Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.
Digital Object Identifier xxx

Empirical evidence of upward social comparison in a prisoner's dilemma game

Matthieu Nadini[1], Peerayos Pongsachai[1], Chiara Spinello[1], Daniel Burbano Lombana[1] (Member, IEEE), and Maurizio Porfiri[1,2] (Fellow, IEEE)
[1]Department of Mechanical and Aerospace Engineering, New York University Tandon School of Engineering, Brooklyn, NY, 11201, USA
[2]Department of Biomedical Engineering, New York University Tandon School of Engineering, Brooklyn, New York 11201, USA
Corresponding author: Maurizio Porfiri (e-mail: mporfiri@nyu.edu).

The authors acknowledge financial support from the National Science Foundation under grant No. CMMI-1561134.

ABSTRACT A large body of work has offered compelling evidence of the influence of social context on individual decision-making, but the reasons why individuals tend to cooperate with others remain elusive. The prisoner's dilemma constitutes a powerful, yet elementary, social game to study the drivers underlying cooperation. Here, we empirically examined a prisoner's dilemma game where small groups of participants played with controlled, virtual players over a series of rounds. Toward investigating how individual decisions on cooperation are influenced by others, the virtual players were engineered so that they would have a higher cumulative score than some participants and a lower cumulative score than others. Our results corroborate upward social comparison theory, whereby only participants who had a lower cumulative score than cooperating virtual players displayed an increased tendency to cooperate. Overall, our experimental findings indicate that the players' cumulative score plays a critical role within the prisoner's dilemma game, thereby offering a mean for increasing cooperation. For practitioners, this finding sheds light on how players' cumulative score alone modulates decision-making processes toward choices that are suboptimal for the individual, but optimal for the entire group.

INDEX TERMS Cooperation, Cumulative Score, Small group, Social game, Virtual players.

I. INTRODUCTION

WHY people cooperate is a long-standing question, which bears important ramifications in biology, political sciences, and economics [1], [2]. Cooperation is defined as the process where two or more individuals exchange benefits or work together to obtain a higher reward than they would obtain working alone for a selfish purpose [1]. Potential drivers for cooperation in humans have been extensively investigated from a game-theoretical perspective [3]–[5].

Among the most popular games is the prisoner's dilemma, a non zero-sum game where two players simultaneously choose to either cooperate or defect and receive appropriate scores (or payoffs) according to their choices [6]. In both one-time and iterated prisoner's dilemmas, the Nash equilibrium [7]–[9] that maximizes individual payoffs corresponds to the case where both players defect, even though the group payoff is maximized when all players cooperate. This dichotomy between individual and group benefits is found in many versions of the iterated prisoner's dilemma, where the number of players simultaneously participating in the iterated prisoner's dilemma ranges from two [10] to more than a thousand [11].

A pressing challenge is to identify the social drivers that could lead to cooperation. For example, an empirical study on the iterated multiplayer prisoner's dilemma demonstrated that cooperation emerges through the tit-for-tat strategy [12], where individuals choose to cooperate only when others have cooperated in the previous round. Other social drivers of cooperation include peers' reputation [13], game's fairness [14], friendship [15], [16] or acquaintanceship [17], and heterogeneity of players' strategy [18]. These studies indicate the importance of a social context in individual decision-making toward cooperation [6], [19]–[21].

A less studied aspect in the iterated prisoner's dilemma is the relationship between players' cumulative score and their choices. An improved understanding of this relationship might offer insight into dynamics of many complex systems where decision-making is linked to performance. In competitive sports, for instance, high performing teams are known to attract more sponsorship than others [22]. Likewise, higher quality institutions attract more investments [23]. A similar scenario is also documented in economics, where past studies have proposed that subjects with higher worth tend to be in privileged positions, which may further attract additional worth – this concept can be summarized in the "rich get richer" hypothesis [24]. Despite the empirical evidence of the role of performance on individual decision-making, the role of the players' cumulative score in the prisoner's dilemma is yet to be experimentally studied.

Will players choose to cooperate even if it is not the optimal choice for them when top-performing individuals cooperate? And, what will happen, instead, when low-performing individuals cooperate? To the best of our knowledge, these questions are yet to be answered and, more in general, the relationship between cumulative score and choices in the prisoner's dilemma has only been investigated from a theoretical point of view in [25]. Therein, it has been shown that only those individuals with high cumulative scores can influence the peers' answers over time. In this mathematical model, individuals play the iterative prisoner's dilemma against their peers and imitate the strategy of the peer who has the highest cumulative score. Players are also given the possibility to rewire the network of interactions, whereby they can punish or reward their peers by removing or sustaining interactions with them. Theoretical results suggest that only the individuals with high cumulative score can enhance or promote cooperation in the group. Such a possibility is supported by upward social comparison theory [26]–[28], in which peers would mimic the choices of better performing individuals.

In order to empirically investigate the relationship between players' cumulative scores and their choices, we propose the use of virtual players. The possibility of integrating virtual players affords the opportunity to create unique social experiments, in which the experimenter has full access to several independent variables that would be impossible to manipulate otherwise.

For instance, in [29], the use of virtual players was shown to be successful in sustaining cooperation in the iterated prisoner's dilemma for a number of rounds. In [30], the authors explored virtual players that could interact with players through different communication means (text, voice, or face and voice), concluding that cooperation is favored by interaction through text. More recent research on communication pathways of virtual players has also been conducted by [31], where facial expressions of avatars were found to have a subtle effect on cooperation in the prisoner's dilemma. In [32], it was found that virtual players can promote cooperation better than humans. However, cooperation decreases when information about the nature of the peer is revealed. Likewise, in [33], it was determined that people tend to cooperate more with people rather than robots in the prisoner's dilemma game.

Our past research [34], [35] also offers evidence about the potential of virtual players in social experiments, where we found virtual players to be conducive to increasing players' commitment and engagement. Interestingly, the use of virtual players is not limited to humans; for example, in [36], results on pigeons playing an iterated prisoner's dilemma against virtual players, indicate that pigeons are impulsive decision-makers that tend to defect.

Through the use of virtual players, we sought to examine whether individuals would be willing to choose a strategy that is suboptimal for them (cooperation), but optimal for the group, if a high-performing individual was systematically cooperating. To delve into the role of social context and in contrast to previous research [29]–[32], we systematically varied the number of virtual players in the group and their strategy. These factors have been proposed to be determinants of individual decision-making processes in the iterated prisoner's dilemma game [37]–[39], but never have they been examined in conjunction with virtual players.

