OXFORD

The Effect of Trusting and Trustworthy Environments on the Provision of Public Goods

Sergio Lo Iacono[1,*] and Burak Sonmez[2]

[1]Department of Sociology, University of Utrecht, Padualaan 14, Utrecht 3584 CH, Netherlands and
[2]Department of Sociology, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

*Corresponding author. Email: s.loiacono@uu.nl

Abstract

Trusting and trustworthy environments are argued to promote collective action, as people learn to rely on their fellow citizens and believe that only few individuals will free ride. To test the causal validity of this mechanism, we propose an experimental design that allows us to create different trusting and trustworthy conditions simply by (i) manipulating the incentive structure of an iterated binary trust game and (ii) allowing information to flow among participants. Findings indicate that, given a similar distribution of resources among subjects, trusting and trustworthy environments strongly foster the provision of public goods. This outcome is largely driven by a learning effect: subjects transfer what they assimilate during a sequence of dyadic exchanges to their decision to act for the collectivity. In particular, results showed that what we learn from the community has a relevant effect on our ability to overcome the free-rider problem: we are more likely to act for the collectivity when we learn from the community to be trustful or reliable in our one-to-one interactions. The same applies in the opposite direction: we are more prone to free ride when we learn from the environment to be distrustful or unreliable in our dyadic exchanges.

Theoretical Background

Scholars have often suggested that trusting and trustworthy environments encourage collective action, facilitating the provision of public goods (Sampson, Raudenbush and Earls, 1997; Offe, 1999; Putnam, 2000; Rothstein, 2000; Ostrom, 2003)—especially when sanctioning is an unviable option, and communication is not possible. This claim relies on the idea that frequent trusting and trustworthy behaviours in a community promote the belief that only few fellow citizens will free ride. Hence, contributing to the public good will appear to yield high returns. This can also be described as a *learning effect*: our knowledge of positive past interactions with (or among) other citizens

(acquired through personal experience or observation) fosters our propensity to act as a group and overcome the free-rider problem.

Observational evidence supports this line of reasoning, showing a positive correlation between higher levels of social trust (i.e. trust towards our fellow citizens) and large-N collective actions, such as frequent recycling behaviours, more common neighbourhood watches, and broader tax compliance (Sampson *et al.*, 1997; Putnam, 2000; Hammar, Jagers and Nordblom, 2009; Sønderskov, 2009). However, these studies cannot properly deal with self-selection and reverse causality issues, or assess to what extent alternative factors are responsible for the relationship. For example, the correlation

might be driven by the fact that high-trusting communities tend to be wealthier (Knack and Keefer, 1997), potentially lowering the costs of collective action.

Experimental research, on the other hand, has often investigated the impact of the learning effect on cooperative actions. Several studies demonstrate that subjects' contributions in iterated public goods games (PGGs) tend to depend on the knowledge of other players' past actions (Andreoni, 1988; Kurzban *et al.*, 2001), which remains relevant even in the long-run (Mason, Suri and Watts, 2014), or across different generations of players (Schotter and Sopher, 2003; Chaudhuri, Graziano and Maitra, 2006). Similarly, research in sociology indicates that information regarding prior cooperative behaviours affects future exchanges of the same kind, while also pointing out the relevance of social embeddedness (Buskens and Weesie, 2000; Barrera and Buskens, 2007; Hofstra, Corten and Buskens, 2015; Iacono, 2018). For instance, Buskens and Raub (2002) and Buskens, Raub and van der Veer (2010) report that trusting and trustworthy deeds towards strangers in repeated interactions are significantly influenced by the learning effect at the dyadic and small networks level. In this sense, both individual (i.e. what we learn from our own experiences) and group (i.e. what we learn from the community) learning appear to have a crucial role in shaping cooperation.

Learning effects are also shown to spill over different games. Indeed, evidence from economics indicates that people learn across social exchanges, extrapolating what they assimilated in a specific dilemma to similar situations (Cooper and Kagel, 2003, 2008, 2009; Rick and Weber, 2010; Cason, Savikhin and Sheremeta, 2012; Mengel and Sciubba, 2014; Peysakhovich and Rand, 2016; Duffy and Fehr, 2018; Liu *et al.*, 2019). Cason *et al.* (2012) illustrate the existence of spillover effects in sequentially played coordination games, while Cooper and Kagel (2003, 2008, 2009) argue that meaningful context and team play have a significant impact on positive cross-game learning. More broadly, Grimm and Mengel (2012) show that individuals learn to play strategically equivalent games in the same way, and Liu *et al.* (2019) point out the relevance of the cognitive load of the interaction in producing spillovers.

Noticeably, the learning process described in this branch of the literature focusses on how people gradually understand what the optimal rational strategy is and converge to the equilibrium—i.e. strategy learning (Rick and Weber, 2010). However, as the structural conditions of the games differ, learning spillovers are harder to occur since transferring strategies might be inapplicable or even misleading. This is empirically shown, for example, in Mengel and Sciubba (2014) who 'found that playing a structurally different game hurts convergence to Nash equilibrium, while playing a structurally similar game leads to better (faster) convergence' (Mengel and Sciubba, 2014: p. 384)—see also Duffy and Fehr (2018). The notion that repeated trusting and trustworthy one-to-one interactions encourage the emergence of collective action (even in one-shot instances) is a challenging claim for the strategy learning approach, since it implies that subjects will transfer optimal strategies across social dilemmas that are structurally different between each other (e.g. sequential vs simultaneous interaction, different number of players, repeated vs one-shot, and so on).

Psychological research, on the other hand, emphasizes the relevance of other aspects of the learning process across games and contexts, such as the ability of individuals 'to obtain meaningful cognitive representations of higher-order concepts, rules, and relationships' (i.e. meaningful learning) (Rick and Weber, 2010: p. 716)—see also Peysakhovich and Rand (2016) for developments in this direction in the spillover literature.[1] Rusch and Luetge (2016), for instance, found that subjects use successful coordination with their partners as a cue of reliability to guide their behaviour in strategically different social interactions with the same partners. According to this line of thought, subjects should tend to generalize a wide spectrum of concepts across domains (not only the optimal strategy) on the basis of their individual experiences.

In a similar vein, sociological contributions in signalling theory indicate that people use information assimilated from one specific environment to cooperate (or not) in diverse situations where that information is relevant (Posner, 2000; Przepiorka and Berger, 2017). Fehrler and Przepiorka (2013) show that charitable giving works as a signal of trustworthiness in exchange games, leading people to rely on individuals who appeared to be generous. Also, Gambetta and Przepiorka (2014) illustrate that natural generosity is more effective in promoting trust than strategic generosity. On a similar note, Berger (2019) suggests that buying decisions in one domain (e.g. buying green products vs non-green products) can work as a signal to decide whether to trust or not (see also Gambetta and Székely, 2014; Przepiorka and Liebe, 2016; Przepiorka and Berger, 2017). In this sense, what we learn about other people's reputation (through direct experience or observation) in one context can operate as a sign or a signal of trustworthiness in other contexts, leading to the emergence of cooperation in dissimilar social exchanges.

