

## SOCIAL SCIENCES

# The origins of human prosociality: Cultural group selection in the workplace and the laboratory

Patrick Francois<sup>1,2\*</sup>, Thomas Fujiwara<sup>2,3,4</sup>, Tanguy van Ypersele<sup>5</sup>

**Human prosociality toward nonkin is ubiquitous and almost unique in the animal kingdom. It remains poorly understood, although a proliferation of theories has arisen to explain it. We present evidence from survey data and laboratory treatment of experimental subjects that is consistent with a set of theories based on group-level selection of cultural norms favoring prosociality. In particular, increases in competition increase trust levels of individuals who (i) work in firms facing more competition, (ii) live in states where competition increases, (iii) move to more competitive industries, and (iv) are placed into groups facing higher competition in a laboratory experiment. The findings provide support for cultural group selection as a contributor to human prosociality.**

## INTRODUCTION

No small part of the spectacular success of the human species is due to our unusually high levels of cooperation among nonrelated individuals. The scale of this cooperation in human nonkin is rare in the animal kingdom, unique among mammals, and strongly at odds with our closest genetic relatives. But the origins and reasons for the continued existence of this prosociality are still an ongoing and important puzzle. Richerson *et al.* (1) and Bowles and Gintis (2) extensively discuss the puzzle of human prosociality.

The variety of theories proposed to explain these phenomena are typically hard to assess empirically. Examples include reciprocal altruism (3), sexual selection (4), and the mismatch hypothesis (5). A reason is that predictions often concern elements of our primordial past, perhaps traceable via the archeological record, or rest on non-observables that are not, for the most part, readily discernible. But a class of theories that can be grouped under the heading of cultural group selection (CGS) provide an exception that we argue will allow us to scrutinize contemporary data for evidence in accord with their predictions. The evidence that we will present is novel and, as Richerson *et al.* (1) argue, rare, in that it is quantitative, although another quantitative study supporting CGS has been argued using observed group extinction rates among tribal groups in Papua New Guinea by (6).

CGS posits that our “social” world coevolved with our social instincts. As a species, we evolved a psychology expecting life to be structured by moral norms, and we developed features designed to learn and internalize norms [see (7) for discussion of the evidence supporting humans as evolved social learners]. By at least 70,000 years ago, most human populations resembled the hunter-gathering societies of the ethnographic record, that is, tribal-scale societies of a few hundred to a few thousand people. Competition across these populations induced selection of group beneficial (prosocial) but individually costly traits (in the form of normative prescriptions or culture). The content of these norms was not fixed, nor were they hard-wired behavioral imperatives, allowing human societies to adapt

norms suited to prevailing conditions. But “selection” occurred as societies with the fitness-enhancing norm/institution combinations proliferated via defeating less successful groups in direct conflict or taking their resources; being imitated by their less successful competitors; or through selective migration and internalization of norms upon migration [see (8) for a more detailed elaboration of these selective forces]. The ones able to generate prosociality “won” the evolutionary battle, and the proliferation of this prosociality today is a reflection of the winners of that battle.

While the narrative in which this explanation for human prosociality is couched is in terms of our prehistoric past, the scope of this paper is to test the predictions of CGS in contemporary settings. Since CGS emphasizes the nonhard-wired features of behavior such as norms, forces of group-level competition should help in sustaining cooperative norms and hence observed prosocial behavior, even in modern contexts.

One way of assessing this implication would be to see whether features that help in sustaining prosociality are more prevalent in groups subject to greater selective pressure—for instance, if more frequent intergroup conflicts increase individually costly group beneficial behavior, such as altruistic punishment [as discussed in (9)]. Instead of testing for a single specific behavior, such as altruistic punishment, another way to proceed would be to see whether groups experiencing more intense intergroup competition exhibit evidence of more prosocial behavior. We attempt that here, in economic settings.

We study individuals who vary in how much competition is experienced across the organizations in which they work. We test to see whether variation in cross-group competitiveness affects a measure of the individuals’ prosociality. We report on a variety of individual data sources, both cross-sectional and at panel (within-person) level to do this. Before turning to the data, we clarify two key aspects: (i) the definition of relevant organizations or “groups” over which CGS may occur and over which competition is to be gauged and (ii) our measure of individual level prosociality.

## MATERIALS AND METHODS

Perhaps the most ubiquitous avenue of group-level competition occurring in contemporary settings is likely to be competition across firms. Individuals within firms need to undertake (at least some) group beneficial but individually costly actions. Moreover, competition across firms affects returns to cooperative versus selfish individual acts and,

Copyright © 2018  
The Authors, some  
rights reserved;  
exclusive licensee  
American Association  
for the Advancement  
of Science. No claim to  
original U.S. Government  
Works. Distributed  
under a Creative  
Commons Attribution  
NonCommercial  
License 4.0 (CC BY-NC).

<sup>1</sup>University of British Columbia, Vancouver School of Economics, Vancouver, British Columbia, Canada. <sup>2</sup>Canadian Institute for Advanced Research, Toronto, Ontario, Canada. <sup>3</sup>Department of Economics, Princeton University, Princeton, NJ 08540, USA. <sup>4</sup>National Bureau of Economic Research, Cambridge, MA 02138, USA. <sup>5</sup>Aix Marseille University, CNRS, Centrale Marseille, Aix-Marseille School of Economics, Marseille, France.

\*Corresponding author. Email: patrick.francois@ubc.ca

we conjecture, should help in selecting the firms most successful in obtaining cooperative efforts from their workers. There is already considerable evidence showing that a degree of norm-based acculturation occurs through workplace interactions; for example, (10), present evidence linking the social identity of employees to the performance of firms. To this, we added the conjecture that workplaces subject to more intense external competition would be more likely to engender prosocial norms of cooperation among their employees.

We used the generalized trust question, or a close variant of this, as our proxy for the prevalence of prosociality in the empirical analyses reported here. “Do you think that, on the whole, people can be trusted or that you can’t be too careful in dealing with people?” As has been documented, this question conjures a “weakly institutionalized” setting: “Answering this question, subjects consult either their own experiences and behaviors in the past or introspect how they would behave in situations involving a social risk” (11).

Survey-based questions of individual trust have been found to reflect variation in the degree to which subjects perceive the degree of prosociality of individuals around them. Laboratory-based validation studies of the generalized trust question suggest a few important features of this question that make it suitable for measuring prosociality. First, in the laboratory, generalized trust reported by individuals seems to be malleable and influenced by specific experiences. Second, beliefs about the trustworthiness of others seem to matter for informing subjects’ potentially costly trusting decisions and correlate with answers to the generalized trust question (12). Thoni (13) summarizes a large literature on this as showing that survey measures of trust are informative with regard to the cooperativeness of others, but there are sometimes discordant findings between play in incentivized games and answers to trust questions (9, 12) and studies showing that other factors are also at play in answers to the trust question (14–16). Third, individuals tend to respond to trustworthiness experiences by increasing their own trustworthiness, which is consistent with individuals being conditional cooperators, evidence of which is amply demonstrated by (17) in public goods experiments and summarized over a more comprehensive set of studies by (18). Conditional cooperators are then willing to follow perceived norms. These are not the only factors influencing responses to the generalized trust question, and we discuss other aspects known to affect responses as we present the main empirical results.