Ultimately, we formulated two hypotheses. First, we hypothesized that only a high-performing cooperating player is able to incentivize others toward cooperation. Second, we hypothesized that individual choice would be modulated by group decision-making, such that an individual would have a higher tendency to cooperate when the majority of the other players chose to cooperate.

To empirically test these two hypotheses, we simultaneously conducted two experiments in which groups of three participants played the iterated prisoner's dilemma with virtual players. The participants were instructed to play with each other and were not informed about the presence of virtual players in the game. We recruited a total of 144 participants in groups of three. For each group, two out of the three participants played with one virtual player, and the remaining participant played with two virtual players. All the 144 participants were part of a competition for monetary rewards, which was introduced as a motivational factor to maintain players' engagement [40], [41]. The game was conducted through a computer where participants interacted only through the information displayed on their computer screens. At the end of the experiment, participants were asked to fill out a survey with general demographic information and questions regarding the experiment itself.

The first experiment (involving two participants and one virtual player) was designed to test whether a high-performing player is able to incentivize other players toward cooperation. The behavior and presentation of the virtual player were varied with respect to the choice (mostly cooperate or mostly defect) and performance (high or low cumulative score). We expected that participants would cooperate more only when the high-performing virtual player would mostly cooperate.

The second experiment (involving one participant and two virtual players) was designed to test whether a participant is more likely to cooperate when the majority of the other players cooperate, despite he/she would gain more by defection. To this end, we programmed the first of the two virtual players to either mostly cooperate or mostly defect, as in the first experiment, and the second virtual player to either copy the choice of the first one from the previous round or do the opposite. To further delve into the underlying drivers of cooperation, we explored correlations between participants' choices in the game and their knowledge about the game and their perception of the other players.

II. MATERIALS AND METHODS

A. EXPERIMENTAL SETTING

The experimental setup consisted of three computer monitors (Acer, V246HQL), three partition walls (Amazon, SKY3028), two Raspberry Pi 3 Model B (RS Components Ltd and Allied Electronics), and five custom-made remote controllers. A sketch of the setup is presented in Fig. 1b. Each controller included a microcontroller (LightBlue Bean+, Punch Through) in a custom-made case (15.7×4.6×3.8 cm, length × width × height), with five pushbuttons labelled as "A", "B", "C", "D", and "E". Different from [42] where all the buttons were needed, for the purpose of this study, we covered buttons "A", "B", and "E" with black tape. The remaining "C" and "D" were used for choosing cooperation and defection, respectively. The controllers were wirelessly connected to the Raspberry Pi's via Bluetooth.

We recruited participants in groups of three. Participants were deceivingly informed that they were playing together, as in Fig. 1a, although they were actually assigned to two groups, as shown in Fig. 1b. One group was composed of participants 1 and 2 with virtual player 3, and the other was comprised of participant 3 with virtual players 1 and 2. Each group was assigned to one of the two experiments: one virtual player with two participants (1V2P) and two virtual players with one participant (2V1P). Experiments 1V2P aimed at testing the first hypothesis that a cooperating player with a high cumulative payoff is able to incentivize others toward cooperation, while experiment 2V1P was designed to test the second hypothesis that a participant would tend to cooperate when interacting with two cooperating players.

B. ITERATED PRISONER'S DILEMMA GAME

We implemented the iterated prisoner's dilemma for both experimental settings 1V2P and 2V1P. Each game comprised of 45 successive rounds. The payoff matrix used in our experiment is given in Fig. 2a, where we set the cooperation payoff $R = \$30$, the "punishment" payoff $P = \$5$, the "temptation" payoff $T = \$50$, and the "sucker" payoff $S = \$0$. Monetary values displayed refer to virtual amounts. The payoffs were set such that $T > R > P > S$ and $2R > T + S$, following the classical iterated prisoner's dilemma [6]. Indeed, the payoff relationships $T > R$ and $P > S$ imply that defection is the dominant strategy. In addition, the relationships $R > P$ and $2R > T + S$ guarantee that mutual cooperation is favored over mutual defection and that mutual cooperation is favored over alternating a defection with a cooperation, respectively [6].

In our experimental setting, we considered an all-to-all network of interactions, where each player simultaneously played two separate games, each against one of the two peers. The same answer was used in both games. For example, in Fig. 2b, we illustrate the case when only one player cooperated, while the others defected for a single round. For this combination, the player who cooperated received \$0 and the two players who defected gained \$55 each, as shown in Fig. 2c.

C. COMPUTER INTERFACE

The computer interface was developed in Python GUI programming to run two different experiments and consisted of two phases: phase 1 included the first 15 rounds and phase 2 the remaining 30 rounds. Two screenshots that exemplify these two phases are shown in Fig. 3a and Fig. 3b, respectively. In both phases, the interface provided general information, such as the possible choices (cooperate and defect), the relative cumulative scores according to the payoff matrix in Fig. 2a, and the seconds left for the players to make their choice out of the maximum of four seconds. If a player failed to choose within four seconds, their choice in the previous round was used. If they failed to choose in the first round of the experiment, the choice (either cooperation or defection) was randomly selected.

During phase 1, players saw all relative percentage scores with respect to the winner (who had the highest cumulative payoff), as shown in Fig. 3a. For instance, if a player had half of the points of the winner, the height of the bar was scaled to 50%. This phase served to build the performance of each player. During phase 2, participants were able to see all players' choices from their previous round, their nominal increment from the two interactions in the previous round, and all cumulative scores from the beginning of that game, as shown in Fig. 3b. This phase was designed to study whether participants' decisions were influenced by the choices of the other players.

D. VIRTUAL PLAYERS' STRATEGIES

In Tables 1 and 2, we summarize the strategies of the virtual players during phase 1 and phase 2 for