Building upon these contributions, we address a crucial question for the trusting environment argument. That is, do people transfer what they learned about fellow citizens' behaviour in one-to-one interactions to their decision to act for the collectivity? We argue that the effect of trusting and trustworthy environments on the provision of public goods rests upon a learning process that concerns primarily what people assimilate about their fellow citizens' behaviour. More precisely, we claim that through dyadic interactions with strangers (whether directly experienced or observed) we learn not only what is the best response in a particular social exchange but also how people in our community tend to conduct themselves. Facing trusting dilemmas with strangers will contribute to create an 'image' of how strangers in our environment behave (Foddy, Platow and Yamagishi, 2009; Foddy and Yamagishi, 2009), allowing us to infer whether they are acting in a way that is trustworthy, selfish, distrustful, and so on. In this sense, each social exchange contributes to build the reputation of our unknown fellow citizens (Foddy *et al.*, 2009; Foddy and Yamagishi, 2009), and grasp what group dynamics (e.g. pro-social) are emerging in the community. If we establish that a relevant proportion of actors in our environment is being cooperative, then we will be likely to resist the temptation of defecting in one-to-one interactions, favouring instead mutually beneficial options. This should then affect collective interactions with the strangers in our community, even if one-shot (Putnam, 2000; Ostrom, 2003). Indeed, it is reasonable to assume that people will use what they assimilate regarding their fellow citizens' behaviour in dyadic interactions for other, structurally different, interactions where the same group of people (with whom a specific group dynamic has been established) is involved. That is, what we learn from other people's behaviour in one-to-one exchanges is likely to play a major role on our decision to act (or not) for the collectivity. For instance, if people learned to cooperate (or defect) in their dyadic exchanges with unknown fellow citizens (on the basis of what they observed in the environment, i.e. group learning, or their individual experiences, i.e. individual learning), then they will be more (or less) prone to contribute to the public good and overcome (or not) the free-rider problem (see Figure 1).

As a practical example, consider a neighbourhood where people often lend or borrow their possessions (e.g. cutlery, pegs to hang the washing, a ball to play, ingredients). Let us assume that such dyadic interactions tend to have a positive outcome, so that our neighbours lend us or return us what they borrowed in the vast majority of the cases. In addition, we see that the same

happens quite frequently among the other neighbours. When asked, later on, to work as volunteers to repair the fencing of the public park down the road, we might reasonably assume that our fellow citizens will offer their time and work to repair the fencing, and we will be likely to do the same. Indeed, even though we do not know most of our neighbours, from what we learned so far (based on direct interactions or observation), we can infer that they are generally inclined to cooperate (see Figure 1B). On the other hand, if in the neighbourhood lending or borrowing possessions is an uncommon practice because people tend to 'forget' or return the possessions in a terrible state, then we will probably make the opposite assumption (see Figure 1A). Thus, further cooperation will be unlikely to occur. Under this perspective, our claim is very straightforward: how strangers behave in their dyadic exchanges and what we learn from them will affect our willingness to cooperate with them even in other types of interactions, including collective action dilemmas.

However, how this learning effect is operating or if it is independent from other factors, such as the level of wealth and inequality in the community, is still empirically unclear. Observationally, high-trusting communities enjoy better conditions (e.g. lower material deprivation and crime rates, higher economic growth), which might be responsible for the higher propensity to collective action. Only in conditions of high inequality, wealthier individuals appear to defect more frequently, as the potential costs of collective action increase (Cherry, Kroll and Shogren, 2005; Côté, House and Willer, 2015; Nishi *et al.*, 2015). If trusting and trustworthy environments create more cooperation because of a learning effect, then even in conditions of low overall wealth but same inequality levels, people should learn from their fellow citizens and contribute to the collectivity.

The lack of experimental research on the relationship between trusting environments and the provision of public goods is clearly rooted in the practical and ethical issues related to the manipulation of trust. Paxton and Glanville (2015) only partially overcame this barrier by creating different trusting environments while relying on deception. In their experiment, subjects were told that they would play with other unseen participants using a computer interface. In fact, partners were simulated and their actions were set *ex-ante* to artificially increase the frequency of cooperative behaviours (Paxton and Glanville, 2015). Notice that Paxton and Glanville (2015) were interested in comparing the change in trusting attitudes in the two conditions, and their design did not allow to assess the effect of the different trusting environments on collective action.
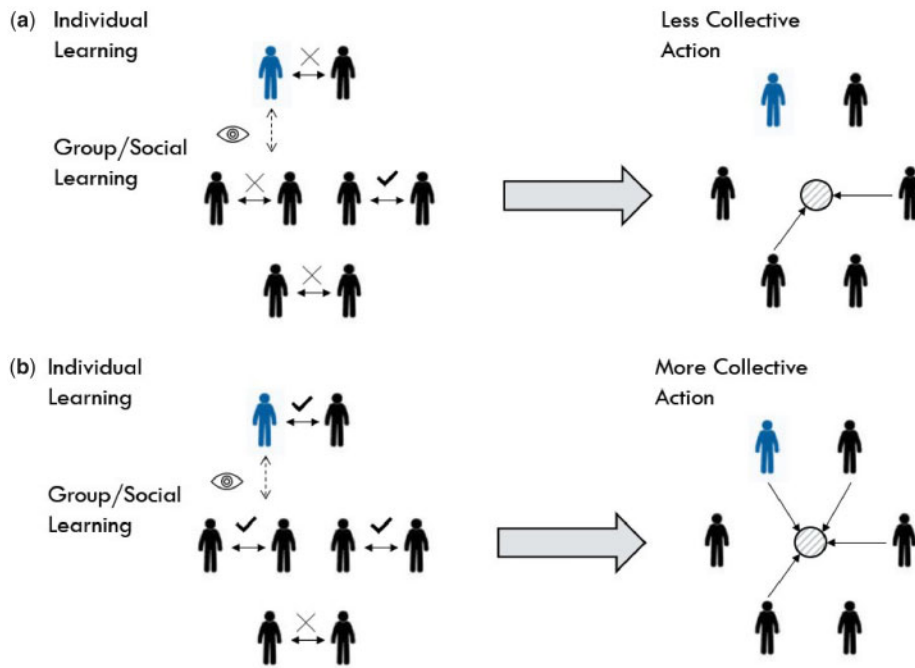
Figure 1. Theoretical mechanism—learning effect and collective action. (A) Distrusting and untrustworthy environment. (B) Trusting and trustworthy environment

Notes: Bi-directional arrows represent dyadic exchanges. Ticks (✓) indicate trusting and trustworthy interactions, while crosses (×) indicate either distrusting or untrustworthy behaviours. Uni-directional arrows indicate contributions to the public good, which is represented by the dashed circle. The figure assumes that the blue subject in condition A will learn to defect in his/her dyadic exchanges with strangers because of the predominance of distrusting and untrustworthy behaviours in the community (experienced both at the individual and group level). This knowledge on fellow citizens is then transferred to his/her decision to free ride when asked to act for the collectivity. The opposite should be true in condition B.