The evidence presented is drawn from four sources: (i) U.S. cross-sectional correlations between competitiveness of industry of employment and individual trust; (ii) U.S. state-level policy changes that altered cross-firm competition at the state level, inducing changes in individual trust; (iii) German panel data evidence showing changes in individuals’ industry of employment competitiveness induced changes in individual trust. All three forms of evidence confirm a strong and statistically significant effect of increased competition across firms on increased individual trust. We discuss precisely how CGS explains this observational data after we present it.

We augment these findings with (iv) evidence drawn from laboratory experiments conducted in France. We placed subjects into groups where group-level rewards are shared across members in a public goods game (PGG) setting. We manipulated the degree of competition across groups in a manner intended to mimic the variation in competition across firms that was observed in the data. We tested to see whether this variation replicated the correlations observed between competition and generalized trust in the data.

It does: Increases in competition across groups lead to increased generalized trust reported by individuals within the groups.

The pattern of subject behavior suggested a likely channel of effect. Cross-group competition increased the frequency of group beneficial behavior as it affected the economic returns to cooperation. Some subjects experienced enhanced group beneficial behavior and formed new groups in which they also exhibited increased group beneficial actions. These subjects responded by answering the generalized trust question more positively, perhaps because of extrapolating their experiences in the experimental setting beyond the laboratory. We returned to the precise interpretation of the empirical findings in light of CGS as they are presented.

## RESULTS

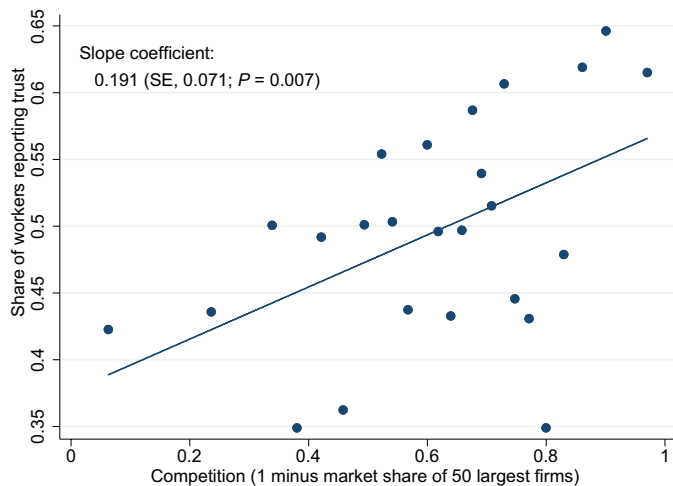
### Cross-sectional evidence in the United States

By its nature, cross-sectional data provide the weakest evidence that we consider here because a correlation between cross-firm competition and worker prosociality may reflect the effects of omitted variables that drive both. However, the labor force module asked of workers in the United States’ General Social Survey’s (GSS) 2004 wave has advantages in mitigating some of these concerns. This wave of the survey extensively focused on the workplace of survey respondents. This allowed us to control for many factors that may be affecting the generalized trust level of individual respondents, as well as rich personal information about respondents that allowed us to control for individual characteristics known to correlate with individual trust. As an example, we were able to include controls for the security of employment to ensure that positive answers to the generalized trust question are not just picking up tolerance for risk.

The percentage of sales covered by the  $k$  largest firms ( $k = 4, 8, 20, 50$ ) in an industry (as defined by the North American Industry Classification System) measured the competitiveness of each worker’s industry of employment. Our reported measure of competition is equal to one minus the sales covered by the 50 largest firms. In other words, the competitiveness of industry  $s$  is the percentage of total sales in  $s$  that is not covered by the largest 50 firms in that industry. Using competition measures based on the shares of the 4, 8, and 20 largest firms yields similar results. The use of these concentration ratios is common in economic analysis of industry structure and market power. Further information on the data is provided in the Supplementary Materials.

Figure 1 shows a binned scatter plot cutting the 612 GSS respondents in our sample into 25 equal-sized bins arranged by industrial competitiveness ( $x$  axis) plotted against share of workers reporting affirmative answers to the generalized trust question in that bin ( $y$  axis), after controlling for individual-level economic and demographic controls. The line is fitted from the unbinned data, so it perfectly matches regressions reported in the Supplementary Materials. The positive slope of 0.191 ( $P$  value of 0.007) is robust across many demanding specifications. In particular, as explained in the Supplementary Materials, including the rich and unusually comprehensive set of workplace controls obtained from the GSS workplace module does not alter this finding.

We provide a suggestion of causality by considering the effect of potential experience. The individuals likely to have had the longest exposure to the labor market were the ones for whom the effect of industrial competitiveness had the strongest association with trust (see the Supplementary Materials).



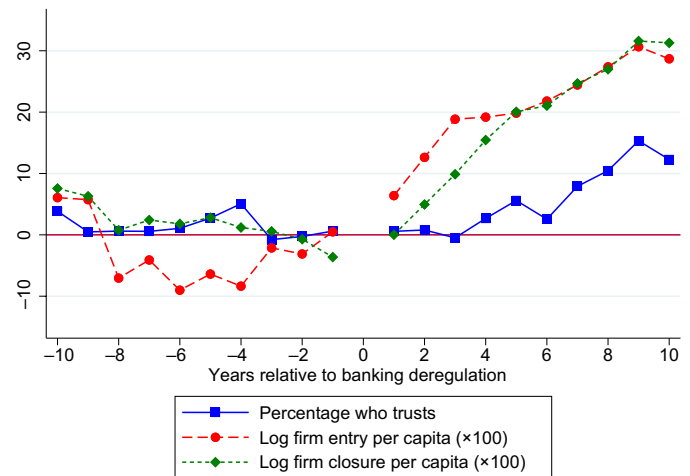
**Fig. 1. Americans who work in more competitive industries are more likely to trust.** This figure shows a binned scatter plot of worker's trust versus the competition in her industry of employment. The sample consists of 612 employed respondents of the 2004 GSS workplace module. The plot is constructed by dividing competition into 25 bins with an equal number of observations each and plotting mean trust indicator versus mean competition within each bin. The best fit line (and reported slope coefficient) is estimated by ordinary least squares (OLS) using the original (unbinned) data. Both the plot and linear fit partial out the determinants of trust control variables (see text for information).

Despite the inclusion of detailed workplace controls, as noted above, this correlative evidence is a long way from evidence of a causal relationship. The possibility of omitted factors potentially affecting both competition and individual trust cannot be discounted. A potential solution is to identify sources of variation that would alter competition between firms—without themselves having direct effects on trust levels. It turns out that this variation has been provided by episodes of U.S. banking deregulation, which we turn to next.

