Social Science Computer Review
http://ssc.sagepub.com/

The Norm-Signaling Effects of Group Punishment: Combining Agent-Based Simulation and Laboratory Experiments

Daniel Villatoro, Giulia Andrighetto, Jordi Brandts, Luis Gustavo Nardin, Jordi Sabater-Mir and Rosaria Conte

Social Science Computer Review published online 11 December 2013
DOI: 10.1177/0894439313511396

Article

Social Science Computer Review
201X, Vol XX(X) 1–20

DOI: 10.1177/0894439313511396
ssc.sagepub.com
SAGE

The Norm-Signaling Effects of Group Punishment: Combining Agent-Based Simulation and Laboratory Experiments

Daniel Villatoro[1], Giulia Andrighetto[2,3], Jordi Brandts[4], Luis Gustavo Nardin[2], Jordi Sabater-Mir[1], and Rosaria Conte[2]

Abstract

Punishment plays a crucial role in favoring and maintaining social order. Recent studies emphasize the effect of the norm-signaling function of punishment. However, very little attention has been paid so far to the potential of group punishment. We claim that when inflicted by an entire group, the recipient of punishment views it as expressing norms. The experiments performed in this work provide evidence that humans are motivated not only by material incentives that punishment imposes but also by normative information that it conveys. The same material incentive has a different effect on the individuals' future compliance depending on the way it is implemented, having a stronger effect when it also conveys normative information. We put forward the hypothesis that by inflicting equal material incentives, group punishment is more effective in enhancing compliance than uncoordinated punishment, because it takes advantage of the norm-signaling function of punishment. In support of our hypothesis, we present cross-methodological data, that is, data obtained through agent-based simulation and laboratory experiments with human subjects. The combination of these two methods allows us to provide an explanation for the proximate mechanisms generating the cooperative behavior observed in the laboratory experiment.

Keywords

norms, punishment, cooperation, agent-based simulation, laboratory experiments

Introduction

Cooperation and norm compliance is a puzzle for both the social scientists and the game theorists. Solutions to the puzzle of cooperation are generally found out in punishment-based social control

[1] Artificial Intelligence Research Institute, Spanish National Research Council, Barcelona, Spain
[2] Institute of Cognitive Sciences and Technologies, Rome, Italy
[3] European University Institute, Florence, Italy
[4] Institute for Economic Analysis, Spanish National Research Council, Barcelona, Spain

Corresponding Author:
Daniel Villatoro, Artificial Intelligence Research Institute, Spanish National Research Council—Campus UAB, Barcelona, Spain.
Email: dvillatoro@iiia.csic.es

(Axelrod, 1987; Fehr & Gächter, 2002; Ostrom, Walker, & Gardner, 1992; Sigmund, 2007). However, it is far from clear how punishment precisely works. With few exceptions, punishment has been considered as a manner of changing people's future conduct by increasing the cost of undesired behavior. In experimental works, punishment has been mainly implemented as the imposition of a material cost on the punished agent; hence, by reducing the expected payoffs of future violations, punishment increases the probability of positive social action.

This approach to punishment is in line with the rational choice theory of crime (Becker, 1968). Recently, this Beckerian view of punishment has been attacked and it has been shown that when punishment is modeled as the mere imposition of material costs, it is not always effective in promoting compliance and may even have detrimental effects (Dreber, Rand, Fudenberg, & Nowak, 2008; Egas & Riedl, 2008; Fehr & Rockenbach, 2003; Gneezy & Rustichini, 2000; Herrmann, Thoni, & Gächter, 2008).

Several scholars have instead emphasized that punishment promotes compliance not only by increasing the costs of undesired behavior but also by signaling norms (Andrighetto et al., 2013; Ellickson, 1991; Galbiati & Vertova, 2008; Heyman & Ariely, 2004; Masclet, Noussair, Trucker, & Villeval, 2003; Noussair & Tucker, 2005; Ostrom et al., 1992; Posner, 2000; Sunstein, 1996; Xiao & Houser, 2005, 2011). In previous works, we have used the term *sanction* to indicate when punishment exploits its norm-signaling function, while *punishment* refers to the practice in which there are changes only in material incentives but the norm-signaling function is not exploited (Andrighetto et al., 2013; Andrighetto & Castelfranchi, 2013; Giardini, Andrighetto, & Conte, 2010). Therefore, in addition to inflict a material damage, as *punishment* does, *sanction* conveys normative information aimed to inform the violator that he or she violated a norm and the violation is not condoned.

By norms, we mean prescribed behaviors shared and enforced by a community (Bicchieri, 2006; Conte, Andrighetto, & Campennì, in press; Conte & Castelfranchi, 1995; Crawford & Ostrom, 1995; Ostrom et al, 1992). As shown in previous works (Andrighetto et al., 2013; Galbiati & Vertova, 2008; Janssen, Holahan, Lee, & Ostrom, 2010; Masclet et al., 2003; Xiao & Houser, 2011), when material punishment is combined with communication of norms in order to form what we call *sanction*, compliance is higher than when they are used separately. Norms inform individuals about how they are prescribed to behave and punishment makes the expected consequences of violating them more certain, thus making norms *salient* in individuals' mind (Andrighetto et al., 2013; Bicchieri, 2006; Cialdini, Kallgren, & Reno, 1990; Xiao & Houser, 2011). As proposed by Andrighetto et al. (2013), we consider that *sanction* activates both the cognitive machinery for instrumental decision making and the norm-psychology humans are provided with (Chudek & Henrich, 2011; Conte & Castelfranchi, 1995; Sripada & Stich, 2006). We claim that successful human cooperation results from the interaction of these two cognitive devices.

Different factors have been proposed to favor and even amplify the norm-signaling effect of punishment, for example, communication, contextual effects, and so on (Andrighetto et al., 2013; Galbiati & Vertova, 2008; Janssen et al., 2010; Masclet et al., 2003; Xiao & Houser, 2011). In this work, we extend these findings, by focusing on the norm-signaling power of *group punishment*. We claim that when inflicted by the entire group, the recipient of punishment (and possibly the public) views it as a norm signaling. In particular, we put forward the hypothesis that with the same material incentive inflicted, group punishment is more effective in enhancing compliance than uncoordinated one, because it is more likely to be interpreted as a sanction. The reason is because the higher the number of punishers, the less likely the observers will interpret the punishing behaviors as dictated by self-interest and, conversely, the more likely they will attribute punishment to impersonal, normative, and possibly, legitimate reasons, that is, as upholding a norm.

In order to test our hypothesis, we designed a laboratory experiment using a public goods game with punishment (Fehr & Gächter, 2002). The experiment consists of four treatments, in which the number of punishers varies. We compare the cooperation levels obtained when the violator is punished by 0, 1, 2, or 3 punishers in a group of four subjects. In the 1, 2, and 3 punishers treatments, the material punishment inflicted to the punished subject is the same and it reduces the punished subjects' payoff to zero. We find that when punishment is inflicted by 1 or 2 punishers, cooperation levels are significantly higher than in the 0 punishers treatment. More interestingly, in the 3 punishers treatment, that is, when the entire group punishes the violator, cooperation levels are significantly higher than in the 1 or 2 punishers treatments. These results provide first evidence that since it amplifies the norm-signaling effect, group punishment performs better than the uncoordinated one in boosting the desired conduct. However, an explanation of the causal factors that generate those data is still missing.

In order to fill this gap and to shed light on the internal mechanisms that allow sanction to boost cooperation more than material punishment, we use the *normative* architecture EMIL Internalizer Agent (EMIL-I-A) (Conte et al., in press; Villatoro, Andrighetto, Conte, & Sabater-Mir, 2011). We replicate the laboratory experiment through agent-based simulation in which we compare agent architectures endowed with different decision and learning algorithms. The agent architectures contrasted are a *reinforcement learning*, which bases its functioning on a cost and benefit analysis, and the normative architecture EMIL-I-A, which is sensitive both to the normative information conveyed by sanction and to the material damage received when punished. The EMIL-I-A architecture allows us to properly explore the power of "moral suasion" of group versus uncoordinated punishment.

The need for cross-methodological research is increasingly felt among the social scientists for the sake of both empirical validation and modeling (Bravo, Squazzoni, & Boero, 2012; Duffy, 2006; Poteete, Janssen, & Ostrom, 2010). While laboratory data show us the impact of manipulated independent variables on human behaviors, agent-based modeling helps us investigate the internal mechanisms that generate such behaviors. The combination of these methods allows us to provide an explanation for the proximate mechanisms behind the observed cooperative behavior.

This article is structured as follows: In the next section, the related work on the norm-signaling function of punishment is presented. Then, in order to test the effectiveness of group punishment in humans, a laboratory experiment is presented followed by its replication through an agent-based simulation. Finally, the conclusions and future work are presented.