This article proposes a way to create trusting and trustworthy environments and evaluate their impact on the provision public goods without the use of deception. To achieve this goal, (i) we change incentives to be trusting and trustworthy by manipulating the payoffs of a binary trust game (hereafter TG) iterated for 20 rounds (Ermisch *et al.*, 2009) and (ii) we let participants be aware of the general level of cooperation in the community during the TGs by allowing full information flow (i.e. participants are shown the percentage of players who trusted or reciprocated at the end of each round). Combining these two elements, we produce self-reinforcing positive circles when incentives are maximized (in a high-trusting environment—T2), and self-reinforcing negative circles when incentives are minimized (in a low-trusting environment—T1)—in accordance with the literature on informational cascades and herd behaviour (Banerjee, 1992; Anderson and Holt, 1997; Çelen and Kariv, 2004; Goeree *et al.*, 2007). Next, we measure individuals' propensity to collective action, employing a one-shot binary PGG, where players

can invest the money earned during the TGs, and no further information on other participants is provided.

To determine if the PGG decision is due to a learning effect rather than variances in the level of economic disadvantage, we manipulate the exchange rate of experimental points (EP) into dollars at the end of the 20 TGs. As a result, people in the high- and low-trusting conditions have a similar distribution of resources at the moment of their PGG decision.

In line with the literature, our main hypothesis is that (H_1) *subjects in a high-trusting environment will be prone to overcome the free-rider problem and contribute more frequently in the PGG than people who experienced a low-trusting environment*. In addition, if people transfer what they assimilated from different environments to their decision to act for the collectivity, we should observe that (H_{2a}) *subjects who show a stronger group learning propensity during the TGs in the high-trusting environment will contribute more in the PGG*; (H_{2b}) *subjects who show a stronger group learning propensity during the TGs in the low-trusting environment*

will contribute less in the PGG; and (*H₂c*) *people who do not show a group learning effect during the TGs will behave similarly in the PGG across different treatments.* Finally, if our own experiences have also a role, then we should expect that (*H₃*) *subjects who learn to cooperate on the basis of their individual experiences during the TGs will be more likely to contribute in the PGG in comparison to subjects who learn to defect on the basis of their individual experiences.*

This article contributes to the literature in several ways. First, it offers a clear experimental test to a well-known theoretical argument in sociology and political science, namely that trusting and trustworthy environments promote collective action, while ruling out alternative explanations (e.g. community wealth and inequality). Second, it assesses the role of different learning process in affecting collective action propensity. In particular, this study distinguishes the impact of learning processes based on individual experiences (i.e. individual learning) from learning processes based on community experiences (i.e. group learning). In doing so, the paper expands the current research agenda in sociology assessing the role of learning in promoting cooperation (e.g. Buskens *et al.*, 2010), while taking into account both individual- and group-level processes (see Testori, Hoyle and Eisenbarth, 2019 for an example of this approach in psychology). Whereas the former has received a great deal of attention in signalling theory (e.g. Fehrler and Przepiorka, 2013) and spillover effect (e.g. Grimm and Mengel, 2012) literature, the latter has been largely overlooked in prior contributions. In this respect, this study complements existing research, emphasizing the relevance of ecological factors in understanding how collective action can emerge.

Material and Methods

Participants and General Setup

We ran a real-time interactive experiment on Amazon Mechanical Turk (AMT) involving 294 US participants (120 women, $M_{age} = 36.95$, SD = 11.55) across 58 sessions (more information on the sample in Supplementary Appendix A). Any given session included a planned number of six players. Subjects occasionally dropped out (most likely because of connection issues) or were not enough to fill the session. In such instances, we employed a standard strategy in web experiments to deal with missing players (see Supplementary Appendix A for further details). Participants earned on average $6.0 (including a showing up fee of $2.5, additional rewards for waiting time and survey completion). The experiment followed the Code of Ethics in Academic Research provided by the European University Institute Ethics Committee. Informed consent was obtained from all participants.

The experiment was programmed in oTree (Chen, Schonger and Wickens, 2016), which provides an integration with AMT. This allowed us to run our game in AMT directly through the oTree interface. Subjects voluntarily participated in the experiment by taking part in a session, which was randomly allocated to either T1 (low-trusting environment, $n = 160$) or T2 (high-trusting environment, $n = 134$). Each session followed the same order: a pre-experiment survey measuring individual trusting and altruistic attitudes; experimental instructions; a waiting page to let all required players to log-in; 2 trial rounds of the TG to allow players to get familiar with the interface and assess their understanding of the instructions; 20 rounds of a binary TG; 1 round of a binary PGG; and a follow-up survey gathering information on subjects' views about the game, as well as their background characteristics and demographics (see Figure 2). Experimental instructions were repeated extensively throughout the session.
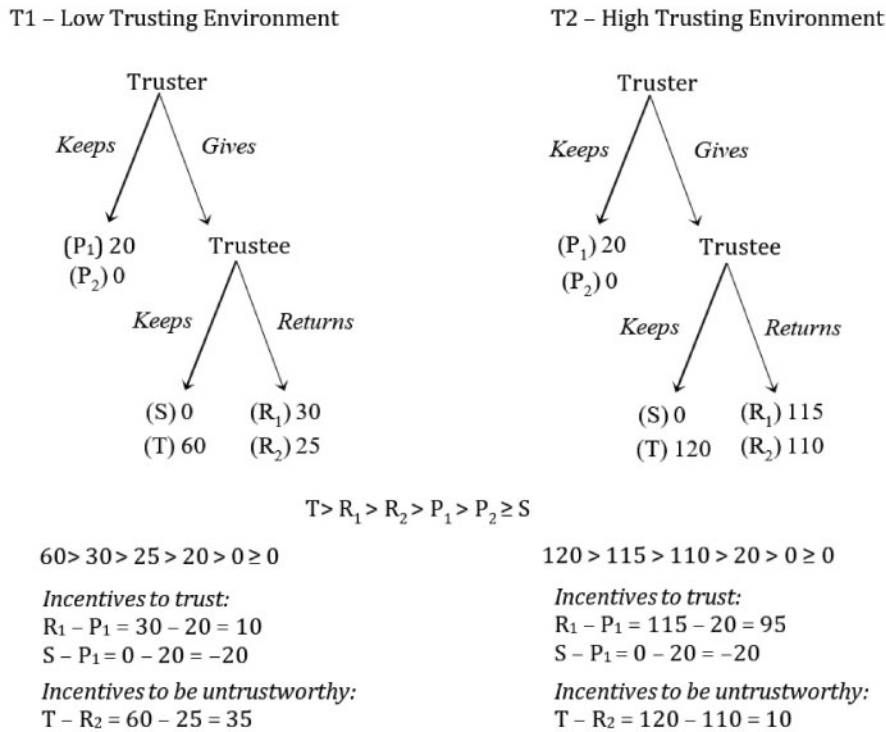
Treatments

Subjects are invited to a session to play 20 rounds of a binary TG and 1 round of a binary PGG. Each session is randomly assigned to one of the following treatments:

Treatment 1—low-trusting environment (T1). The payoffs in the binary TG are set to reduce the likelihood of trusting and trustworthy behaviours (see Figure 3).
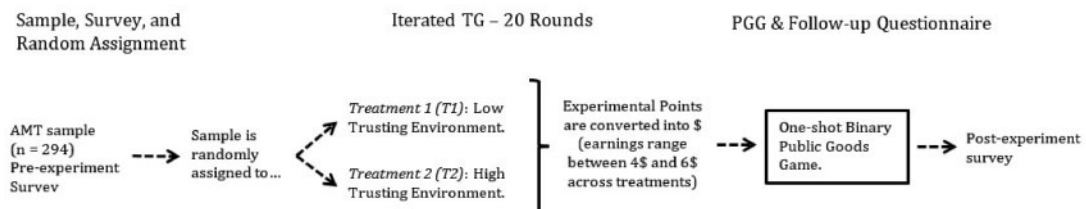
Figure 2. Overview of the experimental design