### Banking deregulation in U.S. states

Starting in the early to mid-1980s, multiple U.S. states lifted long-standing restrictions that prohibited banks from out of state to operate within their borders. Of particular interest for our research design, different states undertook deregulation at different times. Previous research (19, 20) indicates that these reforms can be seen as exogenous shocks to competition across all industries (including the nonfinancial) in a state. This is because banking deregulation increased credit availability, which, in turn, facilitated the creation of new firms and raised the contestability of local markets.

Figure 2 plots an event-study graph showing how deregulation affected trust levels, firm entry, and firm closures. It shows what the effect is of being 10, 9, ..., 2, 1 years before a reform, as well as 1, 2, ..., 10 years after. We normalize all variables to be equal to zero at the date of the year of reform ("year zero"). The red and green lines, rising steadily from each state-level deregulation event, indicate the (log of) firm entry and exit per capita. These are reproduced directly from (20). In addition, as shown by the upward trajectory that commences at the normalized year zero of banking reform (which varies in its calendar time for each state), competitiveness increased with the reforms. It continued to do so until 10 years after the reforms. This is consistent with the posited effects of increased credit availability due to banking reform on competition and is already well



**Fig. 2. Banking deregulation in U.S. states raised firm competition and trust.** This figure plots an event study graph for banking deregulation: the effect of each individual year before and after banking deregulation (normalized to year zero) on the probability of answering positively to the trust question. The sample consists of 17,455 respondents to the GSS in the 1973–1994 period. The estimating equations include state-fixed effects, year effects, individual state trends, and individual controls. The effects for firm entry and closures per capita are obtained from Kerr and Nanda (20) (see text for further information).

known. The pattern for years before the reform reassures that there are no preexisting trends in competition across states that are correlated with the timing of the reforms.

We now augment this finding about firm-level competition with information about individual generalized trust levels obtained from the GSS. As in the rest of this paper, we use a binary indicator of trust, but results are robust to alternatives, as discussed in the Supplementary Materials. Our dataset contains a total of 17,455 individuals in the 1973–1994 waves of the GSS. Leveraging that, we can observe state of residence of GSS respondents; the blue-shaded squares report how the propensity to affirmatively answer the generalized trust question is affected by banking deregulation. The blue-shaded squares again show no pretrend in state-level trust that predicts or preempts the banking deregulation. At time zero, trust is largely unmoved and remains so for the first 3 years. At year 4 after deregulation, state-level trust starts to track up, seemingly increasing hand in hand with the increase in new firm incorporations.

The Supplementary Material provides detailed information on how Fig. 2 is constructed and additional statistical tests. We note that the estimates in the figure control for a host of individual-level correlates of trust, state, and year fixed effects and state-specific linear trends, which control for the effect of state differences that are fixed or vary linearly through time, as well as common nationwide factors that may evolve nonlinearly, such as the business cycle. Estimates of the preferred specification elaborated there imply that a state enacting an interstate banking reform experienced a 1.4 percentage point increase in the share of its population reporting that they "can trust" every year after the reform.

Figure 2 supports a causal interpretation of the effect of banking deregulation increasing trust. It is also consistent with the increases in firm-level competition at the state level, leading to a rise in individual-level trust—precisely as would be posited by CGS.

A separate issue with these estimates is whether banking deregulation increases trust via increased firm competition or through another factor. For example, deregulation may have affected income growth or changed migration patterns. The Supplementary Materials provide further evidence that distinguishes among these possible channels and argues that the evidence is best explained by deregulation affecting trust via firm-level competition.

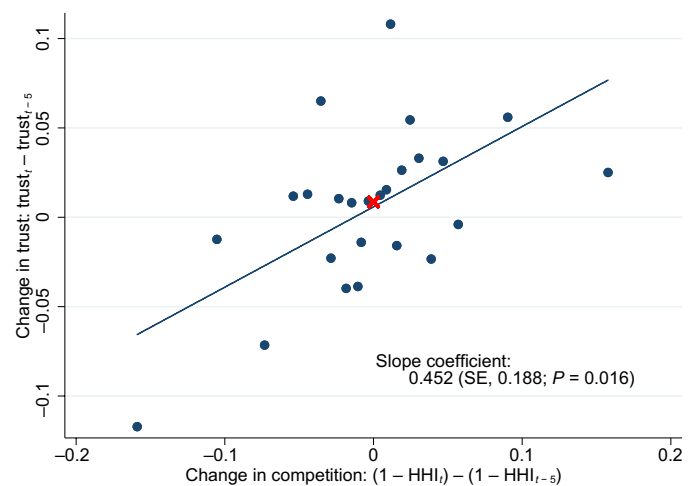
The results discussed in this section identify the effect of competition via an aggregate (state-level) shock to competition. An alternative and complementary strategy is to study whether individuals moving between industries with different levels of competition experience changes in trust. For this, tracking individual workers through time and observing changes in industry of work and trust are needed. To our knowledge, no U.S.-based survey that tracks individuals over a significant length of time has asked the trust question across multiple surveys while simultaneously reporting their industry of work. However, a dataset with these characteristics exists for German workers, and we turn to analyzing this now.

### Movers across industries in the German Socioeconomic Panel

We use the three waves of the German Socioeconomic Panel (SOEP) asking a trust question and including information on industry of employment: 2003, 2008, and 2013. The SOEP is representative of the German population; however, our sample only contains individuals who were employed in at least two consecutive waves of the survey. The sample contains 9103 observations from 6447 unique individuals employed across 50 different industries. Mean trust levels are higher than in the United States: 65% of respondents indicate a positive response on trust, meaning that they totally or slightly agree that “on the whole, one can trust people.”

The SOEP reports employed individual’s industry of work—which we match to a Herfindahl-Hirschman index (HHI) measure of competition obtained from the Orbis database (see the Supplementary Materials for further details). Our measure of competition is one minus the HHI of firms’ operating revenues. It is thus equal to zero in an industry with only one firm (monopoly), and would be equal to one in an industry with an infinite number of small firms. Our competition measure is mostly stable through time and certainly not time-variable enough to identify the effect of a change in competition across individuals who do not change industries. So instead, we explore the effect of changes in firm-level competitiveness by tracking individuals who changed industry (25.4% of the sample). Some individuals moved to jobs in more competitive industries, others stayed put, and others moved to industries with less competition.