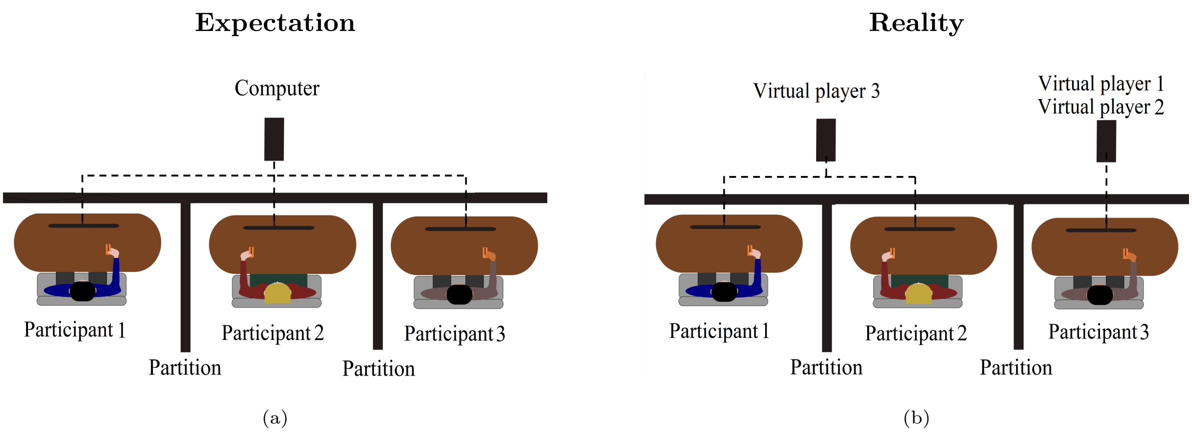

FIGURE 1. Experimental setting: participants' expectation versus implemented defection. Solid dark lines represent partitions that prevented participants from seeing the others. Black dashed lines illustrate the connections between computers and participants. Panel (a) shows the setting that was expected by all three participants, as they were instructed to play the game against each other. Panel (b) depicts the actual experimental settings where two independent games took place simultaneously. In the 1V2P experiment, participants 1 and 2 were connected to the same computer and virtual player 3 substituted participant 3. In the 2V1P experiment, virtual players 1 and 2 substituted participants 1 and 2 in playing with participant 3.

	C	D
C	(\$30, \$30)	(\$0, \$50)
D	(\$50, \$0)	(\$5, \$5)

(a)

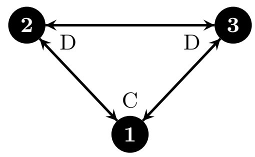

(b)

Label of the focal player	Choice of the focal player	Choice of the other players	Payoff of the focal player
1	C	D and D	\$0
2	D	C and D	\$55
3	D	C and D	\$55

(c)

FIGURE 2. Illustration of the game, through the selected payoff matrix and an example of a single round. Panel (a) illustrates the payoff matrix, whose entries constitute the payoffs two players gain by playing against each other. The first entry in the parenthesis is the reward of the player who chooses the strategy C (cooperate) or D (defect) on the rows, and the second entry is the one of the peer on the columns. Panel (b) depicts the network of interaction, where each player (circle) simultaneously plays two distinct prisoners' dilemma games, against each of the two peers. In the example, one player chooses to cooperate, and the other two to defect. Panel (c) shows the payoff for the three players, illustrating the rationale for the payoff of each of the players – each row isolates a focal player. The player who cooperated gains \$0, while the other two who defected get \$55 each.

the 1V2P and 2V1P experiments, respectively. For the 1V2P experiment, we designed virtual player 3 to behave differently in phase 1 and phase 2. At the end of phase 1, the virtual player could reach the highest cumulative score, by having always defected (winner, indicated by the acronym W), or not achieve the highest cumulative score by having mostly (at least 60%) cooperated (not a winner, indicated by NW). Analogously, in phase 2, virtual player 3 was programmed to either mostly (80%) cooperate (indicated by C) or mostly (80%) defect (indicated by D). Therefore, the four strategies were obtained by combining the two in phase 1 with the two in phase 2, that is, WC, WD, NWC, and NWD. For instance, WC identifies the strategy in which the virtual player emerged as the winner of phase 1 and mostly cooperated in phase 2.

In the 2V1P experiment, we designed virtual player 1 to emerge as the player with the highest cumulative score at the end of phase 1, while virtual player 2 was designed to systematically fall behind. In phase 2, virtual player 1 was programmed to either mostly (80%) cooperate (indicated by C) or mostly (80%) defect (indicated by D). Virtual player 2 either imitated virtual player 1, who was the winner of phase 1, (indicated by W) or chose the opposite answer to that of virtual player 1 in the previous round (indicated by NW). Overall, four strategies exemplified the difference in behavior of two virtual players in phase 2, that is, WC, WD, NWC, and

FIGURE 3. Computer interfaces created for the game. Panel (a) shows a screenshot of the interface the three participants used in phase 1, while panel (b) presents a screenshot of phase 2. On the left hand side of both panels, users were reminded to press "C" to cooperate and "D" to defect. Therein, users were also displayed the scores they would get from interacting with one of their peers. On the bottom of both panels, we presented the remaining time participants had to answer. On the right hand side of panel (a), players could see their relative cumulative score with respect to the winner, who has 100%. On the right hand side of panel (b), participants could see the choices of all the players from the previous round, their payoff from the two interactions in the previous round (orange), and their cumulative payoff since the beginning of that game (black).

NWD. For instance, WC identifies the strategy in which the virtual player 2 copied virtual player 1 who mostly cooperated.

In order to favor more instinctive decisions [43], we gave to players only four seconds to choose between cooperation and defection in each round. Further, players had five seconds while transitioning from phase 1 and phase 2, and 10 seconds between each game. A whole experimental trial lasted 20 minutes. In both experiments, virtual players' strategies were predetermined, and thus, not influenced by the participants' choices.

E. EXPERIMENTAL PROCEDURE

We recruited 48 groups of three participants from a University (144 participants in total). Each group was escorted to a private room, where the experiment took place. Before the experiment, each group took part in a tutorial session, where they were instructed about the goal of the experiment: "make as much (virtual) money as possible". To better engage the participants, the experimenter informed the players that the top six participants would be awarded with one $25 gift card each. In the tutorial session, the experimenter explained to players the game rules and payoff matrix (Fig. 2), with examples of screenshots of the interface (Fig. 3). The participants were told that they were going to be playing four games of multiple rounds each, without disclosing the exact number of rounds, as such information may favor defection [44].