Related Work

Bowles and Polania-Reyes (2012) provide an exhaustive overview of the recent work on the norm expressive theory of punishment. This theory has been largely supported by legal theorists who claim that a well-designed punishment mechanism should clearly express disappointment for norm violation and provide instructions about the appropriate conduct (Ellickson, 1991; Posner, 2000; Sunstein, 1996). Different factors have been proposed to favor and even amplify the norm-signaling effect of punishment, for example, communication, contextual effects, and so on (Andrighetto et al., 2013; Ellickson, 1991; Galbiati & Vertova, 2008; Heyman & Ariely, 2004; Masclet et al., 2003; Noussair & Tucker, 2005; Ostrom et al., 1992; Posner, 2000; Sunstein, 1996; Xiao & Houser, 2005, 2011).

Noussair and Tucker (2005) report data from a public goods game experiment in which incentives have been combined with verbal messages of disapproval. They show that when both incentives and verbal messages of disapproval are available, contributions are higher than when only one of the two systems is available. The authors incidentally suggest that peer pressure appears to be the factor that leads those who receive incentives combined with verbal messages of disapproval to contribute more. Similarly, Galbiati and Vertova (2008) and Andrighetto et al. (2013) show that when obligations and incentives are used together, cooperation is strongly reinforced: The combined effect of incentives and obligations is significantly more positive on contributions than the impact of either incentives or obligations alone. Recently, Janssen, Holahan, Lee, and Ostrom (2010) present an

Table 1. Group Punishment Experiment's Payoff Matrix.

	0*C*	1*C*	2*C*	3*C*
C	5	10	15	20
D	10	15	20	25

Note. C stands for contribute and *D* for defect.

experimental environment in which costly punishment can be combined with communication to study social–ecological systems. In their setup, participants can communicate extensively and decide whether or not to adopt a punishment system and how much the fines should be. They also allow for the different temporal orders of the availability of punishment, communication, and both, and their results depend to some extent on the order of the treatments. Among other results, they find that communication with punishment does not lead to as long-lasting cooperative behavior as communication without punishment, a result at odds with (Andrighetto et al., 2013; Galbiati & Vertova, 2008; Noussair & Tucker, 2005). In addition to these experimental results, ethnographic evidence indicates that in human ecologies punishment rarely takes place as the mere imposition of a material cost, but naturally combines normative information and material punishment (Boehm, 1999; Ellickson, 1991; Guala, 2012; Ostrom, 1990). As suggested by Boehm (1999), verbal communication is highly effective in transmitting information about the prescribed conduct and expressing condemnation for norm violation. The role of verbal communication in promoting compliant conduct has been extensively proven inside and outside the laboratory, but normative information is not always transmitted through verbal messages. Xiao and Houser (2011) provide experimental evidence that contextual cues also amplify the norm-signaling effect of punishment, showing that when punishment is publicly, but anonymously, implemented it promotes cooperation more effectively than when the same incentive is privately implemented. Xiao (2013) also provides experimental evidence indicating that if people know that the enforcer can benefit monetarily by punishing, they no longer view punishment as signaling of a norm violation, and consequently the punitive act is less effective in promoting compliance than when it is perceived as a disinterested act. However, Faillo, Grieco, and Zarri (2013) show that when punishment is meted out by punishers who are perceived as legitimate, it yields substantial benefits to cooperation and welfare compared to discretionary punishment. Finally, Boyd, Gintis, and Bowles (2010) emphasize the role of coordinated punishment in conveying a message of peer condemnation. In this work, we extend these findings by testing the norm-signaling power of group punishment.

Group Punishment: A Laboratory Experiment

In order to test the effect of group punishment in promoting cooperation, we conducted a laboratory experiment reproducing a standard public goods game (see Fehr & Gächter, 2002). In a public goods game, a group of subjects receives a certain initial endowment and decides whether to contribute or not to a public fund that returns a payoff to each of the subjects. As shown in Table 1, the structure of the payoffs in the public goods game makes it a classical social dilemma, as payoffs are such that it is individually rational to abstain from investing in the group fund, thus the group best strategy would be investing in the group fund because this yields a bonus. The more people invest in the group fund, the larger their share of the bonus. Therefore, based on the material self-interest hypothesis, the prediction is that everybody should free ride. Punishment has been shown to be an effective way of overcoming the type of dilemma posed by the public goods game (Axelrod, 1987; Fehr & Gächter, 2002; Ostrom et al., 1992; Sigmund, 2007; Yamagishi, 1986). Our experiment consists of performing this game allowing subjects to use punishment. More specifically, it consists of four treatments, which differ with respect to the number of punishing subjects. The experiment is built as follows: In the four treatments, 20 groups of 4 subjects interact over 40 rounds. The first nine rounds (1–9) are identical across treatments. At every round, each subject *i* independently chooses whether or not to invest its private endowment in a group fund. The resulting payoff for each subject is calculated according to Table 1.

After having contributed or not to the group fund, all the members of the group are informed about their individual earnings and the decisions of each participant of the group and they have the possibility to punish those who did not contribute (Fehr & Gächter, 2002). Punishment consists in reducing the payoff of the punished subject to zero and it also inflicts a cost to the punisher.

From Round 10 to 40, the four treatments differ with respect to the number of punishing subjects: 0 punisher (Treatment 1), 1 punisher (Treatment 2); 2 punishers (Treatment 3); and 3 punishers (Treatment 4).

What is important is that the material damage inflicted on the punished agent in Treatments 2, 3, and 4 is *identical* and the way the experiment has been implemented prevents the occurrence of reputational effects, as participants cannot identify one another. Thus, both the material and the symbolic incentives imposed in Treatments 2, 3, and 4 are the same. As claimed in the Introduction section, we suggest that if subjects react differently in Treatments 2, 3, and 4, this seems to be due to the growing probability that the punishment is interpreted by the target as a sanction, that is, as an act of conveying a normative message of peer condemnation.

What is special to our setup is that each group of four subjects was composed of one human subject and three confederate preprogrammed virtual players. The reason for putting each human subject in a group with three virtual preprogrammed players is to be able to observe humans in a completely controlled situation. Below we specify the confederate virtual players' behaviors. Human subjects knew that the experiment would last 40 rounds, but they were not informed of the fact that they were playing with confederate virtual players (at the end of the experiment, participants were asked to reply to a questionnaire and only 3 of the 80 subjects replied that they might be playing with virtual preprogrammed players; the remaining participants replied saying they did not realize it).

Experimental Design

The game consists of 40 rounds. Each round is structured in three stages:

1. *First stage:* Participants decide simultaneously and without communication whether or not to contribute to the public good.
2. *Second stage:* Participants are informed about their individual earnings and about the decisions and earnings of the other participants in their group. After having been informed, participants are given the opportunity to punish each other simultaneously.
3. *Third stage:* After the participants make their punishment decisions, they are informed about the punishment decisions and resulting payoff of the other participants in their group.

In order to observe whether group punishment, that is, punishment inflicted by the entire group, has a different effect on cooperation than individual punishment, that is, when punishment is uncoordinated, four treatments have been designed, which differ in the number of punisher confederate virtual players:

- *0 Punishers treatment:* Each human subject is assigned to a group with *three nonpunisher* confederate virtual players.
- *1 Punisher treatment:* Each human subject is assigned to a group with *one punisher* confederate virtual player and *two nonpunisher* confederate virtual players.

Table 2. Punisher Confederate Agents Behaviors.

	First Stage	Second Stage
Round < 10	$\Pr(C) = 0.5$	$\Pr(P, A_x) = 0.25; x \in \{C, D\}$
Round $\geq$ 10	$\Pr(C) = 0.9$	$\Pr(P, A_D) = 0.9$

Table 3. Nonpunisher Confederate Agents Behaviors.

	First Stage	Second Stage
Round < 10	$\Pr(C) = 0.5$	$\Pr(P, A_x) = 0.25; x \in \{C, D\}$
Round $\geq$ 10	$\Pr(C) = 0.9$	$\Pr(P, A_x) = 0.0; x \in \{C, D\}$

- *2 Punishers treatment:* Each human subject is assigned to a group with *two punisher* confederate virtual players and *one nonpunisher* confederate virtual player.
- *3 Punishers treatment:* Each human subject is assigned to a group with *three punisher* confederate virtual players, that is, the whole group.

As said, the material damage imposed on the punished agent in the treatments with 1, 2, or 3 punishers is identical (i.e., it reduces the punished subjects' payoff to zero). Punisher and nonpunisher confederate virtual players are programmed to contribute and punish in a particular way (see Tables 2 and 3). During the first nine rounds, all confederate virtual players' behavior is programmed to be exactly the same. They have 50% probability to contribute to the public good and 25% probability to punish. Such behavior is based on previous experimental data (Fehr & Gächter, 2002), in which was observed that in the first round, roughly 50% of humans cooperate, and they punish 25% other participants in their group, independently on how the latter acted at the first stage (even though we are aware of the possible consequences of this assumption, the confederate virtual players mimic for the first nine rounds the behavioral dynamics observed in humans in similar conditions).