Figure 3 is a visual summary of the results. Each blue-shaded circle in the graph is a binned average, constructed using only the respondents who moved across industries. We cut the  $x$ -axis variable (change in competition between two SOEP waves) into 25 bins of equal size, ordered from negative to positive changes. We plot this against the average change in trust between SOEP waves per bin. We estimate the regression line based on the original (unbinned) data. There is again a positive relationship between competitiveness of sector and individual trust. We highlight the distinction between the earlier Fig. 1 and the results represented here in Fig. 3. The former indicates that individuals in more competitive industries report higher trust. The latter indicates that workers who move from less competitive to more competitive industries are more likely to increase their reported trust levels. The red X in Fig. 3 denotes the average change



**Fig. 3. German workers who move to more competitive industries become more likely to trust.** This figure shows a binned scatter plot of worker’s change in trust versus change in the competition in her industry of employment. The sample consists of 9103 employed respondents of the German SOEP in the years 2003, 2008, and 2013. Changes are relative to trust and competition 5 years before. The blue-shaded circles are based on 2309 respondents who change industry of employment between survey years. They are calculated by dividing competition change in 25 bins with an equal number of observations each and plotting mean change in the trust indicator versus mean change in competition within each bin. The best fit line (and reported slope coefficient) is estimated by OLS using the original (unbinned) data. The red X is based on 6794 respondents who do not change industry of employment (and hence do not experience change in competition) (see text for further information).

in trust for those who did not change industries. Their change in trust is similar to those who moved across industries of comparable competitiveness (that is, zero).

The slope of 0.45 indicates that a 1-SD change in the competition measure increases the probability a worker responds affirmatively to the trust question by 1.7 percentage points (p.p.). Expressed alternatively, a worker who moves from a hypothetical industry, where three firms have 33% market share each, to one where four firms have 25% market share each, becomes 3.8 p.p. more likely to respond positively to the trust question.

In the Supplementary Materials, we present three pieces of evidence that further support our interpretation of Fig. 3 as a causal effect of competition on trust. First, we provide evidence of no pre-existing trends in the trust levels of movers; those who move to more competitive industries were not experiencing higher growth in trust before the move. Second, we show that our result cannot be explained by changing income; those moving to more competitive sectors do not experience higher income growth. Third, we show that differential trends in trust levels that are correlated with observable characteristics also cannot explain our results.

Overall, the results in Fig. 3 show that the individuals who changed jobs and ended up in more (less) competitive industries increased (lowered) their levels of trust. This strongly suggests a causal effect of sectoral-level competition on individual-level trust. The next section explains how CGS can account for all these findings.

### DISCUSSION

There is considerable experimental evidence, referenced earlier, supporting the conclusion that people are conditional cooperators: They



condition actions based on their beliefs regarding prevailing norms of behavior. They cooperate if they believe their partners are also likely to do so, and they are unlikely to act cooperatively if they believe that others will not.

The environment in which people interact shapes both the social and economic returns to following cooperative norms. For instance, many aspects of groups within the work environment will determine whether cooperation can be an equilibrium in behavior among group members or whether it is strictly dominated by more selfish actions. Competition across firms can play two distinct roles in affecting this. First, there is a static equilibrium effect, which arises from competition altering rewards from cooperative versus selfish behavior, even without changing the distribution of firms. Competition across firms punishes individual free-riding behavior and rewards cooperative behavior. In the absence of competitive threats, members of groups can readily shirk without serious payoff consequences for their firm. This is not so if a firm faces an existential threat. Less markedly, even if a firm is not close to the brink of survival, more intense market competition renders firm-level payoffs more responsive to the efforts of group members. With intense competition, the deleterious effects of shirking are magnified by large loss of market share, revenues, and, in turn, lower group-level payoffs. Without competition, attendant declines in quality or efficiency arising from poor performance have weaker, and perhaps nonexistent, payoff consequences. These effects on individuals are likely to be small in large firms where any specific worker's actions are unlikely to be pivotal. However, it is possible that employees overestimate the impact of their actions or instinctively respond to competition with more prosocial attitudes, even in large teams (21).

Competition across firms does not typically lead to a unique equilibrium in social norms but, if intense enough, can sustain a cooperative group norm. Depending on the setting, multiple different cooperative group equilibria differentiated by the level of costly effort can also be sustained. For example, if individuals are complementary in production, then an individual believing co-workers to all be shirkers and thus unable to produce a viable product will similarly also choose to exert low effort. An equilibrium where no one voluntarily contributes to cooperative tasks is sustained, and such a workplace looks to have noncooperative norms. In contrast, with the same complementary production process, and a workplace where all other workers are believed to be contributing high effort, a single worker will optimally choose to exert high effort as well to ensure viable output. In that case, a cooperative norm is sustained. When payoffs are continuous in both the quality of the product and the intensity of the competition, then the degree of cooperative effort that can be sustained can be continuously increasing in the intensity of market competition across firms. We have formalized this in an economic model that we include in the Supplementary Materials.

Competition's first effect is thus to make it possible, but not necessary, for group-level cooperative norms to arise as equilibria. The literature has shown that there are many other ways to stabilize cooperative norms as equilibria, such as institutional punishment, third-party punishment, or reputations. Cross-group competition may also enhance these other well-studied mechanisms for generating cooperative norm equilibria, but with or without these factors, it has a general effect of tilting the set of equilibria toward those featuring cooperative norms.

The second effect of market competition is a dynamic selection effect. This is the effect most usually emphasized in the literature on

CGS. Competition selects among the array of equilibrium norms displayed, those firms that converge on the best ones. Firms featuring cooperative norms should be able to outcompete those unable to sustain cooperation, and hence producing low-quality output. The more intense the competition, the greater the selective pressure, implying that the better firms expand more quickly and the weaker firms decline and shut down more rapidly. This selective effect has been argued to operate in a number of ways: via firm decline and exit, via migration from less successful to more successful firms that are expanding and hiring, and via mimicry (organizations selectively imitating and copying the behaviors or norms that prevail in their more successful competitors). Richerson and Boyd (22) precisely discuss how mimicry allows selection to still occur at the group level.

The final link in our explanation is how cooperative norms are related to affirmative answers to the generalized trust question. This question has been widely studied, and the extensive literature on it suggests that a number of factors can affect its answers. However, a factor that is consistently important is an individual's beliefs about the likely trustworthiness of anonymous others, and it is this component that would be moved by the forces of competition. Trustworthiness beliefs will increase under intense competition because of both the static equilibrium and the dynamic selection effects. In answering the trust question, and reflecting on their beliefs about the trustworthiness of others, subjects are informed by their life experiences, a major one of which is the performance of people around them in the workplace. Working in competitive sectors, subjects experience more cooperative behavior and accordingly respond more positively as to whether the imagined anonymous other can be trusted.

A limitation of the observational data is that it is impossible to interrogate this correlation between trust answers and competition further. In addition to beliefs about trustworthiness, answers to the generalized trust question have been shown to be affected by an individual's own preferences for behaving in a trustworthy manner (as they introspect about how they would act in a situation of social risk), their risk preferences, and their aversion to betrayal. Of course, if they are conditionally cooperative norm followers, then an introspective assessment of how they would behave in a situation of social risk will be related to previous experiences and beliefs. Nonetheless, with observational data alone, we cannot test which specific aspects of subjects' experiences affect the generalized trust answers, nor are we able to definitively establish a causal link. Despite the strength of the panel results, it remains at least a theoretical possibility that reverse causation or omitted variables are driving the observed correlation.

This is a marked advantage of the laboratory setting. There, we will both be able to more certainly assert a causal relationship and to explore the reasons for it in more detail.

