 and A2 transfer their points.

This will lead to the following payoffs:

Payoff of A1 =	Payoff from Step 1 + 10 points from Step 2 - 3*Deduction points from B1
Payoff of A2 =	Payoff from Step 1 + 10 points from Step 2 - 3*Deduction points from B2
Payoff of B1 =	Endowment of 70 points - Deduction points to A1
Payoff of B2 =	Endowment of 70 points - Deduction points to A2

How will you decide?

- You will be in the role of **B1**.
- Your assigned **participant A1 is from another platoon**.
- The participant **A2 is from another platoon**. He got assigned to a participant **B2 from your platoon**.

None of the participants will ever find out to whom he was assigned. We guarantee total anonymity. When all the participants reached a decision, we will calculate the points and the resulting monetary payoffs in the following way:

4 points = CHF1

The amount will be delivered to you by mail.

Everything clear?

Before you decide, answer the following questions. The question make sure that all the participants understand the instructions.

If you have questions, please contact the staff.


1. In Step 1, A1 and A2 keep their points. In Step 2, neither B1 nor B2 assign any deduction points. Please calculate the resulting points for all participants. State all the steps in getting to the result.
2. In Step 1, A1 and A2 transfer their points. In Step 2, neither B1 nor B2 assign any deduction points. Please calculate the resulting points for all participants. State all the steps in getting to the result.
3. In Step 1, A1 keeps his points and A2 transfers his points. In Step 2, B1 assigns 2 deduction points and B2 assigns 5 deduction points. Please calculate the resulting points for all participants. State all the steps in getting to the result.
4. In Step 1, A1 transfers his points and A2 transfers his points. In Step 2, B1 assigns 1 deduction points and B2 assigns 4 deduction points. Please calculate the resulting points for all participants. State all the steps in getting to the result.

Please contact the staff when you are done with the questions or if you have questions.

Decision Sheet


- You will be in the role of **B1**.
- Your assigned **participant A1 is from another platoon**.
- The participant **A2 is from another platoon**. He got assigned to a participant **B2 from your platoon**.

In the following figure are the participants from the other platoons shaded.




Please decide about the assignment of the deduction points for all possible cases. Only the cases that really happen will determine your payoff and the payoff of the other participants. In each of the cases, you can assign between 0 and 10 deduction points.


Case 1:




Case 2:



Case 3:




Case 4:



Please let the staff know when you decided.

D Appendix Tables and Figures



Notes: B-players could deduct between 0 and 10 points. Each deduction point costs B-players 1 point and A1-players 3 points. Deduction points were made conditional on whether A1-players cooperated or defected using the strategy method. Error bars show standard errors of the mean.

Figure A1: Punishment in the case A2-player defected

Table A1: Punishment as a Function of Group Membership

Behavior of A1: Environment:	(1)	(2)	(3)	(4)
	A1 defects		A1 cooperates	
	Neutral	Comp	Neutral	Comp
A1 outgroup (γ_1)	-0.103 (0.661)	3.249*** (0.549)	0.018 (0.524)	1.336*** (0.470)
A2 ingroup (γ_2)	1.445** (0.645)	0.308 (0.558)	0.297 (0.534)	-0.485 (0.469)
A2 defects (=1)	-2.774*** (0.387)	-1.477*** (0.269)	-1.084*** (0.376)	-1.705*** (0.401)
Constant	4.746*** (0.622)	2.127*** (0.420)	2.345*** (0.460)	2.768*** (0.438)
R^2	0.141	0.205	0.028	0.091
# of observations	223	264	223	264
# of individual	112	132	112	132
<i>Tests across equations (environments):</i>				
Test that γ_1 differs	$p < 0.01$		$p = 0.06$	
Test that γ_2 differs	$p = 0.18$		$p = 0.27$	
Test that γ_1 and γ_2 differ	$p < 0.001$		$p = 0.10$	

Note: Dependent variable: # of deduction points. OLS estimates. Robust standard errors clustered on the individual in parentheses. p -values in cross-equation tests are all two-sided.

Level of significance: * : $0.05 \leq p < 0.1$, ** : $0.01 \leq p < 0.05$, *** : $p < 0.01$.