Starting from the 10th round, punishers and nonpunishers behave differently: (a) all confederate virtual players have 90% probability to contribute; (b) nonpunishers never punish; and (c) punishers have 90% probability to punish defectors.

Our choice of a punishing probability of 90% for each of the confederate punisher virtual players is based on several considerations. First, we wanted to have strong punishers, but we did not want to choose a probability of 100%, since we believe that such consistent punishing behavior would have revealed to the humans that they were playing with artificial agents. The large body of experimental research with humans suggests that humans are very rarely completely consistent in their choices. Second, we maintain the treatment differences in the expected number of punishers; consequently, there are small differences in the probability of being punished. Thus, as the number of punishing virtual players increases in the group, the probability of being punished (as each punisher confederate player has 90% probability to punish) increases as well. In other words, the defector's probabilities of being punished in each of the treatments are 90%, 99%, and 99.9%, respectively, for the treatments with 1, 2, and 3 punishers. We think that given the general dynamism of the experimental environment, subjects were not able to perceive the small differences in the probability of being punished.

We therefore suggest that such different punishment probabilities cannot explain the difference in cooperation observed in the three treatments with punishment (see Figures 1 and 2) and in the next main section we provide data that support our claim. As claimed in the Introduction section, we

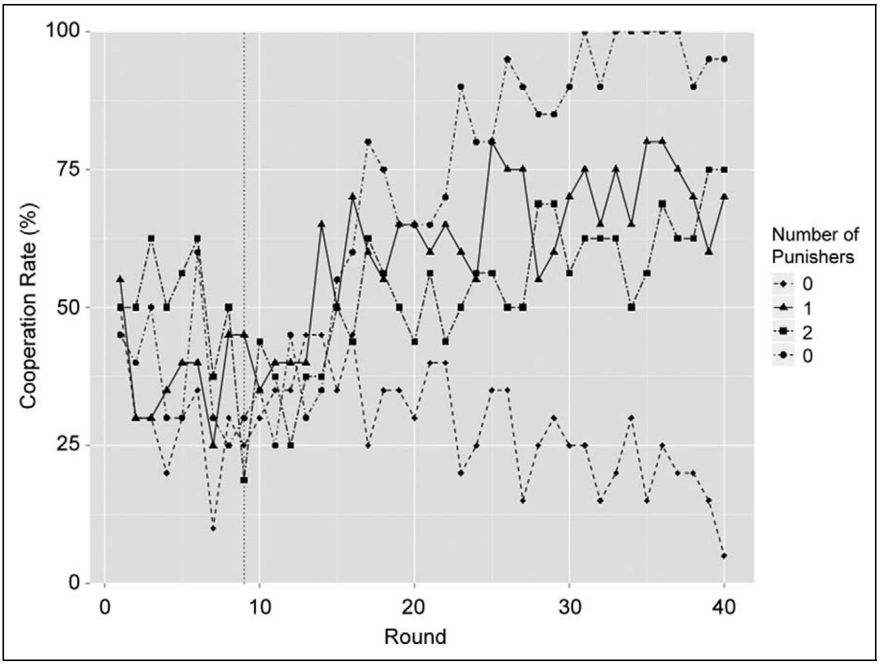

Figure 1. Average human contribution rate.

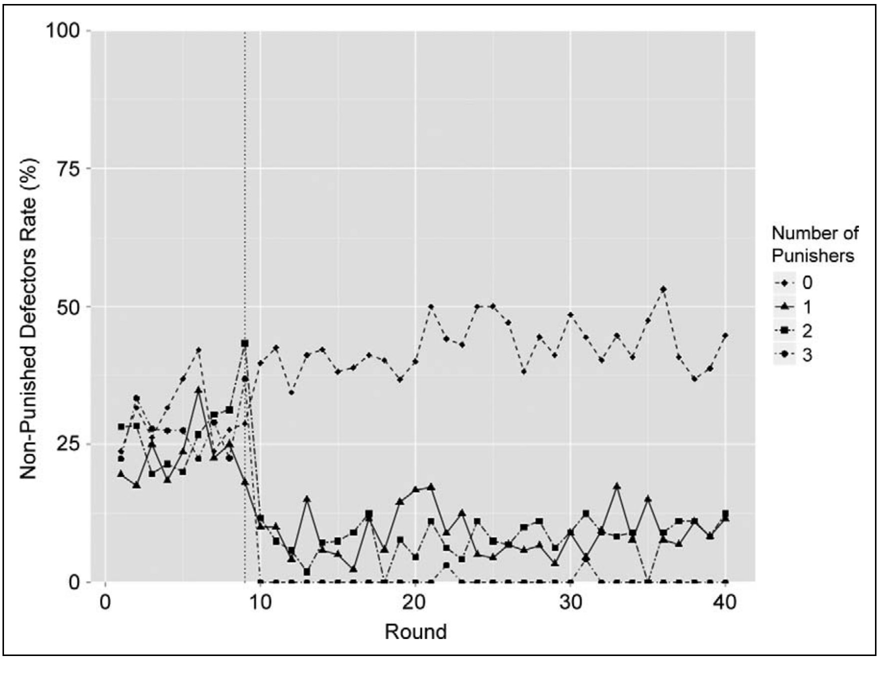

Figure 2. Nonpunished human defectors rate.

suggest that if human subjects react differently in treatments with 1, 2, and 3 punishers, this can be attributed to the different ways in which humans interpret the received punishments. In particular, we postulate that group punishment is interpreted as a sanction, that is, as an act of conveying a normative message of peer condemnation and this has an impact on the norm psychology humans are provided with.

Procedures

A total of 80 human subjects voluntarily participated in four separate experimental sessions conducted in the Experimental Economics Lab of the Universitat Autònoma de Barcelona in December 2010. However, due to platform failure in two treatments the data from two subjects were lost; only 76 samples were considered for analysis. As these are independent observations, such failure does not influence the results. Considered that the experiment consisted of four treatments and each treatment was played by 20 groups of the 4 subjects, a total of 320 between real (80) and virtual players (240) participated in the experiment.

The experiment was programmed using the HIHEREI (Brito, Pinyol, Villatoro, & Sabater-Mir, 2009) platform, which allows the design and implementation of laboratory experiments for human subjects with the possibility of including confederate virtual players. The human participants were students from a wide range of fields of study and interacted anonymously. Subjects were not allowed to participate in more than one experimental session. Prior to the first round of each session, human participants were told that they had been randomly arranged in groups of four subjects and that groups would be the same for the entire experiment (the partner protocol used in our experiment is similar to Denant-Boemont, Masclet, & Noussair, 2007). The total number of subjects in the room was always a multiple of four in order to avoid participants from guessing they were playing with virtual players. Participants could not communicate with each other during a session. At the end of the game, human participants were paid depending on their performance.

Results

In Figure 1, the human average contribution rates obtained in the four treatments are shown. In the four treatments from Round 1 to 9, human subjects interact with confederate players who have a 50% probability to contribute to the public good and 25% probability to punish other participants in their group, independently on how they acted at the first stage. During these first nine rounds, the contribution level decreases, but it does not collapse completely as a consequence of the 25% probability of being punished.

From Round 10, it is possible to observe the effect of the four different experimental treatments on the contribution level. In the 0 punishers treatment, the contribution level collapses. Considering that the incentive scheme is structured in such a way that defecting is the dominant strategy (see Table 1), this is not surprising (see also Fehr & Gächter, 2002). On the other hand, in all the three punishing treatments, the contribution level increases with respect to the first nine rounds and is higher than the one obtained in the 0 punishers treatment (Mann–Whitney test [$\alpha = 5\%$] from Round 10 to 40 with $p = 8.633 \times 10^{-11}$ for 1 punisher treatment; $p = 1.481 \times 10^{-9}$ for 2 punishers treatment; and $p = 1.422 \times 10^{-9}$ for 3 punishers treatment).

The interesting result is that even though the material damage inflicted on the punished subject is identical, the fact of being punished by the entire group (i.e., 3 punishers treatment) leads to a higher contribution level than being punished by just one or two group members. Using the Mann–Whitney test ($\alpha = 5\%$), we show that from Round 10 to 30 the contribution level among the three treatments is statistically different ($p = .00175$ for 1 and 2 punishers treatments; $p = .04722$ for 1 and 3 punishers treatment; and $p = .00142$ for 2 and 3 punishers treatments). From Round 31 to 40, however, only the contribution level obtained in treatments with 1 and 3, and 2 and 3 punishers continues to be statistically different ($p = .00034$ for 1 and 3 punishers treatment; $p = .00032$ for 2 and 3 punishers treatment), while the contribution levels obtained in the treatments with 1 and 2 punishers are not statistically different ($p = .06055$ for 1 and 2 punishers treatment).