## MATERIALS AND METHODS: LABORATORY EXPERIMENT

Our aim was to place subjects in settings where rewards were allocated on the basis of group-level outputs while altering the competitiveness that the groups experienced across treatments. By observing differential levels of competition exposure across individuals and designing the treatments so that equilibria varied under competitive and non-competitive arms, we could trace the effects on subjects directly: whether subjects in the competitive treatment were induced to increase their generalized trust. This would allow us to explore the consequences of the static equilibrium effect of competition directly. A limitation of the laboratory, however, is that the setting is artificial

(not an actual workplace) and short term (about 1 hour), so it was not possible for us to explore differential exit or success of groups based on their norms. One major channel of the dynamic selection effect of CGS was thus not present in our experiments. However, by allowing for multiple groups to be formed across rounds of the experiment, we were able to explore dynamic selection that would occur via selective imitation of successful groups.

We undertook our experiments starting in the fall of 2015 and ending in early 2016. Previous experimental work (23–27) has already shown that subjects placed together in groups and asked to contribute to a collective good—the canonical PGG—can have their contributions to the game substantially increased by putting them in group competitive settings. But do the effects of increasing competition also induce higher levels of trust? And if so, is this happening because of effects that could be attributable to CGS as we have argued for the observational data? We explored these questions in a pool of subjects from the Paris School of Economics.

Subjects played the PGG in two different treatments. The first (control) was the standard PGG. Twenty individuals were endowed with 10 euros per round and could decide how much they would contribute to a collective good that would benefit all members of their (two-member) group equally. By giving up  $x$  of her own private endowment, the amount of the collective pool (shared equally by both) would increase by 1.5 times  $x$ , thus benefiting the subject by only  $0.75x$  and therefore being a net cost to the subject. If a participant's objective is to maximize monetary reward, the dominant strategy is thus to contribute nothing in this game, and both individuals in each group doing so are the unique Nash equilibrium of the game.

Individuals were matched anonymously into groups, asked to make a contribution choice, and told the outcome and contribution of the other player they were paired with at the end of the round. In the next round, they were rematched into another group and played again. The rematching was with another anonymous individual, with whom they had not been previously matched, and the nonrepeated nature of the setting was made clear. This one-shot interaction was repeated 19 times per session, and subjects were rewarded on the basis of their payoffs computed in one randomly chosen round of the session (see the Supplementary Materials for details of the game and the full set of experimental instructions).

Before playing, subjects filled out a questionnaire regarding their particulars—education, occupation (if they had one, most were students), age, and gender. After playing, subjects were asked a number of questions drawn from the GSS—one of which was the generalized trust question.

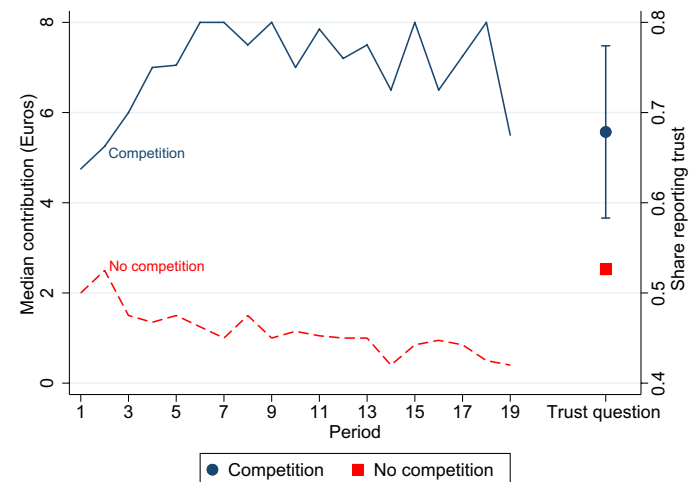
## RESULTS: LABORATORY EXPERIMENT

The red dashed line in Fig. 4 depicts the median contributions of players over the multiple rounds. As in almost every other experimental version of the PGG, the figure shows a declining pattern of contributions. Individuals started out contributing at a median level around 2 euros of their endowment—and this gradually tracked downward throughout the rounds, ending with a median well below 1 in round 19. This may be evidence of individuals learning the optimal strategy in the game, although other experiments focused on explaining these patterns lead one to doubt this interpretation (18, 28). This declining pattern is not our focus here, so we do not address it further.

We placed the remaining subjects in a “competitive” treatment. Here, the amount they received from the collective pool depended not only on the joint individual contribution and their partner's but also on the size of their joint contribution relative to that of a randomly allocated comparator group. If, and only if, their joint contribution equaled or exceeded that of their comparator group, did they receive their share of the collective account. We exactly computed the collective account as in the control group; total contributions were multiplied by 1.5 and shared equally by both members.

Contributions under the competitive treatment were less certain to create benefits, both for the group and for any individual contributing, since payoffs from contribution were now conditional on “winning” against the comparator group. All players contributing zero remains a Nash equilibrium of this game. But this competitive treatment also gave rise to equilibria with contribution levels that far exceeded the standard PGG of the control. Any positive level of contribution became a symmetric Nash equilibrium of this game. For example, if a subject expected all other players in the game to contribute the full amount, then contributing any less than that would lead to zero payment from the collective pool. However, by contributing the full amount of 10 euros, the pair's collective account would have 30 euros. If the other group did the same, then, since no group dominated, each subject in both groups was paid 15 euros, yielding this as another equilibrium. The same reasoning can be easily shown to support any other symmetric contributions as Nash equilibria of this game.

As in the control treatment, subjects were rematched anonymously into new groups after each round. The pair was also rematched (again anonymously) with a different randomly allocated comparator group, and the game was repeated for 19 rounds. We informed subjects about the contributions of their partner in the previous round and about the total contribution of their comparator group in the previous



**Fig. 4. Introducing competition in public good laboratory game increases contributions and propensity to trust.** The blue solid (red dashed) line plots the median contribution by participants in the competitive (noncompetitive) experimental sessions across the 19 rounds of the experiment. The red-shaded square plots the share of participants in the noncompetitive sessions who reported positive trust. The blue dot is constructed by adding the treatment effect of being in the competitive session from a regression with controls for age and gender. The whiskers show the 95% confidence interval. Sample consists of 220 experimental participants (100 in five competitive sessions and 120 in noncompetitive sessions) (see text for further information).

round too, before making their current round decision. We administered the same pre- and post questionnaires as in the control (standard PGG), so that we also measured generalized trust levels after participation.