After the tutorial, participants signed a consent form and proceeded to do the experiment. Since we did not disclose to participants the presence of virtual players until the end of the game, we ensured that the two experiments - one where a participant played with two virtual players, and the other where two participants played with one virtual player - started at the same time, thereby giving them an impression that they would play the same game (Fig. 1a). To avoid bias in the game, participants were provided only with their own number, displayed in the interface, without knowing others'. This setting was also adopted to reinforce the impression that they would play with each other. Each group played all four virtual player strategies in a random order, listed in Tables 1 and 2. The order of the virtual player configurations (WC, WD, NWC, and NWD) was balanced across the 48 groups. No form of communication was allowed

One virtual player with two participants (1V2P)

Virtual player 3	End of phase 1	Phase 2
WC	Highest cumulative score	Mostly cooperate
WD	Highest cumulative score	Mostly defect
NWC	Not the highest cumulative score	Mostly cooperate
NWD	Not the highest cumulative score	Mostly defect

TABLE 1. Summary of the programmed strategies for virtual player 3, who played with two participants (1V2P). The first column of each table illustrates the acronyms used for the experimental conditions. The second column shows the expected cumulative score of the virtual players at the end of phase 1. In order to reach the highest cumulative score at the end of phase 1, virtual player 3 always defected in phase 1. On the contrary, in order to ensure a cumulative score less than the highest one, virtual player 3 defected 40% of the rounds in phase 1 when its cumulative score was not the highest one; in a round where the virtual player had the highest cumulative score it automatically cooperated. The third column shows the virtual players' strategy during phase 2. When virtual player 3 mostly cooperated, it defected only 20% of the rounds, while when it mostly defected, it defected 80% of the rounds. In every experiments, the virtual players' fraction of defection were fixed and we randomly picked without replacement the rounds when the virtual players defected.

during the experiment. At the end of a game, there was a 10 second pause to ease the transition toward a new game [45].

After the participants played the four games, they filled out a survey. Other than general demographic information, we asked the participants to answer a question concerning the presence of virtual players: "Did you suspect that some of your opponents were computer programmed?". Possible answers were: (a) "Not at all", (b) "Maybe", (c) "Yes, I suspected that one of the opponent was computer programmed", and (d) "Yes, I suspected that two of the opponents were computer programmed". Willing participants also filled out personal information for competing for the six $25 gift cards.

After they filled out the survey, the experimenter informed the participants that they were not playing together, as in Fig. 1a, but against virtual players, as in Fig. 1b. The experimenter also informed the participants that four gift cards were given to the top four participants in the 1V2P experiment and two gift cards to the top two in the 2V1P experiment, since in the 1V2P experiment there were twice the participants than in the 2V1P experiment. Upon signing a debriefing form, the experiment was concluded. A video of the experiment is included in the supplementary information.

F. DATA ANALYSIS

All statistical analyses were performed with the statistics software R (version 3.6.1). The overall comparisons of the participants' fraction of defection between the conditions of interest were carried out using the Wilcoxon signed-rank test [46], with a significance level $\alpha = 0.05$. We used one-tailed tests, as we expected the following directional relationships: in condition WD participants would defect more than in condition WC; in condition NWD participants would defect more than in condition NWC; and participants who were suspecting about virtual players would defect more than participants who did not have this suspicion. All 144 participants completed the game under all four virtual player configurations summarized in Tables 1 and 2. Among them, 96 participants played in the 1V2P experiment, and 48 participants played in the 2V1P experiment.

In both 1V2P and 2V1P experiments, we focused on phase 2 and computed the participants' fraction of defection in each of the four conditions (WC, WD, NWC, and NWD). We performed a pairwise comparison between conditions WC and WD to examine whether the virtual player who emerged as the winner of phase 1 was able to increase the participants' tendency to cooperate, by opting to mostly cooperate, rather than mostly defecting. A similar analysis was performed between NWC and NWD conditions.

In the 1V2P experiment, where a pair of participants interacted with one virtual player, we performed the classification in Table 3 to ease the statistical analysis. Specifically, we compared the two participants' cumulative scores at the end of phase 2 to create two classes: winners and losers. The winner was the participant who had the higher cumulative score in the pair at the end of phase 2, and the loser was the one with the lower cumulative score. We discarded experiments where the two participants tied at the end of phase 2, yielding the following number of pairs: $n_{\mathrm{WC}} = 41$ for condition WC, $n_{\mathrm{WD}} = 42$ for condition WD, $n_{\mathrm{NWC}} = 38$ for condition NWC, and $n_{\mathrm{NWD}} = 45$ for condition NWD.

Two virtual players with one participant (2V1P)

Virtual player 1	End of phase 1	Phase 2	Virtual player 2	End of phase 1	Phase 2
WC	Highest cumulative score	Mostly cooperate	WC	Not the highest cumulative score	Imitate virtual player 1
WD	Highest cumulative score	Mostly defect	WD	Not the highest cumulative score	Imitate virtual player 1
NWC	Highest cumulative score	Mostly cooperate	NWC	Not the highest cumulative score	Differ from virtual player 1
NWD	Highest cumulative score	Mostly defect	NWD	Not the highest cumulative score	Differ from virtual player 1

TABLE 2. Summary of the programmed strategies for virtual players 1 and 2, which played with one participant (2V1P). The first column of each table illustrates the acronyms used for the experimental conditions. The second column shows the expected cumulative score of the virtual players at the end of phase 1. In order to reach the highest cumulative score at the end of phase 1, virtual player 1 always defected. On the contrary, in order to ensure a cumulative score less than the highest one, virtual player 2 defected 40% of the rounds in phase 1 when its cumulative score was not the highest one; in a round when the virtual player had the highest cumulative score it automatically cooperated. The third column shows the virtual players' strategy during phase 2. When virtual player 1 mostly cooperated, it defected only 20% of the rounds, and when it mostly defected, it defected 80% of the rounds. Virtual player 2 could either imitate virtual player 1, by deterministically copying its answer in the previous round, or differ from virtual player 1, by choosing the opposite answer to virtual player 1 in the previous round. In every experiment, the virtual players' fraction of defection were fixed, and we randomly picked without replacement the rounds when the virtual players defected.

Beyond separating participants based on their cumulative score in phase 2, we also classified 1V2P pairs on the basis of their suspicion regarding the presence of virtual players. By considering that a participant did not suspect about the virtual player only if he/she answered "Not at all" to the related question ("Did you suspect that some of your opponents were computer programmed?"), we classified pairs into not suspecting and suspecting. A pair did not suspect about the presence of virtual players if none of the participants had a suspicion. Otherwise, a group was considered to be suspecting about the presence of virtual players if at least one participant had a suspect. For 2V1P experiments, we followed an equivalent classification, based on the perception of the single participant.

To delve into the mechanisms underlying the potential influence of the virtual players on the participants' decision, we conducted an additional qualitative analysis. For the 1V2P experiment, we computed the following metrics: i) the fraction of rounds, during phase 2, in which the virtual player's cumulative score was higher than the cumulative score of the participant who was winning in the pair, and ii) the fraction of rounds, during phase 2, in which the virtual player's cumulative score was higher than the cumulative score of the participant who was losing in the pair. For the 2V1 experiment, we scored the fraction of rounds, during phase 2, in which the cumulative score of any of the two virtual players was ahead of the participant (2V1P).