In order to support our hypothesis that the different punishment probabilities do not have a significant effect on subjects' contribution behavior, we analyzed the nonpunished defectors level (i.e., the proportion of subjects in the population that does not get punished after defecting), as depicted in Figure 2. Using the Mann–Whitney test ($\alpha = 5\%$), we show that from Round 31 to 40 the nonpunished defector level among the three treatments is statistically different ($p = .00151$ for 1 and 2 punishers treatments; $p = .00596$ for 1 and 3 punishers treatments; $p = .0124$ for 2 and 3 punishers treatments).

The previous two results provide support for our hypothesis: while from Round 31 to 40, the nonpunished defectors level in 1 and 2 punishers treatments is significantly different (Mann–Whitney test with $\alpha = 5\%$, data from Round 31 to 40, $p = .001511$), the contribution levels obtained in the two treatments are not significantly different (Mann–Whitney test with $\alpha = 5\%$, data from Round 31 to 40, $p = .06055$), suggesting that even though defectors are punished with different probabilities the humans behave similarly when the punishment act inflicted does not convey norm signaling (as in the treatments with 1 and 2 punishers). For this reason, we hypothesize that the large differences in the contribution levels are not an effect of a difference in the defector's probability to be punished.

Discussion

As the *same* material incentive is imposed in all the three treatments with punishment, we hypothesize that the explanation for the difference on the cooperation rates has to be found in additional information that the punished subjects receive. The result shown in Figure 1, that is, the fact that the same material incentive has a different effect if imposed only by uncoordinated punishers or by the entire group, is at odds with the Beckerian perspective that claims that punishment achieves deterrence *only* by modifying the relative costs and benefits obtained in a given situation, in such a way that wrongdoing becomes a less attractive option (Kreps, Milgrom, Roberts, & Wilson, 1982).

According to the Beckerian perspective, the three punishment treatments should be expected to yield similar results in terms of contribution. On the contrary, our data show that when the entire group punishes, the level of cooperation is significantly higher than when uncoordinated players inflict punishment (i.e., in the 1 and 2 punishers treatments).

Our hypothesis is that when inflicted by the entire group, the punishment recipient views punishment as expressing a message of peer condemnation, thus playing a sanctioning function. The normative information conveyed by sanction has the effect of informing the punished subject (and possibly the public) that the target behavior violated a norm and that the violation is not condoned. The punishment part of sanction, that is, the material cost inflicted, makes the expected consequences of violating norms more certain thus making them *salient* in the mind of the subjects. Salient norms lead, ceteris paribus, to higher compliance. We therefore claim that the cooperation observed in the treatment when the entire group punishes is the outcome of the positive interaction between the instrumental decision making and the norm psychology humans are provided with. The norm psychology is a cognitive machinery to detect and reason upon norms, characterized by a salience mechanism devoted to keep track of how much a norm is prominent within a group (see Andrighetto et al., 2013). The tandem work of cost avoidance and normative goals enables a higher and more durable compliance level with respect to an enforcement mechanism, such as uncoordinated punishment, that impacts cost avoidance goals only. The experimental data presented here seem to give support to this hypothesis.

Table 4. Norm Salience Mechanism: Cues and Weights.

Social Cue	Acronym	Weight
Norm compliance/violation	C	$w_c = (+/-)\ 0.99$
Observed norm compliance	O	$w_o = (+)\ 0.33$
Nonpunished violators	NPD	$w_{npd} = (-)\ 0.66$
Punishment observed/applied/received	P	$w_p = (+)\ 0.33$
Sanctioning observed/applied/received	S	$w_s = (+)\ 0.99$
Explicit norm invocation observed/applied/received	E	$w_e = (+)\ 0.99$

Group Punishment: A Simulation Experiment

To test our hypothesis about the decision process underlying the observed behavior, we develop an agent-based model that explicitly incorporates the norm psychology as part of its decision making. The motivation for norm compliance (i.e., for cooperating in this scenario) is modeled as dependent of both the salience of the norm and the instrumental decision making. Simulations reproduce the public goods game used in the laboratory and the experimental conditions: 0 punishers, 1 punisher, 2 punishers, and 3 punishers. The goal of the simulation is to check that the result obtained by agents in the simulation resembles that of the humans in the laboratory, supporting therefore that the explanation given in terms of norm psychology and norm salience is a plausible explanation for the observed human behavior. Agent-based modeling proves particularly apt to test our model, as it allows implementing separately the material damage and the normative information conveyed by sanction, and explore their respective and combined effects on motivating individuals to comply with the norm.

In order to capture the normative cues conveyed by group punishment, our simulation experiments are populated by agents able to extract the normative information contained in different types of punishments (i.e., group vs. uncoordinated one) and to include it into their decision making. The agent architecture used for such task is EMIL-I-A (Conte et al., in press; Villatoro et al., 2011).

This normative architecture allows agents to adapt to unstable environmental conditions, thanks to their ability to recognize norms, detect their salience and dynamically update the norm salience value on the basis of the social information they gather from the environment. More specifically, the EMIL-I-A architecture allows agents (a) to recognize norms, (b) to generate new normative representations and to act on them, and (c) to infer the normative information (explicitly or implicitly) conveyed by different types of punishment. The agent-based model is a dynamic one in which the propensity to follow the norm changes over time depending on the behavior observed during the interaction and, in this sense, goes beyond the purely static social preference models.

Salience Mechanism

The innovative aspect of this agent architecture is the presence of a norm salience mechanism indicating to the agents how prominent a certain norm is. The norm salience is a parameter, which is endogenous and dynamically updated at the end of every simulation round by each agent according to its personal decisions as well as the normative and the social information gathered by observing and communicating with others (see Andrighetto et al., 2013, for a detailed description).

The actions of others provide information or signs, such as the level of compliance, the amount of nonpunished norm violations, the frequency and typology of punishment, the consistency of the actions performed by other agents, through which individuals can infer how important and active, thus how salient, a norm is within a group. The cues that influence the norm salience updating and their respective weights are summarized in Table 4 ("Social Cue" and "Weight" columns), which are based on Cialdini, Kallgren, and Reno (1990).

At the end of each interaction, agents count the occurrences of each of the cues within their group and calculate the norm salience according to the function

$$\mathrm{Sal}_t^N = \mathrm{Sal}_t^N + \frac{1}{\phi}(w_c + O \times w_O + \mathrm{NPD} \times C \times w_{\mathrm{NPD}} + P \times w_P + S \times w_S + E \times w_E),$$

where Sal_t^N represents the salience of the norm N at time t, ϕ the normalization value, w_x the weights specified in Table 4, and finally O, NPD, C, P, S, E indicate the registered occurrences of each cue.

The resulting salience ($\mathrm{Sal}_t^N \in [0; 1]$, which 0 represents the minimum salience and 1 the maximum) is subjective for each agent, thus providing heterogeneity across agents as well as flexibility and adaptability to the system. This norm salience mechanism enables agents to monitor and track the strength of the norms in their social environment and to dynamically adapt to them. For example, in an unstable social environment, if a norm decays, our agents are able to detect this change, ceasing to comply with it and adapting to the new state of affairs. Moreover, if norm enforcement suddenly decreases, agents are less inclined to violate those norms that they perceive as highly salient. Perceiving a norm as highly salient is a reason for an agent to continue complying with it even in the absence of punishment. This guarantees a sort of inertia, making agents less prone to change their strategy in a pure reactive way (Villatoro et al., 2011).

Decision Making and Strategy Update

EMIL-I-A agents' decisions are modeled as dependent of both the salience of the norm and the instrumental decision making. More specifically, their decision making consists of the following drives: (1) *individual drive* (ID): this motivates agents to maximize their individual utility independently of what the norm prescribes and (2) *normative drive* (ND): this motivates agents to comply with the norm, independently of instrumental considerations.

The ID approximates the instrumental decision-making processes. It pushes agents to maximize their own personal utility regardless of what the norm prescribes and is updated according to a winner-stay-losers-change algorithm (Sutton & Barto, 1998). The more an action increases an agent's payoff, the higher the probability it will be chosen. The ID directs the choice toward contribution (C) only when the benefit of defecting is lower than the benefit of contributing. Agents' payoff depends on their actions, and they are lowered according to the costs sustained when imposing punishment or sanction and when receiving them.