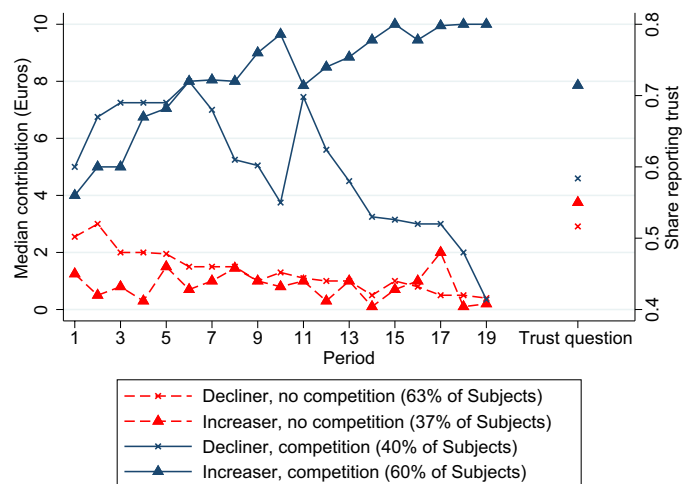
As the blue solid line in Fig. 4 shows, competition induced higher levels of contribution in the PGG across all rounds. The effect of competition is 0.384 (SE, 0.062;  $t = 6.16$ ;  $P < 0.001$ ;  $n = 180$ ). The specification controls for age and gender. There is also a markedly different experiment progression effect; median contributions jump in round 1 of the experiment to being more than twice as high in the competitive treatment than in the control and do not exhibit the customary decay.

The level being higher in round 1, before subjects have any experience of play, is consistent with what we have termed the static equilibrium effect, subjects inferring the possibility of Nash equilibria at higher levels of contribution in the competitive setting. Because of the complexity of computing equilibria (and that we provided no instructions on how to compute them), it is possible that many did not understand the equilibrium structure of the game. Hence, it is plausible that this first round difference is due to simply putting subjects into the competitive setting and is therefore not deliberative. It has been argued that group competitive settings can cue individual-level group cooperative set of responses as a type of priming effect (8).

But a competitive prime cannot explain the changes in play observed as the experiment progressed. Figure 4 shows that median contributions in the competitive treatment started below 5 euros, tracking up markedly over the first few periods; from there, they remained fluctuating around 7 euros. The pattern of decline exhibited in the standard (control) PGG did not appear.

Recall that this was not a repeated game played with the same group. Groups were created afresh across each round. Moreover, the median obscures considerable heterogeneity across individuals in how this progression happened. In particular, some subjects exhibited a positive trend in their contributions as the experiment progressed; we denote these “increasers,” and their median values for contributions across rounds are shown with blue-shaded triangles in Fig. 5. Others exhibited a declining trend (denoted “decliners”), denoted with blue X marks. The Supplementary Materials further discuss how these groups are defined. Decliners predominated in the standard PGG of the control (63% of subjects were decliners). In contrast, increasers were the largest group in the competitive sessions (45%, compared to 40% being decliners). These subjects started out similarly to the decliners in the competitive treatment but significantly increased their contributions across the competitive rounds, strikingly converging to a median of full contribution by the end.

To understand why subjects who started similarly can vary in their progression of contributions through the game, we explore the effects of the random matching of individuals (both as partners and competitors) as the experiment progresses. One of the channels of dynamic selection in theories of CGS is that groups selectively imitate the norms and/or practices of successful groups. This mimicry can lead to the diffusion of beneficial norms even in the absence of selection directly based on fitness (that is, via conflict and exit, which is not present here). Consistent with this, we conjecture that subjects might have been induced into becoming increasers when they experienced higher levels of partner (and competitor) contributions previously. The Supplementary Materials discuss how, by isolating variation arising from the random allocation of the ordering of partners, we can explore whether individuals who experienced higher levels of partner



**Fig. 5. Experimental subjects who increase their contributions throughout the competitive experimental sessions report higher trust.** The blue-shaded (red-shaded) connected triangles plot the median contributions by “increaser” participants in the competitive (noncompetitive) experimental sessions across the 19 rounds of the experiment. Increasers are defined as those who, on average, increase (or keep constant) their contributions as the experiment progress (see main text and the Supplementary Materials for further details). The blue (red) connected X marks plot the median contribution by “decliner” participants in the competitive (noncompetitive) experimental session across the 19 rounds of the experiment. Decliners are defined as those who, on average, decrease their contributions as the experiment progress (see main text and the Supplementary Materials for further details). The red X marks plot the share of decliners in the noncompetitive sessions who reported positive trust. The red-shaded triangle is constructed by adding the effect of being an increaser in a noncompetitive session, and the blue-shaded triangle (X mark) is constructed by adding the effect of being an increaser (decliner) in a competitive session (from a regression with controls for age and gender). Sample consists of 220 experimental participants (100 in five competitive sessions and 120 in noncompetitive sessions) (see text for further information).

(and competitor) contributions also reacted by increasing their own contributions in newly formed groups. This is the case. The average of lagged partner (and competitor) contributions positively predicts a subject’s own contribution in the next round.

The effect of previous partner contributions can occur for multiple reasons, such as misplaced reciprocity or “warm glow.” However, individuals increased their contributions when they experienced higher contributions from competitors as well. Such a pattern cannot be due to a competitive prime because it happens within the competitive treatments (not a comparison between competitive and noncompetitive treatments). It is also not evidence of reciprocal behavior (as groups were drawn afresh each round), nor a warm glow (as it increased the likelihood of losing the competition). It is, however, consistent with one dynamic selection channel of CGS: mimicry of the actions or norms in successful groups leading to diffusion of those norms into the broader population. Individuals form new groups and contribute more heavily when experiencing competing groups able to obtain higher contributions in the past. Moreover, this is not a general feature of the competitive treatment: Subjects experiencing competitors who contributed low amounts, on average, tended to lower their own contributions in newly formed groups subsequently, so it was not competition per se that did it.

We next checked whether the association between competition and trust found in the three survey datasets we analyzed also held in our experiment. It did. Subjects (51%) in the noncompetitive control

group answered the trust questions affirmatively (5 or higher on the provided 0 to 10 scale), and this is represented by the red-shaded square in Fig. 4. Subjects in the competitive session were 14.6 p.p. more likely to answer the trust question positively (blue-shaded circle in Fig. 4). The difference is statistically significant ( $P = 0.011$ ) and robust to multiple specifications and methods of inference, as discussed in the Supplementary Materials.

However, recall that there was variation among subjects in the competitive treatment that affected their own pattern of contributions across rounds, with the increasers seeming to be induced to higher contributions by being matched with relatively high contributors early on. It turns out that these same individuals were the ones who were also induced into affirmative answers to the generalized trust question. Formally, one can predict whether an individual would answer the generalized trust question affirmatively by knowing whether (in the random allocation of subjects across rounds) this individual was matched (either as a partner or as a competitor) with relatively high- or relatively low-contributing subjects. So, to reiterate, it is not only the case that individuals put into the competitive treatment and contributing more straightaway (that is, in round 1) reported higher levels of trust. It was also the case that individuals who experienced high levels of competitor and partner contributions through the random matching of the experiment both increased their own contributions in new groups that they formed in subsequent rounds and were significantly more likely to affirmatively answer the generalized trust question when the experiment ended. The instrumental variables strategy that underlies this conclusion is detailed in the Supplementary Materials. This is indicated by the breakdown of trust answers by differing types depicted at the right part of Fig. 5. The blue-shaded triangle, corresponding to the increasers in the competitive treatment, drives the difference in averages between competitive treatment and the PGG of the control.