T1 – Low Trusting Environment

T2 – High Trusting Environment

$$T > R_1 > R_2 > P_1 > P_2 \geq S$$

60 > 30 > 25 > 20 > 0 ≥ 0   120 > 115 > 110 > 20 > 0 ≥ 0

Incentives to trust: *Incentives to trust:*
$R_1 - P_1 = 30 - 20 = 10$ $R_1 - P_1 = 115 - 20 = 95$
$S - P_1 = 0 - 20 = -20$ $S - P_1 = 0 - 20 = -20$

Incentives to be untrustworthy: *Incentives to be untrustworthy:*
$T - R_2 = 60 - 25 = 35$ $T - R_2 = 120 - 110 = 10$

Figure 3. Payoffs of high- and low-trusting environments. Left panel: T1—low-trusting environment and right panel: T2—high-trusting environment

Treatment 2—high-trusting environment (T2). The payoffs in the binary TG are set to increase the likelihood of trusting and trustworthy behaviours (see Figure 3).

Figure 3 shows that the incentives to trust are 9.5 times higher in T2 than T1 (95/10 = 9.5), while the incentives to be untrustworthy are 3.5 times lower in T2 than T1 (35/10 = 3.5). In addition, T1 has lower incentives to cooperate than the standard version of the game (Ermisch *et al.*, 2009). Notice that both treatments maintain the essential premises of a trust situation (Buskens *et al.*, 2010), as they respect the condition $T > R_1 > R_2 > P_1 > P_2 \geq S$.

Thus, trusting and trustworthy behaviours are encouraged in dyadic interactions in T2 and discouraged in T1. Furthermore, to give subjects a perception of the general trend in the community, we let participants know the percentage of players who trusted or reciprocated in the community at the end of each round (screenshots in Supplementary Appendix E). After the iterated TG, subjects were asked to play a standard one-shot binary PGG, which measures individuals' propensity to collective action. This design loosely represents the social dilemma faced by our hypothetical individuals in the neighbourhood example presented in the theoretical section. Indeed, in T2 people will be likely to engage in one-to-one interactions mimicking lending and returning exchanges, whereas in T1 this will be unlikely to happen. Furthermore, the decision to contribute (or not) to the public good reproduces, in a way, the decision to volunteer (or not) to repair the fencing of the public park. As mentioned earlier, our main expectation is that subjects used to trusting and trustworthy exchanges with strangers (e.g. lending and returning possessions with unknown neighbours) in their environment (e.g. neighbourhood) will be more prone to contribute to the collectivity (e.g. repair the fencing of the park down the road).

Binary TG (Stage 1)

All subjects played 20 rounds of a binary TG, which entailed two roles: a Truster and a Trustee. Within each round, the same steps were followed. First, the Truster was given an endowment of 20 EP and had the choice to send or keep the endowment. If the Truster sent her EP, the researcher multiplied the sum sent by a fixed amount. Then, the Trustee could decide whether to keep the sum received or return part of it to the Truster. After, players visualized their individual payoff, as well as a summary table displaying the percentage of people who sent or

returned EP in their session for the current round (see screenshots in Supplementary Appendix E). The action of the Truster implies a trusting behaviour, while the action of the Trustee implies a trustworthy behaviour.

In line with previous studies, roles were randomly assigned to subjects at the beginning of the session and kept fixed throughout the game. Subjects' identities were always anonymous. No information on opponent's individual history was provided. This was done to inhibit control effects. Similarly, subjects were matched with a different partner each round to avoid retaliation or direct reciprocity. Participants were informed that the partner would change each round, and had no knowledge of the identity of the opponent, or the number of players involved in the session. Therefore, though each session was composed by a planned number of six players, interactions were entirely anonymous and involved strangers. This was further reinforced by the web setup of the study, which allowed us to re-create actual interactions among unknown fellow citizens, increasing the ecological validity of the experiment. Finally, to minimize the 'end of the game' effect rounds' numbers were not displayed.

Public Goods Game (Stage 2)

After the iterated TGs, subjects were asked to play a one-shot binary PGG with incomplete information. In the PGG, subjects could either invest half of their endowment in a public pot, or keep the entire endowment. The money invested in the public pot was multiplied by 1.5 and then re-distributed equally among all members of the group regardless of their individual contribution. After the PGG decision, participants were informed of the result of the game, and asked to complete a short follow-up survey. The PGG involved all players in the session (i.e. six players). Participants were told that they were playing with all other subjects in their session, and were informed of the multiplying factor. However, they were not aware of the number of players involved in the session. This was done to better reproduce real-life collective action dilemmas, where the number of actors involved in the interaction is unclear. In this sense, the web setting of the experiment was central to simulate a more lifelike situation and improve the ecological validity of our study. Notice that, given the multiplying factor of 1.5 and group size greater than 1, the Nash equilibrium of the PGG is to contribute 0 to the public pot, in line with the standard PGG. Thus, independent of participants' understanding of group size, the rational choice was not to contribute. In addition, subjects were unaware of group size in *both* treatments; therefore,

differences in PGG decisions between the two conditions cannot be due to this factor.

The individual endowment for the PGG was determined by subjects' decisions during the TGs and it was expressed directly in dollars. This contributed to create a clearer separation between the two games from the participants' perspective: while the TG was played with EP in repeated dyadic interactions, the PGG entailed a single important group decision involving all the money the subject earned until that point. To avoid an asymmetrical distribution of resources across treatments, we applied different exchange rates of EP into dollars for the two conditions while defining an upper and lower bound.[2] As a result, PGG endowments varied between \$4 and \$6 within each session, reflecting an extremely low level of inequality among players (Gini coefficient$_{T1}$ = 0.06, SD = 0.02; Gini coefficient$_{T2}$ = 0.06, SD = 0.04). This was done to prevent different levels of inequality from affecting decisions in the PGG, a tendency clearly illustrated in previous research (Côté *et al.*, 2015; Nishi *et al.*, 2015).

Individual and Group Learning

To assess whether differences in the PGG decisions between T2 and T1 can be attributed to a learning effect, we construct two behavioural measures of learning based on Buskens and Raub's (2002) and Buskens *et al.*'s (2010) work: an Individual Learning Index and a Group Learning Index. The first one estimates how people learn from their own past experiences in the iterated TG. This measure is based on the consistency between subject's decision at round j and the opponent's decision at round $j - 1$. When subject cooperated at round j and his/her opponent cooperated at round $j - 1$, learning experience $l_{i,j}$ takes a value of 1 to indicate that subject is learning to cooperate. If subject defected at round j and his/her opponent defected at round $j - 1$, learning experience $l_{i,j}$ takes a value of -1 to indicate that subject is learning to defect. When subjects behaved in a way that is incoherent with their own past experience (i.e. defecting after opponent's cooperation or cooperating after opponent's defection), $l_{i,j}$ takes a value of 0 to indicate that no learning occurred:

$$\text{Individual Learning Index}_i = \sum_{j=1}^{x} \left(\frac{1}{2}\right)^{x-j} l_{i,j}$$

$$l_{i,j} = \begin{cases} -1 & \text{Learning Defection} \\ 0 & \text{No Learning} \\ 1 & \text{Learning Cooperation,} \end{cases} \quad (1)$$

where x is the total number of learning occasions (i.e. rounds where both players made a decision). The

measure is weighted in such a way that last rounds have more influence than early rounds, under the assumption that subjects remember recent experiences with more ease. The Individual Learning Index ranges between −2 and 2 (see Supplementary Appendix D for an alternative operationalization of individual learning).