The ID algorithm can be summarized as follows: Each agent checks its current payoff and compares it with the one achieved in the previous period. If the current payoff is greater than the previous one, an agent changes its strategy with a proportional probability: winners-stay-losers-change strategy (if the marginal reward is 0 and the last time step reward are the same, agents change their strategy with an inertial value in the same direction it last changed its probability). In formal terms, the ID is updated in the following way:

$$\mathrm{Orientation} = \begin{cases} 1 & \text{if Contributed} \\ -1 & \text{if Defected} \end{cases}$$

$$\mathrm{ID}_t = \begin{cases} \mathrm{Orientation} \times (R_t + R_{t-1}/R_{\mathrm{Normalizer}}) & \text{if } R_t \neq R_{t-1} \\ \mathrm{Orientation} \times R_{\mathrm{Inertial}} & \text{if } R_t = R_{t-1} \end{cases}.$$

where ID_t represents the value of the ID, Orientation is a value that represents whether agents' strategy change toward defection or contribution, R_t represents the payoff obtained at time t, $R_{\mathrm{Normalizer}}$ is a normalization value to ensure that ID_t is in the range [0,1], and R_{Inertial} represents an inertial value to ensure changes in the ID.

The ND models the motivation to comply as dependent on norm salience. Norm salience is a parameter updated by each agent at every round according to the information gathered by observing the behavior of the other agents and by communicating with them. The higher the number of individuals complying with the norm and enforcing it, the more salient the norm is perceived to be. Norms, depending on their perceived salience, affect the decision making of agents and elicit the compliant conduct accordingly. The more salient a norm is perceived to be, the higher its strength on the ND. In formal terms, the updating of the ND is given by

$$\mathrm{ND} = \mathrm{Sal}_t^N.$$

The contribution probability changes over time depending on the values that the ID and the ND of each agent take. The contribution probability varies across agents, thus generating heterogeneity within the population. The tendency to contribute is always positively affected by the ND and possibly by the ID, if contribution returns higher payoff than defection. In this case, the two drives complement each other. Conversely, it will be negatively affected by the ID, when defection returns higher payoff than contribution. In this second case, one drive goes against the other. The resulting probability to contribute in a specific time t is calculated by

$$p^t_{\mathrm{Contribution}} = p^{t-1}_{\mathrm{Contribution}} + (\mathrm{ID} \times \mathrm{IW} + \mathrm{ND} \times \mathrm{NW}),$$

where, $p^t_{\mathrm{Contribution}}$ is the probability of contributing at time t and, ID and ND are, respectively, the individual and the ND (whose values are adapted at each round), and the IW and the NW are the individual weight and normative weight. IW and NW are fixed at 0.5 for all the simulations, and represent the importance that each agent assigns to the ID and ND. By varying these values different types of agent can be designed, varying from the more individualistic to the more normative types of agents.

In particular, receiving punishment has two distinct effects on EMIL-I-A agents' decision making: (a) affects the agent's ID (thus reducing the probability that the action that triggered the punishing reaction will be performed again) and (b) it conveys normative information impacting also the norm's salience and consequently the ND.

Simulation Design

As in the present laboratory experiment, in the simulation experiments, agents are also arranged in groups of four. Each group consists of one agent endowed with the normative architecture EMIL-I-A and three confederate players (not endowed with the EMIL-I-A architecture and that follow the same behavioral rules of the confederate virtual players). EMIL-I-A agents play the same public goods game that human subjects played and the same four different treatments (0 punishers, 1 punisher, 2 punishers, and 3 punishers) are reproduced. The EMIL-I-A agents' initial Norm Salience is set to 0. When an EMIL-I-A agent receives punishment from two agents, this multiplicity of punishment has a stronger impact on the norm salience than when inflicted by one single agent (see Table 4). Finally, when punishment is unanimously imposed by the entire group, EMIL-I-A interprets it as an act that conveys a strong normative message of peer condemnation. In addition to run the simulation experiments using EMIL-I-A agent architecture, a second set of simulations experiments has been performed in which EMIL-I-A agents have been replaced by reinforcement learning (RL) agents. The purpose of this second set of experiments is contrasting the behavior of our EMIL-I-A architecture against an architecture whose decisions are purely taken on an instrumental basis (i.e., RL agents) in order to evidence that EMIL-I-A better represents the way humans take advantage of norm signaling conveyed by group punishment while the RL cannot.

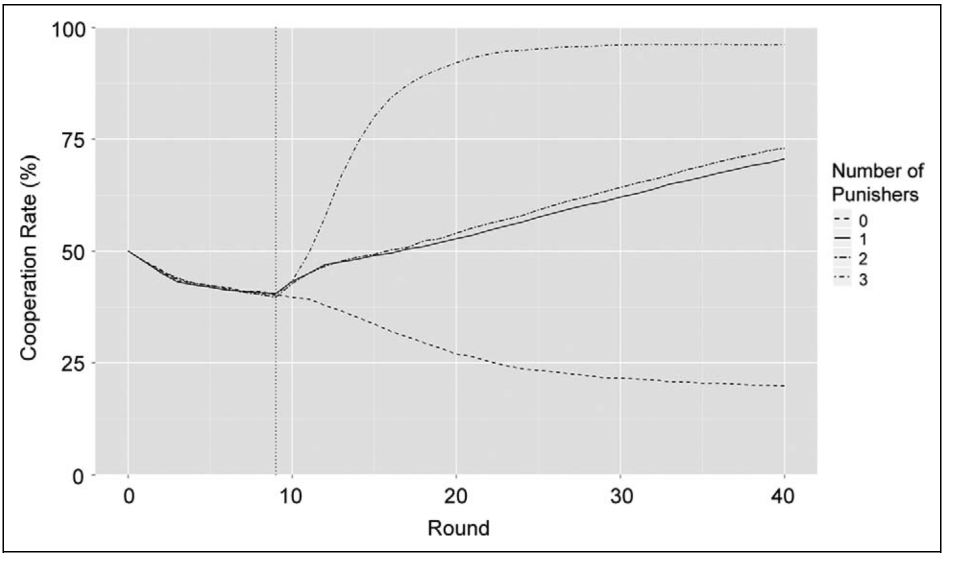

Figure 3. Simulation results with EMIL-I-A agents.

From the different RL architectures available, we decide to use a Q-Learning (Watkins & Dayan, 1992) for two reasons: (1) it better represents the learning process of a utility-based human because it will maintain its strategy while obtaining benefits; otherwise, it changes its behavior to learning and (2) its basic parameter configuration allows a simpler analysis, yet with a powerful potential of applications. Q-Learning algorithms base its functioning on an exploration/exploitation approach. Agents explore randomly accordingly to an exploration rate that in this work is fixed to 5% to maintain a similarity to the EMIL-I-A configuration. When agents decide not to explore, they will choose the action with the highest Q value. These Q values are initialized to 0 and updated using a Q-Update function. The Q-Update function for estimating the utility of an action a is

$$Q^t(a) \leftarrow (1-\alpha) \times Q^{t-1}(a) + \alpha \times \text{reward},$$

where, α represents the learning rate, reward is the current interaction payoff, and $Q^t(a)$ is the utility estimate of action a after selecting it t times.

Results

Figure 3 shows the average EMIL-I-A contribution rates obtained in the four treatments based on 5,000 simulation replications (the raw data are available upon request to the authors, as well as the R Statistics scripts developed to process them and generate the figures). It is interesting to notice that the contribution dynamics achieved in the simulation experiment with EMIL-I-A agents are similar to the ones obtained in the experiment with human subjects (as shown in Figure 1): The 3 punishers treatment is the one in which a higher and faster contribution level is achieved. Even though the same material damage is imposed in all the treatments, group punishment affects differently agents' minds with respect to the other two treatments with punishment. The larger the amount of punishers, the stronger the effect it has on the norm salience of the punished agent. While the norm salience increases proportionally to the number of punishers, the ID is affected in the same way in all the treatments. The combined effect of the ND and the ID results in a similar dynamics to those obtained by the human subjects.