If experiencing high-contributing subjects in the experiment raised trust levels via a warm glow from higher payoffs, then we would expect that this would only occur when matched with high-contributing partners, and not with high-contributing competitors (the latter lowers one's payoffs). However, if being matched with high contributors informed subjects about what successful groups tended to do and which cooperative norms were present within the subject pool, then it should not have mattered whether one experienced high contributions via partners or via competitors in previous rounds. Since we found an effect of experiencing higher contribution from both partners and competitors on trust, this suggests that increased trust levels induced by the competitive treatment were not driven by warm glow experience. Instead, the CGS-based explanation is that subjects who experienced relatively high contributions by both partners and competitors reacted to this in two ways. First, they increased their own contributions in subsequent rounds, although they would not be matched with the same partner again. This could be because they believed that future partners would also contribute more and (consistently) because they believed that success would come from these higher contributions. Second, it also changed their attitude toward the "anonymous other" as reflected in their response to the generalized trust question. They were more likely to think that others could be trusted, as they themselves were also induced into acting in ways that were more trustworthy (by contributing more). As found in previous laboratory studies of the generalized trust question, like (29), subjects seemed to extrapolate from the trustworthiness of their partners, and even their competitors in the experiment, to the wider context imagined by the generalized trust setting.

## CONCLUSION

Increased competition across firms exposes subjects to increased group beneficial behavior on the part of their co-workers and increases their own such behavior. In competitive markets, firms unable to elicit this cooperative behavior are likely to be outcompeted by firms that are more successful in doing so, leading to the proliferation of firms exhibiting cooperation. Workers in these settings experience, and themselves internalize, more cooperative norms. They then report more positive answers to the generalized trust question, which explains the cross-sectional and panel correlations we have reported here.

Competition across groups in an otherwise standard PGG conducted in the laboratory induces more group beneficial contributions from individuals within the groups. This happens immediately upon being put in this environment. In addition, for a subset of individuals who, by chance, are matched with more generous partners (and competitors), there is a progressively induced increase in their own group beneficial contributions. Subjects experiencing these cooperative contributions (either via their anonymous partners or competitors) are more likely to affirmatively answer the generalized trust question, which imagines a setting beyond the laboratory context.

Our competitively treated subjects in the laboratory do seem to have raised their beliefs regarding the possibility of a cooperative interaction, at least in the laboratory (relative to the noncompetitive treatment). Perhaps beliefs beyond that, as indicated by their responses to the generalized trust question, have also been similarly altered. However, we acknowledge that permanent effects flowing from this limited laboratory exposure seem implausible. Nonetheless, the laboratory does demonstrate that cross-group competition can alter actions and seemingly beliefs of subjects, in a manner that is consistent with CGS. If exposure to this competition is repeated, for example, as would occur in longer-term interactions arising from the workplace, then this evidence suggests that workplaces could be important conduits for these cooperative prosocial behaviors in general.

## SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/4/9/eaat2201/DC1>

Section S1. Cross-sectional evidence in the United States

Section S2. Banking deregulation in U.S. states

Section S3. German SOEP

Section S4. Laboratory experiment

Section S5. Experimental instructions

Section S6. A simple model where competition affects effort provision

Section S7. Experimental instructions for noncompetitive sessions (English translation)

Section S8. Experimental instructions for competitive sessions (English translation)

Table S1. Summary statistics: Cross-sectional data.

Table S2. Summary statistics for extended workplace variables.

Table S3. Sectoral concentration and trust (GSS workplace module data).

Table S4. Effects of sectoral concentration on trust—Interactions with experience (GSS workplace module data).

Table S5. Summary statistics—Banking regulation/GSS data.

Table S6. Effect of banking deregulation on trust (banking deregulation/GSS data).

Table S7. Effects of oil reserve value (banking deregulation/GSS data).

Table S8. Effect of changing industry of employment on changes in trust (German SOEP).

Table S9. Effect of competition in experimental sample.

Table S10. First stage: Effect of partners' and competitors' first period contributions (laboratory experiment).

Table S11. Learning from other players in the experiment.

Table S12. Effect of experimental experience on trust.

Fig. S1. Falsification test—Changes in competition uncorrelated with previous trends in trust.

Fig. S2. Falsification test—Changes in competition uncorrelated with changes in income.

Fig. S3. Distribution of answers to trust question in experimental sample.



Fig. S4. Distribution of answers to trust question in experimental sample—Cumulative distribution.

Fig. S5. Histogram of contributions in experiment by period.

Fig. S6. Distribution of trends in contributions.

Fig. S7. Experimental instructions (noncompetitive sessions).

Fig. S8. Experimental instructions (competitive sessions).

References (30–39)