Differently, the Group Learning Index identifies subjects who progressively adapted to the environment during the iterated TG, assimilating the dominant behaviour in the group. This is calculated as the weighted distance between the individual decision at round j and the session average decision at round $j-1$, so that last rounds have more weight than early rounds:

$$\text{Group Learning Index}_i = \sum_{j=1}^{y} \left(\frac{1}{2}\right)^{y-j} (1 - \mid \mu_{j-1} - i_{\text{decision}j}\mid),$$

(2)

where y is the total number of learning occasions (i.e. rounds where player made a decision), μ is the mean of the session decisions at round $j-1$, and $i_{\text{decision}j}$ is the decision of player i at round j (1 = trusting or trustworthy decisions; 0 = distrusting or untrustworthy decisions). The Group Learning Index takes higher values when the distance between subjects' decision and the group average in past round is small (e.g. subject cooperates when the majority of the group cooperated in the previous round), while it takes lower values when such distance is large (e.g. subject defects when the majority of the group cooperated in the previous round). The measure ranges between 0 and 2.

In line with our hypotheses, we expect the different payoff structures in T1 and T2 to generate opposite cooperative behaviours and group dynamics in the TG. These should be assimilated by subjects (either through individual- or group-level experiences) and transferred to the PGG where the same group of people (with whom such group dynamics have been established) is involved. Thus, we predict that people in T2 (high-trusting environment) will be more willing to contribute to public goods than people in T1 (low-trusting environment), and that this effect will be largely driven by what people learned during the iterated TG.

Results

Trusting Environments and Collective Action
We begin by showing the average contribution in the binary PGG by treatment in Figure 4. As it can be observed, people in the low-trusting condition contributed significantly less (M = 0.35; SE = 0.04) than people
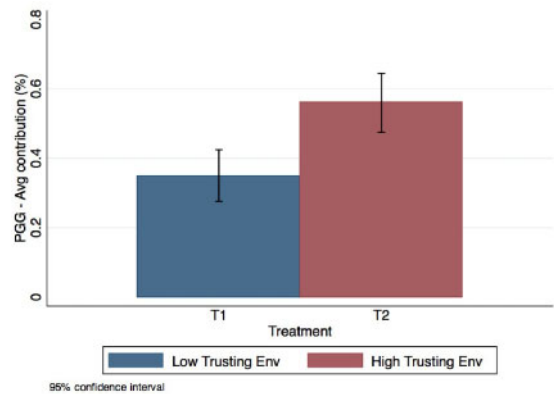
Figure 4. PGG contributions by treatment

in the high-trusting condition (M = 0.56; SE = 0.04), illustrating a substantial effect of trusting environments on collective action; $z = -3.60$, $P < 0.001$. Notice that, at the moment of the PGG decision, the distribution of resources was essentially the same and no significant difference in terms of inequality existed between the two conditions (Gini coefficient$_{T1}$ = 0.06, SE = 0.002; Gini coefficient$_{T2}$ = 0.06, SE = 0.003); $t(221.28) = -1.00$, $P = 0.319$. This is valid for both objective and subjective indicators of inequality (see Supplementary Appendix C for more details).

In addition, results show that people in T2 contributed more even though they were on average poorer than people in T1 at the moment of their PGG decision. As a matter of fact, the average of the PGG endowments in the low-trusting condition (M = \$5.07; SE = 0.05) was about \$0.30 higher than the ones of subjects in the high-trusting condition (M = \$4.76; SE = 0.08); $t(239.44) = 3.29$, $P = 0.001$. However, individual endowment had no impact on PGG contributions ($r_{\text{pb}} = -0.009$, $P = 0.875$; see also Table 1). In other words, given similar levels of economic disparities within the community, participants who experienced higher levels of trust and trustworthiness during the TGs showed a stronger propensity to act for the collectivity, even if they were slightly poorer on average. These results provide a first confirmation of our main hypothesis (H$_1$), suggesting that trusting and trustworthy environments promote collective action controlling for the level of wealth and inequality in the community.

To provide a more robust assessment of the treatment effect, we show results from a multilevel logistic regression estimating the probability of contributing in the PGG (see Supplementary Appendix C for robustness

Table 1. Multilevel logistic regression estimating PGG contributions

DV—PGG contribution	Model 1	Model 2	Model 3	Model 4
Treatment (ref: Low—T1)				
High trust Env—T2	0.938** (0.311)	0.915** (0.319)	0.989** (0.331)	0.942** (0.318)
Role (ref: Truster)				
Trustee	0.733** (0.258)	0.764** (0.266)	0.711** (0.271)	0.755** (0.280)
Individual endowment		−0.097 (0.183)	−0.111 (0.187)	−0.056 (0.192)
Understanding of the game (ref: Guessed correct n)				
Underestimated n			−0.339 (0.339)	−0.467 (0.418)
Overestimated n			0.408 (0.402)	0.337 (0.415)
Demographic covariates	No	No	No	Yes
N subjects/sessions	294/58	294/58	291/58	291/58
SD at the session level	0.641 (0.224)	0.676 (0.232)	0.699 (0.238)	0.564 (0.277)

Notes: Multilevel logistic regressions with standard error in parentheses. Demographic covariates include Age, Race (i.e. White, Black, Hispanic or Latino, Asian, Other), Gender, and Education. See complete results in Supplementary Appendix C. **$P < 0.01$.

checks). This model allows us to (i) control for factors within the design that might influence the relationship (e.g. role); (ii) adjust for baseline covariates potentially unevenly distributed among groups (e.g. gender); and (iii) properly account for sessions' clustering (i.e. individuals nested in sessions).

Results in Table 1 indicate that the treatment effect is robust to different specifications ($P < 0.01$ across all models), while pointing out the impact of other covariates. In particular, Model 1 indicates that the past role of players is a relevant factor in the PGG, as participants who played as Trustee during the TGs were more likely to contribute ($P < 0.01$) (see Supplementary Appendix D.2 for more details). On the other hand, players' endowments had no significant effect (Model 2), in line with prior studies indicating that individual wealth does not matter in PGG decisions when there is low inequality (Cherry *et al.*, 2005; Côté *et al.*, 2015; Nishi *et al.*, 2015). In our design, we did not inform participants of the number of players involved in the session, but we informed them that they were playing in a group comprising other players (thus, given the multiplying factor of 1.5 and a group size greater than 1, the rational choice for all subjects was to free ride); see Material and Methods section. This was done to better simulate real-life situations. In addition, as participants were unaware of group size in *both* conditions, the different results in the two treatments cannot be due to this factor. Nevertheless, since different understandings of group size between treatments might generate different PGG behaviours, we account for subjects' understanding of group size. Data show that participants believed that a similar number of players was involved in the session across treatments (Guessed group size$_{T1}$ = 8.66, SE = 0.53; Guessed group size$_{T2}$ = 7.76, SE = 0.55; $t(289)$ = 1.17, $P = 0.242$). We control for this aspect in Model 3, illustrating that overestimating or underestimating the number of players does not affect subjects' contribution or treatments' impact. Finally, in Model 4 we include standard demographics in the regression (see Supplementary Appendix C for more details). As it can be noticed, the introduction of these covariates slightly changes the coefficients while the significance of the treatments effect remains unaltered. Finally, when calculating the average marginal effect (AME) of the treatment on the dependent variable, it emerges that the probabilities to contribute in the PGG for subjects in the high trust condition are about 19.7 percentage points higher than those in the low trust condition, other covariates held constant (AME = 0.197, SE = 0.065; $z = 3.03$, $P < 0.01$—Model 4). This suggests that the magnitude of the effect is discreetly strong, further supporting the hypothesis that trusting and trustworthy environments can importantly foster collective action.