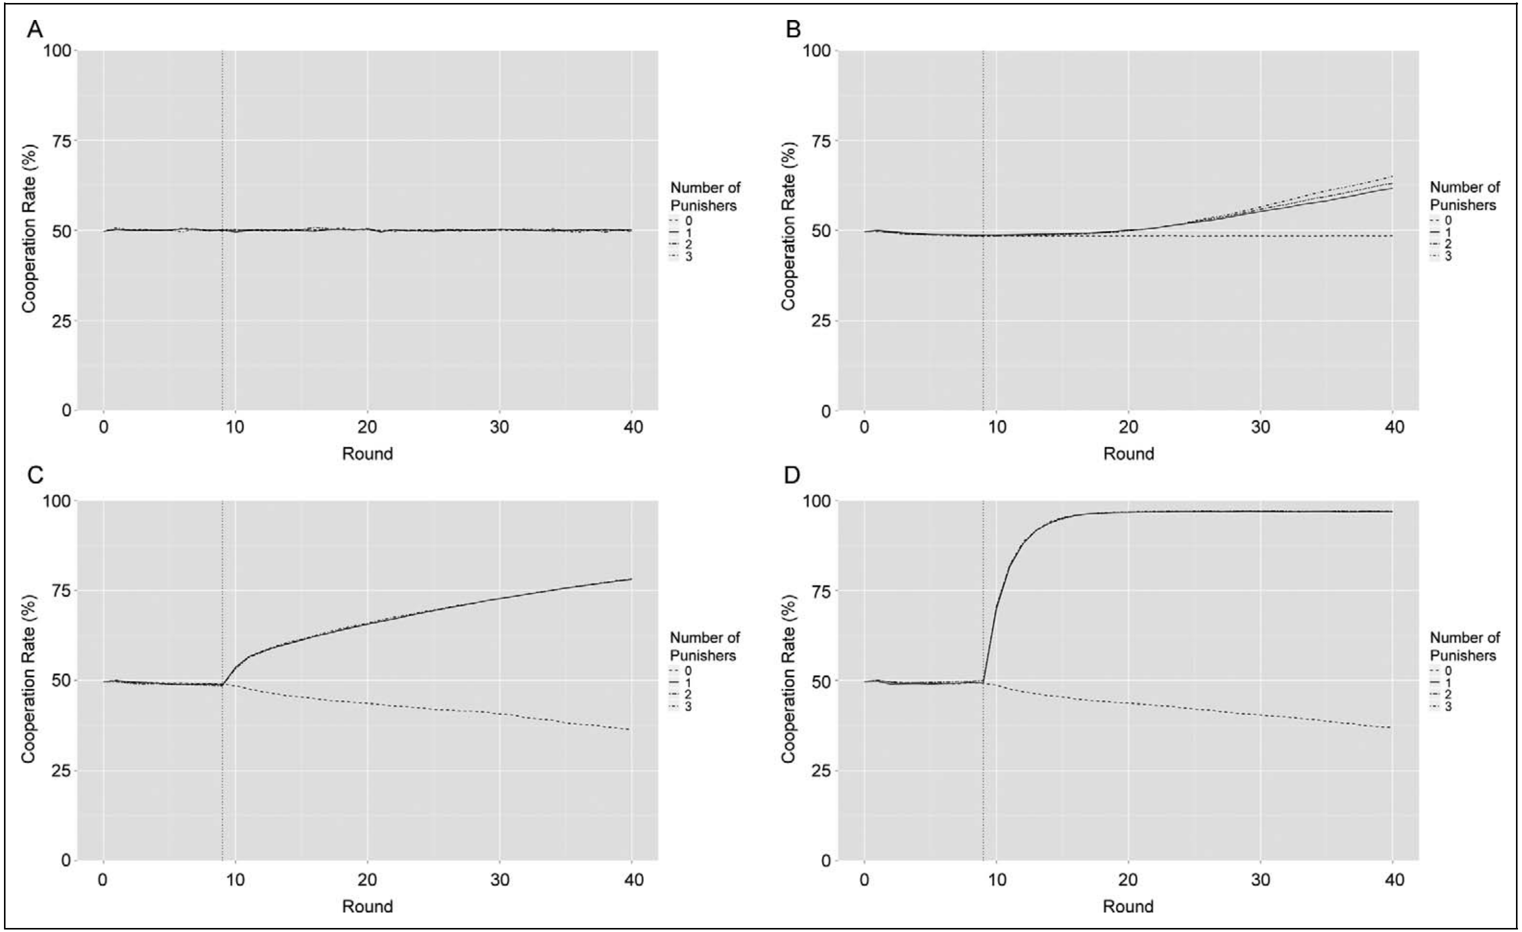

Figure 4. Simulation results with Q-Learning agents.

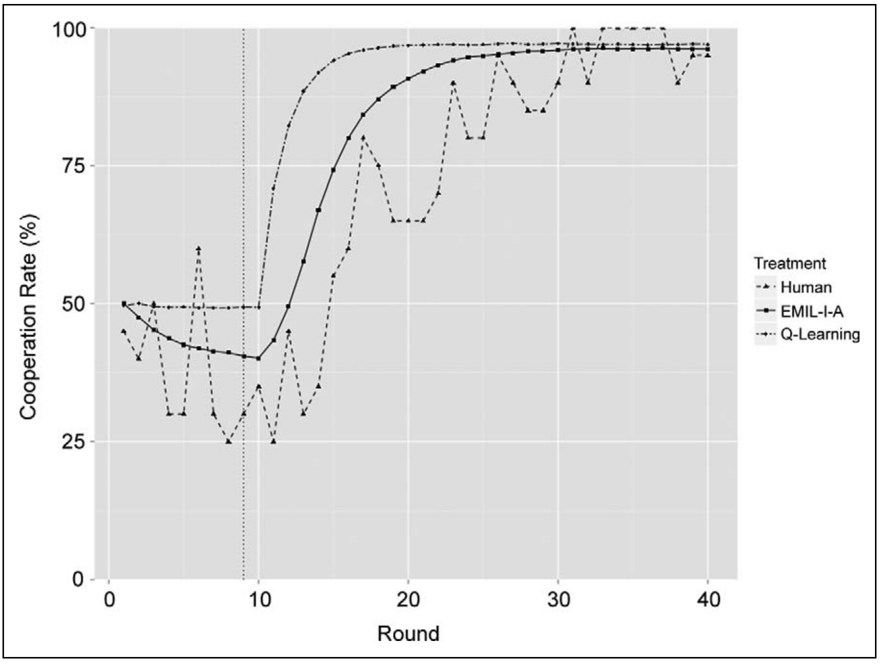

Figure 5. Comparative simulation results.

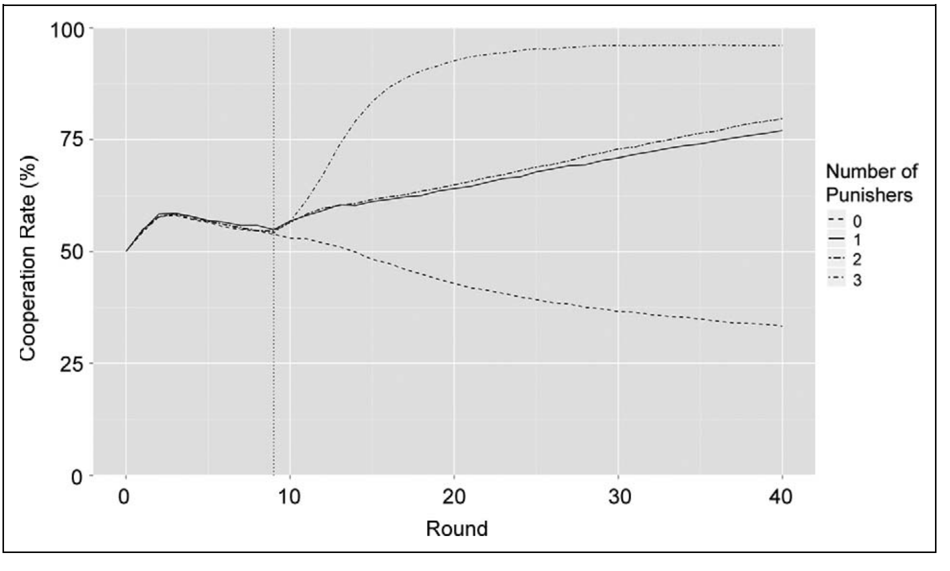

Figure 6. Simulation results with EMIL-I-A agents (initial salience = 0.5).

In order to show the advantage of EMIL-I-A in comparison to agents whose decisions are purely taken on an instrumental basis (i.e., RL agents), we have performed simulations using RL Q-Learning agents in replacement to the EMIL-I-A agents. In Figure 4, the contribution levels of RL agents with different learning probabilities are shown.

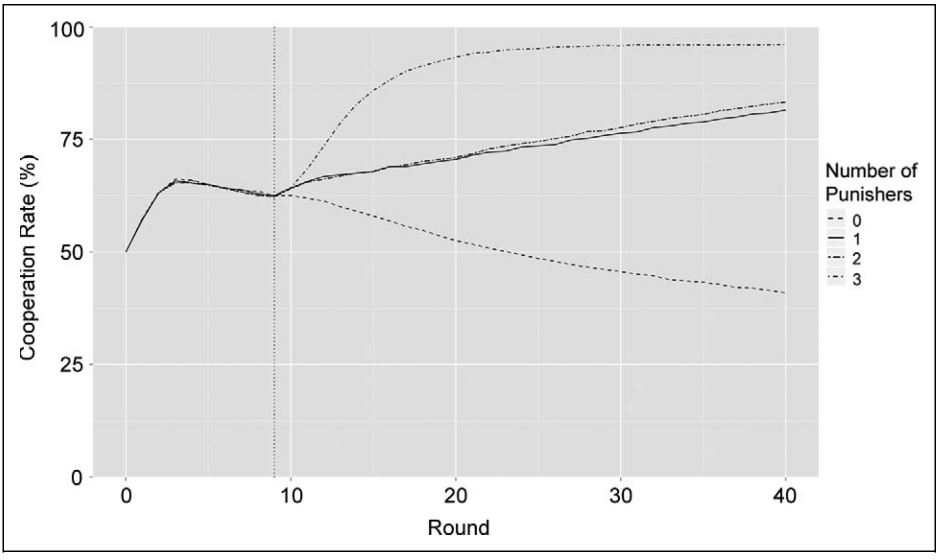

Figure 7. Simulation results with EMIL-I-A agents (initial salience = 0.8).