## REFERENCES AND NOTES

- P. Richerson, R. Baldini, A. Bell, K. Demps, K. Frost, V. Hillis, S. Mathew, E. K. Newton, N. Naar, L. Newson, C. Ross, P. E. Smaldino, T. M. Waring, M. R. Zefferman, Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behav. Brain Sci.* **39**, e30 (2016).
- S. Bowles, H. Gintis, *A Cooperative Species: Human Reciprocity and Its Evolution* (Princeton Univ. Press, 2013).
- E. Hoffman, K. McCabe, V. Smith, Behavioral foundations of reciprocity: Experimental economics and evolutionary psychology. *Econ. Inq.* **36**, 335–352 (1998).
- K. Aoki, Altruism may be sexy: Comment on cultural group selection, coevolutionary processes and large-scale cooperation. *J. Econ. Behav. Org.* **53**, 37–40 (2004).
- J. Tooby, L. Cosmides, Conceptual foundations of evolutionary psychology, in *The Handbook of Evolutionary Psychology*, D. M. Buss Ed. (Wiley, 2005).
- J. Soltis, R. Boyd, P. J. Richerson, Can group-functional behaviors evolve by cultural group selection? An empirical test. *Curr. Anthropol.* **36**, 473–494 (1995).
- P. J. Richerson, R. Boyd, J. Henrich, Gene-culture coevolution in the age of genomics. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 8985–8992 (2010).
- J. Henrich, *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating Our Species, and Making Us Smart* (Princeton Univ. Press, 2016).
- E. Fehr, U. Fischbacher, The nature of human altruism. *Nature* **425**, 785–791 (2003).
- B. E. Ashforth, S. H. Harrison, K. G. Corley, Identification in organizations: An examination of four fundamental questions. *J. Manag. Stud.* **34**, 325–374 (2008).
- E. Fehr, On the economics and biology of trust. *J. Eur. Econ. Assoc.* **7**, 235–266 (2009).
- P. Sapienza, A. Toldra-Simats, L. Zingales, Understanding trust. *Econ. J.* **123**, 1313–1332 (2013).
- C. Thöni, Trust and cooperation: Survey evidence and behavioral experiments in *Social Dilemmas: New Perspectives on Trust*, P. A. M. Van Lange, B. Rockenbach, T. Yamagishi, Eds. (Oxford Univ. Press, 2017).
- N. Bardsley, R. Sausgruber, Conformity and reciprocity in public goods provision. *J. Econ. Psychol.* **26**, 664–681 (2005).
- A. Falk, Gift exchange in the field. *Econometrica* **75**, 1501–1511 (2007).
- I. Bohnet, R. Zeckhauser, Trust, risk and betrayal. *J. Econ. Behav. Organ.* **55**, 467–484 (2004).
- U. Fischbacher, S. Gächter, E. Fehr, Are people conditionally cooperative? Evidence from a public good experiments. *Econ. Lett.* **71**, 397–404 (2001).
- U. Fischbacher, S. Gächter, Social preferences, beliefs and the dynamics of free riding in public good experiments. *Am. Econ. Rev.* **100**, 541–546 (2010).
- S. E. Black, P. E. Strahan, Entrepreneurship and bank credit availability. *J. Finance* **57**, 2807–2833 (2002).
- W. R. Kerr, R. Nanda, Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship. *J. Financ. Econ.* **94**, 124–149 (2009).
- A. Gneezy, D. M. Fessler, Conflict, sticks and carrots: War increases prosocial punishments and rewards. *Proc. R. Soc. B* **279**, 219–223 (2012).
- P. Richerson, R. Boyd, The evolution of human ultrasociality, in *Indoctrinability, Ideology, and Warfare: Evolutionary Perspectives*, I. Eibl-Eibesfeldt, F. K. Salter, Eds. (Berghahn Books, 1998).
- M. N. Burton-Chellew, A. Ross-Gillespie, S. West, Cooperation in humans: Competition between groups and proximate emotions. *Evol. Hum. Behav.* **31**, 104–108 (2010).
- A. Gunnthorsdottir, A. Rapoport, Embedding social dilemmas in intergroup competition reduces free-riding. *Organ. Behav. Hum. Decis. Process.* **101**, 184–199 (2006).
- T. Markussen, E. Reuben, J.-R. Tyran, Competition, cooperation and collective choice. *Econ. J.* **124**, F163–F195 (2014).
- J. C. Cardenas, C. Mantilla, Between-group competition, intra-group cooperation and relative performance. *Front. Behav. Neurosci.* **9**, 33 (2015).
- E. Reuben, J.-R. Tyran, Everyone is a winner: Promoting cooperation through all-can-win intergroup competition. *Eur. J. Polit. Econ.* **26**, 25–35 (2010).
- J. Andreoni, Why free ride?: Strategies and learning in public-goods experiments. *J. Public Econ.* **37**, 291–304 (1988).
- P. Paxton, J. L. Glanville, Is trust rigid or malleable? A laboratory experiment. *Soc. Psychol. Q.* **7**, 194–204 (2015).
- J. Altonji, T. E. Elder, C. R. Taber, Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *J. Polit. Econ.* **113**, 151–184 (2005).
- D. S. Karlan, Using experimental economics to measure social capital and predict financial decisions. *Am. Econ. Rev.* **95**, 1688–1699 (2005).
- L. Schechter, Traditional trust measurement and the risk confound: An experiment in rural Paraguay. *J. Econ. Behav. Organ.* **62**, 272–292 (2007).
- D. J. Cohen, L. Prusak, *In Good Company: How Social Capital Makes Organizations Work* (Harvard Business School Press, 2001).
- E. Glaeser, D. I. Laibson, J. A. Scheinkman, C. L. Soutter, Measuring trust. *Q. J. Econ.* **115**, 811–846 (2000).
- A. Alesina, E. La Ferrara, Who trusts others? *J. Public Econ.* **85**, 207–234 (2000).
- S. Knack, P. Keefer, Does social capital have an economic payoff? A cross-country investigation. *Q. J. Econ.* **112**, 1251–1288 (1997).
- D. Acemoglu, A. Finkelstein, M. J. Notowidigdo, Income and health spending: Evidence from oil price shocks. *Rev. Econ. Stat.* **95**, 1079–1095 (2013).
- J. P. Haisken-DeNew, M. H. Hahn, PanelWhiz: Efficient data extraction of complex panel data sets: An example using the German SOEP. *J. Appl. Soc. Sci. Stud.* **130**, 643–654 (2010).
- R. Ibragimov, U. K. Müller, Inference with few heterogeneous clusters. *Rev. Econ. Stat.* **98**, 83–96 (2016).

**Acknowledgments:** We have benefited from the comments of seminar participants at a large number of conferences and departments. **Funding:** P.F. is grateful to the Canadian Institute for Advanced Research (CIFAR) and the Social Sciences and Humanities Research Council of Canada. T.F. is grateful for the financial support from the CIFAR. T.v.Y. is grateful for the financial support from the Institut Universitaire de France. **Author contributions:** All authors contributed equally to the paper. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplemental Materials. Additional data related to this paper may be requested from the authors. Some of the data used in the analysis are derived from Sensitive Data Files of the GSS, obtained under special contractual arrangements designed to protect the anonymity of respondents. These data are not available from the authors. Persons interested in obtaining GSS Sensitive Data Files should contact the GSS at [gss@norc.org](mailto:gss@norc.org). The SOEP data can be obtained through an agreement with the DIW (Deutsches Institut für Wirtschaftsforschung) and/or via the PanelWhiz application. The Orbis database (also known as Amadeus for European countries) can be purchased from the Bureau van Dijk Electronic Publishing (BvD).

Submitted 5 February 2018

Accepted 6 August 2018

Published 19 September 2018

10.1126/sciadv.aat2201

**Citation:** P. Francois, T. Fujiwara, T. van Ypersele, The origins of human prosociality: Cultural group selection in the workplace and the laboratory. *Sci. Adv.* **4**, eaat2201 (2018).

## The origins of human prosociality: Cultural group selection in the workplace and the laboratory

Patrick Francois, Thomas Fujiwara and Tanguy van Ypersele

*Sci Adv* 4 (9), eaat2201.  
DOI: 10.1126/sciadv.aat2201

ARTICLE TOOLS	<a href="http://advances.sciencemag.org/content/4/9/eaat2201">http://advances.sciencemag.org/content/4/9/eaat2201</a>
SUPPLEMENTARY MATERIALS	<a href="http://advances.sciencemag.org/content/suppl/2018/09/17/4.9.eaat2201.DC1">http://advances.sciencemag.org/content/suppl/2018/09/17/4.9.eaat2201.DC1</a>
REFERENCES	This article cites 33 articles, 1 of which you can access for free <a href="http://advances.sciencemag.org/content/4/9/eaat2201#BIBL">http://advances.sciencemag.org/content/4/9/eaat2201#BIBL</a>
PERMISSIONS	<a href="http://www.sciencemag.org/help/reprints-and-permissions">http://www.sciencemag.org/help/reprints-and-permissions</a>

Use of this article is subject to the [Terms of Service](#)

---

*Science Advances* (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2018 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).