In both the experiments, we studied how suspecting about the presence of a virtual player can influence the participants' decision-making process. Specifically, we analyzed the fraction of defections across the whole experiment by aggregating the four conditions. For the 1V2P, we averaged the fractions of defections of the two participants in the pair.

III. RESULTS

We first studied how participants responded (Fig. 4) to the virtual player(s)'s strategies (Tables 1 and 2). In the 1V2P experiment (where two participants were interacting with one virtual player), we registered a difference between conditions WC and WD. By classifying participants as winners or losers depending on their final cumulative score within the pair, we discovered that losers defected less in condition WC than in condition WD (Fig. 4a, Wilcoxon test, $W = 665$, $p = 0.037$). Hence, changing the behavior of the virtual player in phase 2 from mostly defecting to mostly cooperating increased the tendency of some participants to cooperate, against their own benefit. This statistically significant result was not mirrored by other differences. In fact, the virtual player did not elicit a change in the defection rate of the winners (Fig. 4a, $W = 864$, $p = 0.513$). No significant difference was found between the difference in the defection rates in NWC and NWD conditions, neither for losers (Fig. 4b, $W = 721$, $p = 0.111$) nor for winners (Fig. 4b, $W = 798$, $p = 0.290$). Likewise, for the 2V1P experiment (where one participant played with two virtual players), we failed to identify significant differences in the defection rates between conditions WC and WD (Fig. 4a, $W = 1014.5$, $p = 0.154$), and between conditions NWC and NWD (Fig. 4b, $W = 1058$,

Definition	Description
Winner	Participant who had the higher cumulative score at the end of phase 2
Loser	Participant who had the lower cumulative score at the end of phase 2
Not suspecting	A pair where none of the participants suspected about the presence of virtual players
Suspecting	A pair where at least one participant suspected about the presence of virtual players

TABLE 3. Classification implemented for analysis of the 1V2P experiment, where one virtual player interacted with two participants.

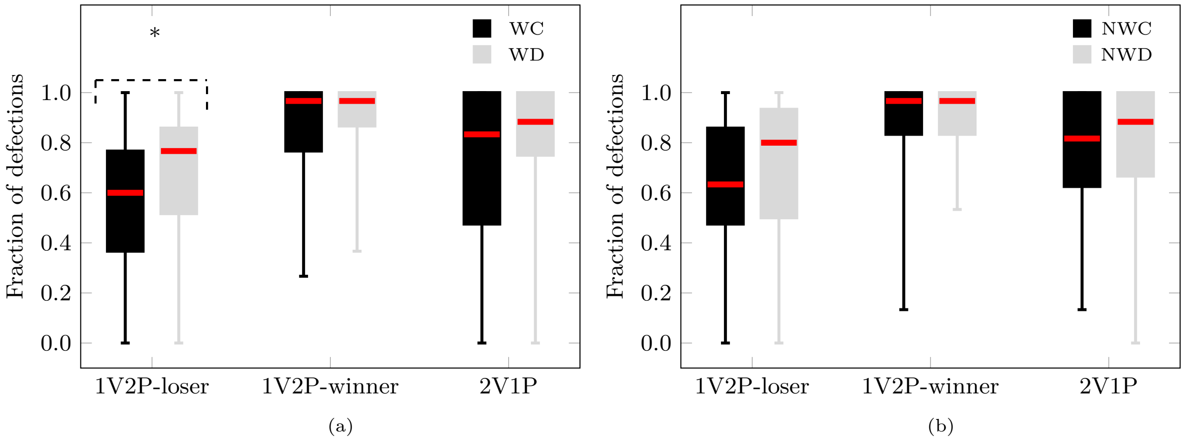

FIGURE 4. How virtual players influenced the participants' choices, measured in terms of the fraction of defections of participants during phase 2. Labels 1V2P-loser and 1V2P-winner refer to the participants who had gained the higher or lower cumulative score within the pair for the 1V2P experiment. For the 1V2P experiment, results are presented in terms of the pair, such that the fraction of defection is the average between the two participants. Panels (a) and (b) offer a comparison between condition WC and WD and NWC and NWD, respectively, for the 1V2P and 2V1P experiments. For each boxplot, the red line is the median, the box delimits the first and third quartiles, and the whiskers identify the minimum and maximum values. An asterisk indicates a statistically significant comparison.

$p = 0.243$).

To offer context into the mechanisms of decision-making by participants, we calculated the fraction of rounds that virtual players were ahead of the participants in phase 2 (Fig. 5). Predictably, in the 1V2P experiment (Fig. 5a), the virtual player 3 was typically ahead of the losers in both conditions WC and WD, but it was ahead of winners only in WD. In the 2V1P experiment (Fig. 5b), virtual player 1 was ahead the participant in both conditions WD and NWD conditions, while virtual player 2 was always behind the participant.

The average age of the participants was 22.58 ± 6.52 (mean $\pm$ standard deviation). 59 participants were female, 83 male, 1 agender, and 1 preferred not to disclose the information. The suspicion about the presence of virtual players in the game significantly influenced the participants' choices (Fig. 6). In the 1V2P experiment, a total of 28 pairs suspected about the presence of a virtual player, while 20 pairs did not suspected about it (corresponding to 38 participants suspected that there was a virtual player in the game, and 58 who did not). The pairs that suspected about the presence of a virtual player defected significantly more than those that did not have such a suspicion (Fig. 6a, $W = 159$, $p = 0.006$). In the 2V1P experiment, a total of 25 participants suspected about the virtual players, while 23 not. Participants who suspected about the presence of virtual players defected similarly to those who did not (Fig. 6b, $W = 322$, $p = 0.765$).