Learning Effect and Collective Action

To assess whether differences in the PGG decisions between T2 and T1 can be attributed to a learning effect, we now turn to the analysis of the two behavioural measures of learning (Buskens and Raub, 2002; Buskens
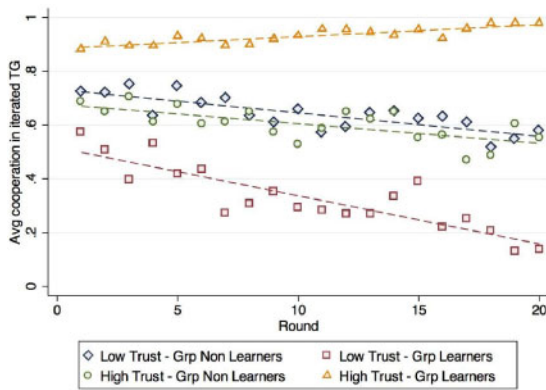
Figure 5. Average TG cooperation at each round by treatment and group learning propensity

Notes: Grp Learners identifies subjects who gradually complied with the environment, whereas Grp Non-Learners refers to subjects who did not. To define the two categories we employed as a cut-off point the mean value of the Group Learning Index (i.e. Grp Learners are above—or at— the cut-off point, while Grp Non-Learners are below the cut-off point). Cooperation in the iterated TG indicates trusting (i.e. give) or trustworthy (i.e. return) decisions. T1 Grp Non-Learners $n = 87$; T1 Grp Learners $n = 73$; T2 Grp Non-Learners $n = 42$; T2 Grp Learners $n = 92$. See Supplementary Appendix B for more details on the evolution of cooperation in the iterated TG.

et al., 2010), namely the group and individual learning indexes (see Material and Methods section).

Figure 5 illustrates how group learning propensities define different behavioural patterns in the high- and low-trusting conditions during the iterated TG. Though players were more cooperative in first rounds, their strategy gradually diverged as the game developed. In particular, it can be noticed that players who showed a stronger learning propensity in T1 adopted a more defective strategy: their average cooperation rate dropped from about 60 per cent in the first round to 30 per cent in the middle part of the game, ultimately plunging at less than 20 per cent in the last rounds ($P < 0.001$—see Supplementary Appendix B for more detailed results in this respect). Conversely, learners in T2 showed a consistent and steady increase in cooperation rates across all 20 rounds ($P < 0.001$—see Supplementary Appendix B), while remaining immune to the 'end of the game' effect. Finally, people who did not adapt to the environment (i.e. group non-learners) behaved similarly across treatments, displaying a slightly negative pattern.

To test whether having assimilated more cooperative or defective strategies during the TGs has an impact on the way people contribute to the public good, we first include the Individual Learning Index in the multilevel logistic regression predicting PGG contributions (Model

5). After, we estimate the impact of the Group Learning Index (Model 6), and interact it with the treatment variable to account for opposite learning tendencies across treatments (Model 7). Finally, we assess the effect of the two indexes simultaneously (Model 8). In these models, we include an additional covariate (i.e. Pure Conditional Cooperators) that identifies participants who complied with the average behaviour of the group in all TG rounds. This is done to control for those subjects who adopted a pure conditional cooperation strategy from the very beginning of the game, potentially biasing our measure of group learning.

Table 2 shows that the way participants adjusted to different trusting environments, and assimilated from their individual experiences were determinant in shaping collective action. In this sense, Model 5 shows that people who learned to cooperate in their one-to-one interactions with strangers on the basis of their own experiences are also more likely to cooperate in the PGG. This result clearly supports H₃. Furthermore, once we properly account for opposite group learning patterns across treatments by introducing the interaction term (Group Learning Index × Treatment) in Model 7, it emerges that PGG contributions depend strongly on whether people assimilated cooperative or defective behaviours from their community during the iterated TG ($P < 0.001$). This is graphically illustrated in Figure 6, which presents the predicted probability of contributing in the PGG by treatment and group learning propensity. As it can be seen, the more subjects adapt to a distrusting environment, the less likely they contribute to the public good. The opposite is true for people with a strong group learning propensity in T2: as participants learn to be cooperative with strangers in a high-trusting condition, they become far more likely to act for the collectivity. This holds true even when we account for individual learning in Model 8 (see Supplementary Appendix C for robustness checks). Indeed, although the Individual Learning index correlates moderately with the Group Learning Index ($r = 0.24$; $P < 0.001$), the interaction term remains positively correlated with PGG contribution ($P < 0.01$). This suggests that what people learn from their community has a significant effect on their propensity to act for the collectivity regardless of what they learn through their own personal experiences in the TG. Model 8 also illustrates that individual learning maintains a positive, though less significant, impact on collective action ($P < 0.05$) when controlling for group learning propensities. To assess the magnitude of these effects, we compute the AME for group and individual learning in Model 8. This shows that a one-unit increase of group

Table 2. Multilevel logistic regression estimating learning effect on PGG contribution

DV—PGG contribution	Model 5	Model 6	Model 7	Model 8
Treatment (ref: Low—T1)				
High trust Env—T2	−0.259 (0.391)	0.890* (0.347)	−4.323*** (1.101)	−3.468*** (1.105)
Role (ref: Truster)				
Trustee	0.739** (0.283)	0.744* (0.294)	0.778* (0.303)	0.736* (0.304)
Individual endowment	−0.194 (0.186)	0.068 (0.198)	0.023 (0.196)	−0.072 (0.196)
Understanding of the game (ref: Guessed correct n)				
Underestimated n	−0.366 (0.418)	−0.459 (0.418)	−0.520 (0.434)	−0.477 (0.432)
Overestimated n	0.416 (0.415)	0.335 (0.416)	0.246 (0.431)	0.299 (0.429)
Pure conditional cooperators	0.070 (0.381)	0.216 (0.435)	−0.605 (0.492)	−0.403 (0.483)
Group Learning Index		−0.073 (0.372)	−1.678*** (0.507)	−1.330* (0.521)
Group Learning Index × High trust Env—T2			3.859*** (0.788)	2.776** (0.878)
Individual Learning Index	0.680*** (0.148)			0.380* (0.173)
Demographic covariates	Yes	Yes	Yes	Yes
N subjects/sessions	291/58	291/58	291/58	291/58
SD at the session level	0.179 (0.674)	0.557 (0.277)	0.444 (0.324)	0.288 (0.457)

Notes: Multilevel logistic models with standard error in parentheses. Demographic covariates include Age, Race (i.e. White, Black, Hispanic or Latino, Asian, Other), Gender, and Education. See complete results in Supplementary Appendix C. *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.
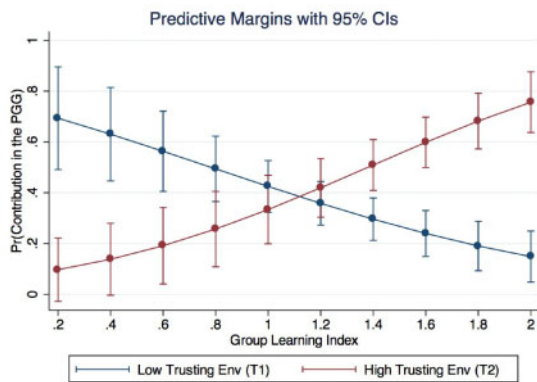
Figure 6. Predicted probability of PGG contributions by treatment and group learning propensity.