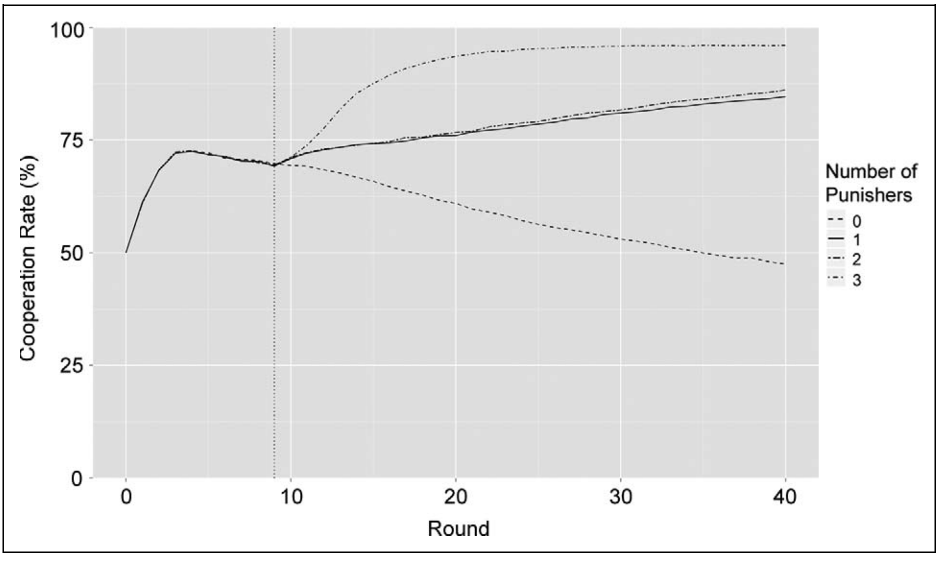

Figure 8. Simulation results with EMIL-I-A agents (initial salience = 1.0).

We can observe that the variation of treatments from 0 to 3 punishers (and the consequent different punishment probabilities) does not affect the contribution rates in a significant way. For example, even considering the 10% learning rate case, which has the highest distinction among treatments, when analyzed using Mann–Whitney test ($\alpha = 5\%$), it does not present any significant difference among treatments ($p = .99438$ for 1 and 2 punishers treatments; $p = .82304$ for 1 and 3 punishers treatments; and $p = .92206$ for 2 and 3 punishers treatments). Moreover, Figure 4 also shows that the RL agents are sensitive to the variation with respect to the value of their learning rate represented by α (i.e., the learning rate). Higher values of α make RL agent learn faster from the most recent actions.

In Figure 5, the contribution levels obtained in the 3 punishers treatment by human subjects, EMIL-I-A agents and RL agents (learning rate 100%) are compared. Data show that the RL and EMIL-I-A obtain contribution levels similar to humans, confirming that the instrumental motivation in humans is very strong. However, analyzing the results presented in Figures 3 and 4, we observe that the EMIL-I-A agents behave differently when punished by 1, 2, or 3 punishers, while the RL agents do not present such a distinction. This difference is also observed in the experiment with human subjects. As observed in the experiments plotted in Figure 6, 7, and 8, EMIL-I-A general behavior is not affected by the initial salience value; the resulting dynamics remain unaltered in the different configurations showing that the initial salience does not affect agent's decision making. Therefore, these results, and other experimental data (Andrighetto et al., 2013; Bowles & Gintis, 2003; Boyd, Gintis, & Bowles, 2010; Li, Xiao, Houser, & Montague, 2009), support our claim that successful human contribution results from the interaction of the norm psychology and the cognitive machinery for instrumental decision making that EMIL-I-A closely replicates.

Conclusion

As claimed by Boyd et al. (2010), in natural social contexts, punishment is usually coordinated, and coordination has the effect of reducing the cost of punishment and the 2-fold benefit of legitimizing the sanction and reducing the likelihood of counterpunishment acts (see also Guala, 2012). In this study, we provide cross-methodological evidence in support of the effectiveness of one specific form of coordinated punishment, that is, group punishment, in promoting cooperation. More generally, our data show that punishment is effective in regulating people's behavior not only through economic incentives but also thanks to the normative information it conveys. Group punishment is a powerful tool through which messages of peer condemnation are expressed.

The comparison of the results from the experiments with humans and from the simulations shows that the simulation model captures the essential features of the human data. The simulation model is not intended just to replicate in silico the experimental findings, but it is an attempt to provide an explicit model of the cognitive mechanisms and processes allowing group punishment to positively promote compliant conduct. By comparing the performance of the normative architecture EMIL-I-A with other classical RL architectures, we have observed that the former reproduces behavioral dynamics more similarly to humans than the latter ones.

This work has served us to understand the relative effect of uncoordinated and group punishment in promoting compliance. The most immediate extension of our work would be to conduct new simulation experiments in which we vary the size of the group to test whether interacting in larger or smaller groups makes a difference in terms of group punishment. Moreover, we want to perform laboratory experiments with humans to understand how the punishers' intentions are perceived by the punished subjects (and possibly the public) and how this affects their future behavior.

Appendix

In order to show that the EMIL-I-A agents' behavior is not affected by the initial *Norm Salience* value (0) that we set on the experiments, here, we present the graphics generated using different initial *Norm Salience* values. The only observed difference is from Round 0 to 9, in which agents have a higher cooperative behavior when initial *Norm Salience* is greater than 0 than the one presented in Figure 3.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: MacNorms project (Intramurales de frontera CSIC—PIF08-007) The Institute of Cognitive Science and Technologies (ISTC-CNR, Rome), the European University Institute, the Spanish Ministerio de Economía y Competitividad (Grant: ECO2011-29847-C02-01), the Generalitat de Catalunya (Grant: 2009 SGR 820 and Grant 2009SGR1434), the Antoni Serra Ramoneda Research Chair (UAB-CatalunyaCaixa) for supporting the authors.

References

Andrighetto, G., Brandts, J., Conte, R., Sabater-Mir, J., Solaz, H., & Villatoro, D. (2013). Punish and voice: Punishment enhances cooperation when combined with norm-signaling. *PLoS ONE*, *8*, e64941.

Andrighetto, G., & Castelfranchi, C. (2013). Norm compliance: The prescriptive power of normative actions. *Paradigmi*, *2*, 152–168.

Axelrod, R. (1987). *The evolution of strategies in the iterated prisoner's dilemma*. Los Altos, CA: Kaufmann.

Becker, G. S. (1968). Crime and punishment: An economic approach. *The Journal of Political Economy*, *76*, 169–217.

Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge, England: Cambridge University Press.

Boehm, C. (1999). *Hierarchy in the forest: The evolution of egalitarian behavior*. Cambridge, MA: Harvard University Press.

Bowles, S., & Gintis, H. (2003). Origins of human cooperation. In P. Hammerstein (Ed.), *Genetic and cultural origins of cooperation*. Cambridge, MA: MIT Press.

Bowles, S., & Polanìa-Reyes, S. (2012). Economic incentives and social preferences: Substitutes or complements? *Journal of Economic Literature*, *50*, 368–425. doi:10.1257/jel.50.2.368

Boyd, R., Gintis, H., & Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science*, *328*, 617–620.

Bravo, G., Squazzoni, F., & Boero, R. (2012). Trust and partner selection in social networks: An experimentally grounded model. *Social Networks*, *34*, 481–492.

Brito, I., Pinyol, I., Villatoro, D., & Sabater-Mir, J. (2009). Hiherei: Human interaction within hybrid environments regulated through electronic institutions. In Decker Sichman, G. Sierra, & C. Castelfranchi (Eds.), *AAMAS '09: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems* (pp. 1417–1418). Budapest, Hungary: International Foundation for Autonomous Agents and Multiagent Systems.

Chudek, M., & Henrich, J. (2011). Culture-gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in Cognitive Sciences*, *15*, 218–226. doi:10.1016/j.tics.2011.03.003

Cialdini, R. B., Kallgren, C. A., & Reno, R. R. (1990). A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. *Advances in Experimental Social Psychology*, *24*, 201–234.

Conte, R., Andrighetto, G., & Campennì, M. (Eds.). (in press). *Minding norms. Mechanisms and dynamics of social order in agent societies*. Oxford Series on Cognitive Models and Architectures. New York, NY: Oxford University Press.