IV. DISCUSSION

The prisoner's dilemma rests upon a dichotomy between individual and group benefits: while the optimal strategy for an individual is to systematically defect, the

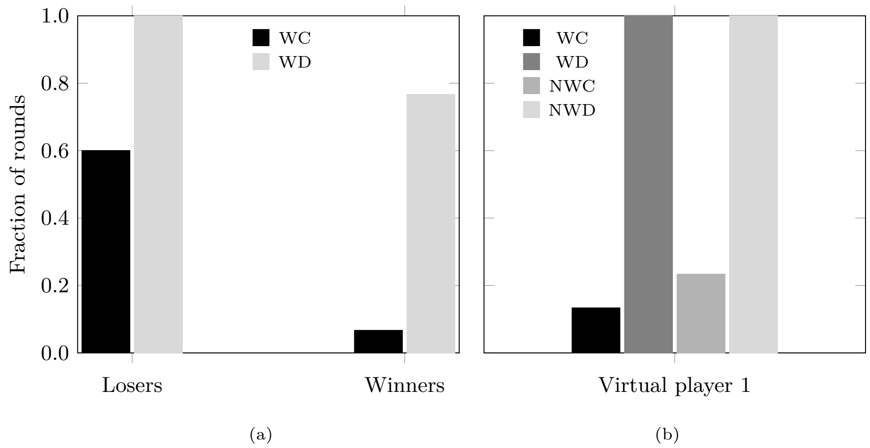

FIGURE 5. Some insight into the mechanisms underlying the influence of the virtual players, measured through the fraction of rounds when the virtual players were ahead of the participants in their cumulative scores. These quantities are scored by counting when the virtual player's cumulative score was higher than either the winning or losing participant in the pair (1V2P experiment), or any of the virtual players was ahead of the participant (2V1P). Panel (a) refers to the 1V2P experiment. Conditions NWC and NWD are not displayed as the virtual player was never ahead of the participants. Panel (b) refers to the 2V1P experiment; virtual player 2 is not displayed as it was never ahead of a participant.

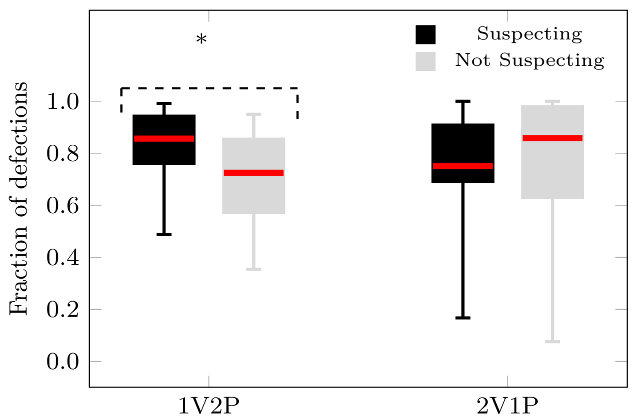

FIGURE 6. How suspecting about the presence of virtual players contributes to the fraction of defections. For the 1V2P experiment, results are presented in terms of the pair, such that the fraction of defection is the average between the two participants and the pair suspects about the presence of virtual players if at least a player does so. Results are presented using boxplots. The red line represents the median, the two ends of the box the first and third quartile, and the two whiskers the minimum and maximum values. An asterisk indicates a statistically significant comparison.

maximum benefit to the group comes from cooperation. The feasibility of modulating individual choices toward collective benefits was the chief research question of this study. Through the use of engineered virtual players, we designed an experimental study to explore two potential mechanisms for enhancing the tendency to cooperate in the iterated prisoner's dilemma game. The first mechanism focuses on the players' cumulative score to establish an association between performance and a specific strategy, which could be then imitated by the participants. The second mechanism taps into group decision-making process to form a reference strategy for the majority that could be adopted by the individual.

We formulated two main hypotheses toward contributing new insight into the role of these two mechanisms. We tested these hypotheses in a series of experiments where virtual players were deceitfully introduced in the iterated prisoners' dilemma. First, we hypothesized that the presentation of a virtual player that was at the same time the top performer and a cooperating player would incentivize participants toward cooperation. Second, we hypothesized that the presentation of multiple cooperating virtual players would promote cooperation by the participants. Our experimental results offer compelling indication in favor of the first hypothesis, but we failed to gather significant evidence regarding the second hypothesis.

Our research can be framed within the theory of social comparison [27], [47], [48], which proposes that several decision-making processes rely on the comparisons that we draw between ourselves and others. Social comparison can be downward or upward, as we compare ourselves with individuals who perform worse or better than us [49]–[51]. Both these forms of social comparison bear important consequences on individual behavior [52]–[54], whereby individuals could apprehend new valuable information as they observe the actions of others and form associations between their performance and actions [55], [56].

In our implementation of the prisoner dilemma, we created a virtual player who could "lead by example" [57]–[59] and offer participants a reference for social comparison [60]–[62]. To test the first hypothesis, we designed an experiment where the virtual player could emerge as the winner in a first phase of the interactions, by systematically defecting, and then it could change its strategy to mostly cooperate or mostly defect for the rest of the rounds. It is likely that upward social comparison was the driver of participants' choices, whereby only the participants with lower cumulative scores were influenced by the strategy of the winning virtual player. Specifically, changing the strategy of the virtual player from mostly defecting to mostly cooperating improved cooperation: participants who interacted with the cooperating virtual player had a higher tendency to cooperate than participants who played with the defecting virtual player.

Integrating with multiple virtual players at a time in the other experiment, we failed to gather support in favor of an effect of group-decision making on individual choice. On the one hand, this may be related to a deliberate strategy of the participant, such that the gain obtained by defecting against two cooperating players counterbalanced upward social comparison. Social comparison might be confounded by the presence of multiple virtual players, whereby the participant might not be able to identify which player to compare to and he/she would rather opt for a strategy that will yield personal victory, that is, defection. On the other hand, it is possible that the presence of more than one virtual player could have challenged the level of deception of the experiment, thereby prompting the participants to identify the presence of virtual players. Although we cannot dispute this second possibility, empirical evidence on the perception of the experiment by the participants does not align with this explanation. In fact, we did not register differences between the strategies of participants who suspected that virtual players were present and those who did not share this suspicion.

An effect of the perception regarding the presence of virtual players was instead detected in the other experiment, when participants were playing in pairs against a single virtual player. Specifically, we found that pairs of participants in which at least one of the two suspected about the presence of a virtual player defected more than pairs who did not have this suspicion. Likely, pairs of participants who suspected about the presence of virtual players could not socially compare [27], [47], [48] themselves with the other players, and tended to opt for the rational, self-centered choice of defecting. On the contrary, pairs of participants who did not suspect about the presence of virtual players attempted to build trust with their peers [63]–[65], fostering mutual cooperation.

The main limitation of this research pertains to the design of the virtual players, which were programmed to randomly defect, without knowledge of the choices being made by the participants. Such a one-directional interaction between virtual players and participants was likely to hamper the appraisal of the deception that was being conducted, whereby real players should be able to adjust their strategy as a function of the participants. The tit-for-tat strategy [12] may constitute an improvement over the present approach, which would, at least, afford some form of bidirectional interaction with the virtual player strategy driven by the participants' previous choices.