Note: The Group Learning Index ranges between 0 and 2, where 0 indicates a weak group learning propensity and 2 a strong group learning propensity.

learning in the high trust environment boosts an average change in the probability to contribute to the PGG of 27.2 percentage points (AME = 0.272, SE = 0.118; $z = 2.30$, $P < 0.05$). Conversely, a one-unit increase of group learning in the low trust environment lowers the probability to contribute by 23.6 percentage points (AME = −0.236, SE = 0.080; $z = −2.96$, $P < 0.01$)—see also Supplementary Appendix C.2. Finally, when individual learning increases by one-unit the chances of contributing increase by 6.8 percentage points (AME = 0.068, SE = 0.030; $z = 2.25$, $P < 0.05$). Such results indicate a quite strong impact of group learning

in comparison to individual learning in shaping PGG contributions, further advocating the idea that the two forms of learning do not necessarily overlap and can influence subjects' collective action propensities in different ways.

More broadly, these findings support hypotheses H_{2a} and H_{2b}, illustrating that people transfer what they assimilated from a sequence of one-to-one interactions in different surroundings to their decision to act for the collectivity. While in trusting and trustworthy environments we learn to rely on others and overcome the free-rider problem, distrusting and untrustworthy environments teach us to defect, reinforcing negative conducts and impeding the provision of public goods. On the other hand, people who did not learn from the environment showed opposite collective action propensities across treatments (see Figure 6), disconfirming H_{2c}. Non-learners are likely to be strong cooperators or defectors who will stick to their strategy independently of incentive structure or peer pressure.

Discussion

This article investigated if trusting and trustworthy environments promote collective action, assessing whether differences in the provision of public goods is due to a learning effect. To do so, we proposed a novel experimental design, which generates different trusting and trustworthy environments simply by (i) manipulating the incentive structure of an iterated binary TG and (ii) allowing information to flow among participants.

Results indicated that, given a similar distribution of resources among subjects, trusting and trustworthy environments foster collective action propensities, strongly increasing the probabilities to contribute to the public good in the high-trusting condition. This holds true even when we controlled for individuals' wealth and role, while adjusting for demographic covariates and data structure. People effectively learn from their community in both high- and low-trusting conditions, driving collective action propensities: the more we adapt to our trusting environment and assimilate its dominant strategy in our one-to-one interactions, the more we act for the collectivity. This works also in the opposite direction: we are more prone to behave as free riders, when we learn from our community to be distrusting and unreliable in our dyadic exchanges. Evidence pointed out that also learning from our own individual experiences matters in shaping our propensity to contribute to the public goods: subjects who learn to cooperate (defect) with strangers on the basis of their own one-to-one interactions are more (less) likely to act for the collectivity. The analysis showed that both individual and group learning are robust to several specifications, maintaining their significant effect even when included simultaneously in the equations. Overall, these findings support convincingly the trusting environment argument, indicating that trusting and trustworthy communities promote collective action through a learning effect (operating both at the individual and the group level), which leads individuals to transfer the knowledge they acquired on their fellow citizens during a sequence of one-to-one interactions to their propensity to contribute to the public good.

Contrary to our expectations, non-group learners showed very different collective action propensities, displaying either strong cooperative or defective strategies that remained unaltered. Unfortunately, our design could not assess what motivations, personal convictions, and strategies were behind their conduct. This point could be addressed in future research. In this sense, it would be particularly interesting to explore whether patterns of behaviour defined by what people learn are more stable and effective than the ones defined by individuals' preferences.

Our work offers an innovative but simple design to create different trusting and trustworthy environments without the use of deception, showing how such conditions can shape group cooperation. With this study, we aimed to address an important gap in the literature and empirically test a common theoretical argument on the provision of public goods. In particular, we highlight how community relations are determinant for a more comprehensive understanding of collective action, separating between individual- and group-based learning. By showing the relevance of ecological factors in shaping contributions to the public good, we expand the current research on learning effects and cooperation, while complementing prior findings in signalling theory and spillover effect, which focus mostly on individual-level processes.

Finally, interestingly enough, our results point out that what we learn from the community might be more relevant than what we learn from our own individual experiences in contributing to the public good. This could be due to how people realize what the general behaviour in the community is, which is essential to decide whether to cooperate in a dilemma involving the entire group. While our individual experiences might represent only a small fraction of all the possible interactions, (truthful) information on the general behaviour in the community provides us a richer understating of how the entire group is behaving. That is, as the sample of experiences on which individual learning is founded is smaller than the one at the basis of group learning, rational actors might give more weight to the latter when deciding whether to contribute to the public good or not. Clearly, further research is required to accurately investigate this aspect, and estimate the relevance of group-level processes in comparison to individual-level ones in promoting collective action.

Data Availability

The datasets generated during and/or analysed during the current study are available from the corresponding author upon reasonable request.

Notes

1   Peysakhovich and Rand (2016) argue that subjects, who internalize cooperative experiences in one setting, are more likely to punish selfishness and be more prosocial in other settings even if there is a stranger-matching protocol—in contrast to Rusch and Luetge (2016).

2   Subjects were not aware of the exchange rate in advance. However, after the completion of the iterated TG, they were shown how many EPs they earned in total, and the conversion in dollars (e.g. 2200 points = \$6).

Supplementary Data

Supplementary data are available at *ESR* online.

Contributions

Acknowledgements

References

Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, **87**, 847–862.

Andreoni, J. (1988). Why free ride? Strategies and learning in public goods experiments. *Journal of Public Economics*, **37**, 291–304.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, **107**, 797–817.

Barrera, D. and Buskens, V. (2007). Imitation and learning under uncertainty: a vignette experiment. *International Sociology*, **22**, 367–396.

Berger, J. (2019). Signaling can increase consumers' willingness to pay for green products. Theoretical model and experimental evidence. *Journal of Consumer Behaviour*, **18**, 233–246.

Buskens, V. and Raub, W. (2002). Embedded trust: control and learning. In Thye, S. R. and Lawler, E. (eds.), *Advances in Group Processes*. Bingley, UK: Emerald Group Publishing Limited, pp. 167–202.

Buskens, V., Raub, W. and van der Veer, J. (2010). Trust in triads: an experimental study. *Social Networks*, **32**, 301–312.