Conte, R., & Castelfranchi, C. (1995). *Cognitive and social action*. London, England: UCL Press.

Crawford, S., & Ostrom, E. (1995). A grammar of institutions. *American Political Science Review*, *89*, 582–600.

Denant-Boemont, L., Masclet, D., & Noussair, C. (2007). Punishment, counter punishment and sanction enforcement in a social dilemma experiment. *Economic Theory*, *33*, 145–167.

Dreber, A., Rand, D. G., Fudenberg, D., & Nowak, M. A. (2008). Winners don't punish. *Nature*, *452*, 348–351.

Duffy, J. (2006). Agent-based models and human-subject experiments. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Vol. 2: Agent-based computational economics. Handbooks in Economics Series* (pp. 949-1011). Amsterdam, the Netherlands: North-Holland/Elsevier.

Egas, M., & Riedl, A. (2008). The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society of London—Series B*, *275*, 871–878.

Ellickson, R. (1991). *Order without law: How neighbors settle disputes*. Cambridge, MA: Harvard University Press.

Faillo, M., Grieco, D., & Zarri, L. (2013). Legitimate punishment, feedback, and the enforcement of cooperation. *Games and Economic Behavior*, *77*, 271–283.

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*, 137–140.

Fehr, E., & Rockenbach, B. (2003). Detrimental effects of sanctions on human altruism. *Nature*, *422*, 137–140.

Galbiati, R., & Vertova, P. (2008). Obligations and cooperative behavior in public good games. *Games and Economic Behavior*, *64*, 146–170.

Giardini, F., Andrighetto, G., & Conte, R. (2010). A cognitive model of punishment. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 1282–1288). Austin, TX: Cognitive Science Society.

Gneezy, U., & Rustichini, A. (2000). A fine is a price. *The Journal of Legal Studies*, *29*, 1–17.

Guala, F. (2012). Reciprocity: Weak or strong? What punishment experiments do (and do not) demonstrate. *Behavioral and Brain Sciences*, *35*, 1–59.

Herrmann, B., Thoni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, *319*, 1362–1367.

Heyman, J., & Ariely, D. (2004). Effort for payment: A tale of two markets. *Psychological Science*, *15*, 787–793.

Janssen, M. A., Holahan, R., Lee, A., & Ostrom, E. (2010). Lab experiments for the study of social-ecological systems. *Science*, *328*, 613–617.

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, *27*, 245–252. doi:10.1016/0022-0531(82)90029-1

Li, J., Xiao, E., Houser, D., & Montague, P. R. (2009). Neural responses to sanction threats in two-party economic exchange. *Proceedings of the National Academy of Sciences of the United States of America*, *106*, 16835–16840.

Masclet, D., Noussair, C., Tucker, S., & Villeval, M.-C. (2003). Monetary and non-monetary punishment in the voluntary contribution mechanism. *American Economic Review*, *93*, 366–380.

Noussair, C., & Tucker, S. (2005). Combining monetary and social sanctions to promote cooperation. *Economic Inquiry*, *43*, 649–660.

Ostrom, E. (1990). *Governing the commons. The evolution of institutions for collective action*. Cambridge, MA: Cambridge University Press.

Ostrom, E., Walker, J., & Gardner, R. (1992). Covenants with and without a sword: Self-governance is possible. *American Political Science Review*, *86*, 404–417.

Posner, E. A. (2000). *Law and social norms*. Cambridge, MA: Harvard University Press.

Poteete, A. R., Janssen, M. A., & Ostrom, E. (Eds.). (2010). *Working together: Collective action, the commons, and multiple methods in practice*. Princeton, NJ: Princeton University Press.

Sigmund, K. (2007). Punish or perish? Retaliation and collaboration among humans. *Trends in Ecology & Evolution*, *22*, 593–600. doi:10.1016/j.tree.2007.06.012

Sripada, C., & Stich, S. (2006). A framework for the psychology of norms. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Culture and cognition* (pp. 280–301). Oxford, England: Oxford University Press.

Sunstein, C. R. (1996). Social norms and social roles. *Columbia Law Review*, *96*, 903–968.

Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Villatoro, D., Andrighetto, G., Conte, R., & Sabater-Mir, J. (2011). Dynamic sanctioning for robust and cost-efficient norm compliance. In T. Walsh, NICTA, & University of NSW (Eds.), *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 414–419). Barcelona, Spain: AAAI Press.

Watkins, C., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279–292. doi:10.1023/A:1022676722315

Xiao, E. (2013). Profit seeking punishment corrupts norm obedience. *Games and Economic Behavior*, *77*, 321–344.

Xiao, E., & Houser, D. (2005). Emotion expression in human punishment behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 7398–7401.

Xiao, E., & Houser, D. (2011). Punish in public. *Journal of Public Economics*, *95*, 1006–1017.

Yamagishi, T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology*, *51*, 110–116.

Author Biographies

Daniel Villatoro obtained his PhD in artificial intelligence at the IIIA-CSIC in Barcelona. His thesis was awarded with the Victor Lesser Distinguished Dissertation Award by the IFAAMAS. During his research on normative multiagent systems, he has visited and actively collaborated with Prof. Sandip Sen (University of Tulsa, USA), Dr. Giulia Andrighetto (ISTC-CNR, Italy), and Prof. Michael Luck (King's College London, UK). Moreover, he has been an active member of the community publishing at specialized journals (*JAAMAS*, *ACM TAAS*, or *PLoS-One*) and organizing workshops (MABS2011 or Citisen2012-13). His research interests focus on modeling human behavior, self-policing multiagent systems, and complex systems; e-mail: dvillatoro@iiia.csic.es.

Giulia Andrighetto is a researcher at the Institute of Cognitive Sciences and Technologies (ISTC-CNR) in Rome and at the EUI Political and Social Sciences Department in Florence. She received her PhD in philosophy from the University of Rome "La Sapienza" in Rome. She has been visiting fellow at the Nuffield College (Oxford), Tilburg University, University of Paris 7, and at the Workshop in Political Theory & Policy Analysis (Indiana University). Her main research interest is in developing a socio and cognitive theory of the dynamics of norms. She uses cognitive modeling, agent-based simulations, and experiments in the lab to validate her hypotheses. She has recently published the book *Minding Norms. Mechanisms and Dynamics of Social Order in Agent Societies* (with R. Conte and M. Campenì) for Oxford University Press; e-mail: giulia.andrighetto@istc.cnr.it.

Jordi Brandts holds a PhD in economics from the University of Pennsylvania. He is a research professor at the Institute for Economic Analysis (CSIC) in Barcelona and at present holds the Serra-Ramoneda/Catalunya Caixa Chair at the Department of Business of the Universitat Autonoma de Barcelona. He is also affiliated with the Barcelona Graduate School of Economics and a Research Fellow of CESifo. His research is experimental in areas such as the study of cooperation, organizational economics, industrial organizational and market analysis, conflict and the effects of communication on strategic interaction. From 2007 to 2011, he was editor in chief of the journal *Experimental Economics*. Since 2013, he serves as advisory editor for games and economic behavior; e-mail: jordi.brandts@iae.csic.es.

Luis Gustavo Nardin is a computer engineering graduate student at University of São Paulo, from where he has obtained his master's degree and a specialization in software engineering. He has obtained his undergraduate degree in computer engineering from the University São Francisco. He has experience in multiagent systems, more specifically reputation, interoperability, and multiagent-based simulation. Currently, he is a research fellow of the Institute of Cognitive Sciences and Technologies and is mainly interested on regulation mechanisms for multiagent systems; e-mail: gustavo.nardin@istc.cnr.it.

Jordi Sabater-Mir is a tenured scientist at the Artificial Intelligence Research Institute (IIIA) of the Spanish National Research Council (CSIC), Barcelona, Spain. He holds a doctorate in artificial intelligence and was a postdoctoral Marie Curie fellow at the Institute of Cognitive Sciences and Technologies (ISTC-CNR) in Rome, Italy. His research is focused on computational models of trust and reputation, agent-based social simulation, and computational cognitive models. He has published more than 100 articles in journals (*JAAMAS*, *Knowledge-Based Systems*, *PLOS-One* . . .) and international conferences (AAMAS, IJCAI, ECAI . . .) and has been the main researcher in several European and national research projects. He has been PC member of the main conferences and workshops in the area of MAS and has organized several workshops in the area of computational trust and reputation systems. Currently he is the president of the European Association for Multi-Agent Systems (EURAMAS); e-mail: jsabater@iiia.csic.es.

Rosaria Conte is a research director at the Institute of Cognitive Sciences and Technologies CNR and vice president of the Scientific Council of CNR. She is the head and founder of the LABSS, Laboratory on Agent-Based Social Simulation. She has published more than 160 works and is recognized as an expert on dependence networks, reputation, and the dynamics of social norms; e-mail: rosaria.conte@istc.cnr.it.