The two experiments presented herein provide empirical evidence about the feasibility of influencing individual choices through upward social comparison, established via the integration of virtual players. Through virtual players that "lead by example," we demonstrate the possibility of promoting strategies which support group benefits, against the individual ones. Leading by example will be effective only when the deception is successful, so that the participants would create a sense of trust toward the virtual player, pursuing actions that benefit the entire group. Although specific to the iterated prisoner dilemma, these findings may translate to other domains of investigation where decision-making can be linked to performance, from higher education to professional sports [22]–[24].

ETHIC STATEMENT

The experimental protocol was approved by the institutional review board (IRB) at New York University (IRB-FY2017-898).

ACKNOWLEDGMENT

The authors acknowledge support from the National Science Foundation under grant No. CMMI-1561134. The authors thank Ms. Elizabeth Krasner for helping in building the platform, Dr. Shinnosuke Nakayama for useful discussion, and Dr. Simone Macri for his help with the statistics.

REFERENCES

[1] P. M. Kappeler and C. P. Van Schaik, Cooperation in Primates and Humans. Springer, 2006.
[2] R. Boyd and P. J. Richerson, "Culture and the evolution of human cooperation," Philosophical Transactions of the Royal Society B: Biological Sciences, vol. 364, no. 1533, pp. 3281–3288, 2009.
[3] R. B. Myerson, Game Theory. Harvard University Press, 2013.
[4] T. Basar and G. J. Olsder, Dynamic Noncooperative Game Theory, vol. 23. Siam, 1999.
[5] J. Riegelsberger, M. A. Sasse, and J. D. McCarthy, "The researcher's dilemma: Evaluating trust in computer-mediated

communication," International Journal of Human-Computer Studies, vol. 58, no. 6, pp. 759–781, 2003.
[6] A. Rapoport, A. M. Chammah, and C. J. Orwant, Prisoner's Dilemma: A Study in Conflict and Cooperation, vol. 165. University of Michigan Press, 1965.
[7] J. F. Nash, "Equilibrium points in n-person games," Proceedings of the National Academy of Sciences, vol. 36, no. 1, pp. 48–49, 1950.
[8] R. B. Myerson, "Refinements of the Nash equilibrium concept," International Journal of Game Theory, vol. 7, no. 2, pp. 73–80, 1978.
[9] C. A. Holt and A. E. Roth, "The Nash equilibrium: A perspective," Proceedings of the National Academy of Sciences, vol. 101, no. 12, pp. 3999–4002, 2004.
[10] M. Milinski and C. Wedekind, "Working memory constrains human cooperation in the prisoner's dilemma," Proceedings of the National Academy of Sciences, vol. 95, no. 23, pp. 13755–13758, 1998.
[11] C. Gracia-Lázaro, A. Ferrer, G. Ruiz, A. Tarancón, J. A. Cuesta, A. Sánchez, and Y. Moreno, "Heterogeneous networks do not promote cooperation when humans play a prisoner's dilemma," Proceedings of the National Academy of Sciences, vol. 109, no. 32, pp. 12922–12926, 2012.
[12] R. Axelrod and W. D. Hamilton, "The evolution of cooperation," Science, vol. 211, no. 4489, pp. 1390–1396, 1981.
[13] R. Cooper, D. V. DeJong, R. Forsythe, and T. W. Ross, "Cooperation without reputation: Experimental evidence from prisoner's dilemma games," Games and Economic Behavior, vol. 12, no. 2, pp. 187–218, 1996.
[14] K. Clark and M. Sefton, "The sequential prisoner's dilemma: evidence on reciprocation," The Economic Journal, vol. 111, no. 468, pp. 51–68, 2001.
[15] B. Majolo, K. Ames, R. Brumpton, R. Garratt, K. Hall, N. Wilson, et al., "Human friendship favours cooperation in the iterated prisoner's dilemma," Behaviour, vol. 143, no. 11, pp. 1383–1395, 2006.
[16] J. C. Waddell and W. Peng, "Does it matter with whom you slay? The effects of competition, cooperation and relationship type among video game players," Computers in Human Behavior, vol. 38, pp. 331–338, 2014.
[17] W. Su, M. N. Potenza, Z. Zhang, X. Hu, L. Gao, and Y. Wang, "Do individuals with problematic and non-problematic internet game use differ in cooperative behaviors with partners of different social distances? Evidence from the prisoner's dilemma and chicken game," Computers in Human Behavior, vol. 87, pp. 363–370, 2018.
[18] J. Grujić, C. Fosco, L. Araujo, J. A. Cuesta, and A. Sánchez, "Social experiments in the mesoscale: Humans playing a spatial prisoner's dilemma," PloS One, vol. 5, no. 11, p. e13749, 2010.
[19] R. Axelrod, "Effective choice in the prisoner's dilemma," Journal of Conflict Resolution, vol. 24, no. 1, pp. 3–25, 1980.
[20] M. A. Nowak, "Five rules for the evolution of cooperation," Science, vol. 314, no. 5805, pp. 1560–1563, 2006.
[21] J. Qin, Y. Chen, W. Fu, Y. Kang, and M. Perc, "Neighborhood diversity promotes cooperation in social dilemmas," IEEE Access, vol. 6, pp. 5003–5009, 2017.
[22] H. M. Ngan, G. P. Prendergast, and A. S. Tsang, "Linking sports sponsorship with purchase intentions: Team performance, stars, and the moderating role of team identification," European Journal of Marketing, vol. 45, no. 4, pp. 551–566, 2011.
[23] O. Bissoon, "Can better institutions attract more foreign direct investment (FDI)? Evidence from developing countries," in International Conference on Applied Economics, pp. 59–70, 2011.
[24] J. Reiman and P. Leighton, Rich Get Richer and the Poor Get Prison: Ideology, Class, and Criminal Justice. Routledge, 2015.
[25] M. G. Zimmermann and V. M. Eguíluz, "Cooperation, social networks, and the emergence of leadership in a prisoner's dilemma with adaptive local interactions," Physical Review E, vol. 72, no. 5, p. 056118, 2005.
[26] L. Wheeler, "Motivation as a determinant of upward comparison," Journal of Experimental Social Psychology, vol. 1, pp. 27–31, 1966.
[27] L. Festinger, "A theory of social comparison processes," Human Relations, vol. 7, no. 2, pp. 117–140, 1954.
[28] E. A. Vogel, J. P. Rose, L. R. Roberts, and K. Eckles, "Social comparison, social media, and self-esteem," Psychology of Popular Media Culture, vol. 3, no. 4, p. 206, 2014.
[29] J. Andreoni and J. H. Miller, "Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence," The Economic Journal, vol. 103, no. 418, pp. 570–585, 1993.
[30] S. Kiesler, L. Sproull, and K. Waters, "A prisoner's dilemma experiment on cooperation with people and human-like computers.," Journal of Personality and Social Psychology, vol. 70, no. 1, p. 47, 1996.
[31] C. Straßmann, A. M. Rosenthal-von der Pütten, and N. C. Krämer, "With or against each other? The influence of a virtual agent's (non) cooperative behavior on user's cooperation behavior in the prisoners' dilemma," Advances in Human-Computer Interaction, vol. 2018, 2018.
[32] F. Ishowo-Oloko, J.-F. Bonnefon, Z. Soroye, J. Crandall, I. Rahwan, and T. Rahwan, "Behavioural evidence for a transparency–efficiency tradeoff in human–machine cooperation," Nature Machine Intelligence, vol. 1, no. 11, pp. 517–521, 2019.
[33] E. B. Sandoval, J. Brandstetter, M. Obaid, and C. Bartneck, "Reciprocity in human-robot interaction: A quantitative approach through the prisoner's dilemma and the ultimatum game," International Journal of Social Robotics, vol. 8, no. 2, pp. 303–317, 2016.
[34] J. Laut, F. Cappa, O. Nov, and M. Porfiri, "Increasing citizen science contribution using a virtual peer," Journal of the Association for Information Science and Technology, vol. 68, no. 3, pp. 583–593, 2017.
[35] R. B. Ventura, S. Richmond, M. Nadini, S. Nakayama, and M. Porfiri, "Does winning or losing change players' engagement in competitive games? Experiments in virtual reality," IEEE Transactions on Games, 2019.
[36] L. Green, P. C. Price, and M. E. Hamburger, "Prisoner's dilemma and the pigeon: Control by immediate consequences," Journal of the Experimental Analysis of Behavior, vol. 64, no. 1, pp. 1–17, 1995.
[37] S. Duca and H. H. Nax, "Groups and scores: The decline of cooperation," Journal of The Royal Society Interface, vol. 15, no. 144, p. 20180158, 2018.
[38] P. Kollock, "Transforming social dilemmas: Group identity and co-operation," Modeling Rationality, Morality, and Evolution, vol. 7, pp. 185–209, 1998.
[39] P. Bonacich, G. H. Shure, J. P. Kahan, and R. J. Meeker, "Cooperation and group size in the n-person prisoners' dilemma," Journal of Conflict Resolution, vol. 20, no. 4, pp. 687–706, 1976.
[40] R. M. Dawes, "Social dilemmas," Annual Review of Psychology, vol. 31, no. 1, pp. 169–193, 1980.
[41] Q. Chen, T. Chen, and R. Yang, "Using rewards reasonably: The effects of stratified-rewards in public goods game," Chaos, Solitons & Fractals, vol. 120, pp. 67–74, 2019.
[42] S. Nakayama, E. Krasner, L. Zino, and M. Porfiri, "Social information and spontaneous emergence of leaders in human groups," Journal of the Royal Society Interface, vol. 16, no. 151, p. 20180938, 2019.
[43] C. Dietrich, "Decision making: Factors that influence decision making, heuristics used, and decision outcomes," Inquiries Journal, vol. 2, no. 02, 2010.
[44] L. E. Fouraker and S. Siegel, Bargaining Behavior. 1963.
[45] N. Koudenburg, T. Postmes, and E. H. Gordijn, "Disrupting the flow: How brief silences in group conversations affect social needs," Journal of Experimental Social Psychology, vol. 47, no. 2, pp. 512–515, 2011.
[46] S. Sidney, "Nonparametric statistics for the behavioral sciences," The Journal of Nervous and Mental Disease, vol. 125, no. 3, p. 497, 1957.
[47] J. Suls and L. Wheeler, Handbook of Social Comparison: Theory and Research. Springer Science & Business Media, 2013.
[48] J. E. Suls and T. A. E. Wills, Social Comparison: Contemporary Theory and Research. Lawrence Erlbaum Associates, Inc, 1991.
[49] B. Major, M. Testa, and W. H. Blysma, "Responses to upward and downward social comparisons," in Social Comparison: Contemporary Theory and Research, pp. 237—-260, Lawrence Erlbaum Associates, Inc, 1991.