Buskens, V. and Weesie, J. (2000). An experiment on the effects of embeddedness in trust situations: buying a used car. *Rationality and Society*, **12**, 227–253.

Cason, T. N., Savikhin, A. C. and Sheremeta, R. M. (2012). Behavioral spillovers in coordination games. *European Economic Review*, **56**, 233–245.

Çelen, B. and Kariv, S. (2004). Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, **94**, 484–498.

Chaudhuri, A., Graziano, S. and Maitra, P. (2006). Social learning and norms in a public goods experiment with inter-generational advice. *The Review of Economic Studies*, **73**, 357–380.

Chen, D. L., Schonger, M. and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, **9**, 88–97.

Cherry, T. L., Kroll, S. and Shogren, J. F. (2005). The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab. *Journal of Economic Behavior & Organization*, **57**, 357–365.

Cooper, D. J. and Kagel, J. H. (2003). Lessons learned: generalizing learning across games. *American Economic Review*, **93**, 202–207.

Cooper, D. J. and Kagel, J. H. (2008). Learning and transfer in signaling games. *Economic Theory*, **34**, 415–439.

Cooper, D. J. and Kagel, J. H. (2009). The role of context and team play in cross-game learning. *Journal of the European Economic Association*, **7**, 1101–1139.

Côté, S., House, J. and Willer, R. (2015). High economic inequality leads higher-income individuals to be less generous. *Proceedings of the National Academy of Sciences*, **112**, 15838–15843.

Duffy, J. and Fehr, D. (2018). Equilibrium selection in similar repeated games: experimental evidence on the role of precedents. *Experimental Economics*, **21**, 573–600.

Ermisch, J. *et al.* (2009). Measuring people's trust. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **172**, 749–769.

Fehrler, S. and Przepiorka, W. (2013). Charitable giving as a signal of trustworthiness: disentangling the signaling benefits of altruistic acts. *Evolution and Human Behavior*, **34**, 139–145.

Foddy, M., Platow, M. J. and Yamagishi, T. (2009). Group-based trust in strangers: the role of stereotypes and expectations. *Psychological Science*, **20**, 419–422.

Foddy, M. and Yamagishi, T. (2009). Group-based trust. In Cook, S. E., Levi, M. and Hardin, R., (eds.), *Whom Can We Trust*. New York: Russell Sage Foundation, pp. 17–41.

Gambetta, D. and Przepiorka, W. (2014). Natural and strategic generosity as signals of trustworthiness. *PLoS One*, **9**, e97533.

Gambetta, D. and Székely, Á. (2014). Signs and (counter) signals of trustworthiness. *Journal of Economic Behavior & Organization*, **106**, 281–297.

Goeree, J. K. *et al.* (2007). Self-correcting information cascades. *The Review of Economic Studies*, **74**, 733–762.

Grimm, V. and Mengel, F. (2012). An experiment on learning in a multiple games environment. *Journal of Economic Theory*, **147**, 2220–2259.

Hammar, H., Jagers, S. C. and Nordblom, K. (2009). Perceived tax evasion and the importance of trust. *The Journal of Socio-Economics*, **38**, 238–245.

Hofstra, B., Corten, R. and Buskens, V. (2015). Learning in social networks: selecting profitable choices among alternatives of uncertain profitability in various networks. *Social Networks*, **43**, 100–112.

Iacono, S. L. (2018). Does community social embeddedness promote generalized trust? An experimental test of the spillover effect. *Social Science Research*, **73**, 126–145.

Knack, S. and Keefer, P. (1997). Does social capital have an economic payoff? A cross-country investigation. *The Quarterly Journal of Economics*, **112**, 1251–1288.

Kurzban, R. *et al.* (2001). Incremental commitment and reciprocity in a real-time public goods game. *Personality and Social Psychology Bulletin*, **27**, 1662–1673.

Liu, T. X. *et al.* (2019). Directional behavioral spillover and cognitive load effects in multiple repeated games. *Experimental Economics*, **22**, 705–734.

Mason, W., Suri, S. and Watts, D. J. (2014). Long-run learning in games of cooperation. In *Paper Presented at the Proceedings of the Fifteenth ACM Conference on Economics and Computation*. Palo Alto, California, USA

Mengel, F. and Sciubba, E. (2014). Extrapolation and structural similarity in games. *Economics Letters*, **125**, 381–385.

Nishi, A. *et al.* (2015). Inequality and visibility of wealth in experimental social networks. *Nature*, **526**, 426.

Offe, C. (1999). How can we trust our fellow citizens. *Democracy and Trust*, **52**, 42–87.

Ostrom, E. (2003). Toward a behavioral theory linking trust, reciprocity, and reputation. *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, **6**, 19–79.

Paxton, P. and Glanville, J. L. (2015). Is trust rigid or malleable? A laboratory experiment. *Social Psychology Quarterly*, **78**, 194–204.

Peysakhovich, A. and Rand, D. G. (2016). Habits of virtue: creating norms of cooperation and defection in the laboratory. *Management Science*, **62**, 631–647.

Posner, E. A. (2000). Law and social norms: the case of tax compliance. *Virginia Law Review*, 1781–1819.

Przepiorka, W. and Berger, J. (2017). Signalling theory evolving: signals and signs of trustworthiness in social exchange. In Ben, J. and Przepiorka, W. (eds.), *Social Dilemmas, Institutions and the Evolution of Cooperation*. Berlin, Germany: De Gruyter Oldenbourg, pp. 373–392.

Przepiorka, W. and Liebe, U. (2016). Generosity is a sign of trustworthiness—the punishment of selfishness is not. *Evolution and Human Behavior*, **37**, 255–262.

Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. New york: Simon and Schuster. [Database]

Rick, S. and Weber, R. A. (2010). Meaningful learning and transfer of learning in games played repeatedly without feedback. *Games and Economic Behavior*, **68**, 716–730.

Rothstein, B. (2000). Trust, social dilemmas and collective memories. *Journal of Theoretical Politics*, **12**, 477–501.

Rusch, H. and Luetge, C. (2016). Spillovers from coordination to cooperation: evidence for the interdependence hypothesis? *Evolutionary Behavioral Sciences*, **10**, 284.

Sampson, R. J., Raudenbush, S. W. and Earls, F. (1997). Neighborhoods and violent crime: a multilevel study of collective efficacy. *Science*, **277**, 918–924.

Schotter, A. and Sopher, B. (2003). Social learning and coordination conventions in intergenerational games: an experimental study. *Journal of Political Economy*, **111**, 498–529.

Sønderskov, K. M. (2009). Different goods, different effects: exploring the effects of generalized social trust in large-N collective action. *Public Choice*, **140**, 145–160.

Testori, M., Hoyle, R. B. and Eisenbarth, H. (2019). How group composition affects cooperation in fixed networks: can psychopathic traits influence group dynamics? *Royal Society Open Science*, **6**, 181329.

Sergio Lo Iacono is a post-doctoral researcher at the University of Utrecht. Current research interests include trust, social cohesion, collective action dilemmas, and social networks. His work has been published in *PLOS ONE*, *Social Science Research*, and *Social Inclusion*.

Burak Sonmez is a social scientist with expertise in social norms, trust, cooperation, immigration, public opinion, and causal inference. He mostly uses computational and experimental methods to investigate human behaviour in different social contexts.