[50] S. E. Taylor and M. Lobel, "Social comparison activity under threat: Downward evaluation and upward contacts.," Psychological Review, vol. 96, no. 4, p. 569, 1989.
[51] R. H. Smith, "Assimilative and contrastive emotional reactions to upward and downward social comparisons," in Handbook of Social Comparison, pp. 173–200, Springer, 2000.
[52] F. X. Gibbons and M. Gerrard, "Predicting young adults' health risk behavior," Journal of Personality and Social Psychology, vol. 69, no. 3, p. 505, 1995.
[53] M. Gerrard, "Health images and their effects on health behavior," Health, Coping, and Well-being: Perspectives from Social Comparison Theory, vol. 63, 2013.
[54] J. Suls and L. Wheeler, "A selective history of classic and neo-social comparison theory," in Handbook of Social Comparison, pp. 3–19, Springer, 2000.
[55] S. E. Asch, "Opinions and social pressure," Scientific American, vol. 193, no. 5, pp. 31–35, 1955.
[56] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal, "Meet, discuss, and segregate!," Complexity, vol. 7, no. 3, pp. 55–63, 2002.
[57] D. Goleman, "What makes a leader," Organizational Influence Processes, vol. 82, pp. 229–241, 2003.
[58] D. Arce M, "Leadership and the aggregation of international collective action," Oxford Economic Papers, vol. 53, no. 1, pp. 114–137, 2001.
[59] L. Rendell, R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland, "Why copy others? Insights from the social learning strategies tournament," Science, vol. 328, no. 5975, pp. 208–213, 2010.
[60] W. de Waal-Andrews and I. Van Beest, "A sweeter win: When others help us outperform them," Journal of Experimental Social Psychology, vol. 74, pp. 218–230, 2018.
[61] P.-y. Nie, C. Wang, and T. Cui, "Players acting as leaders in turn improve cooperation," Royal Society Open Science, vol. 6, no. 7, p. 190251, 2019.
[62] J. G. Reiter, C. Hilbe, D. G. Rand, K. Chatterjee, and M. A. Nowak, "Crosstalk in concurrent repeated games impedes direct reciprocity and requires stronger levels of forgiveness," Nature Communications, vol. 9, no. 1, pp. 1–8, 2018.
[63] J. Hirshleifer, "Competition, cooperation, and conflict in economics and biology," The American Economic Review, vol. 68, no. 2, pp. 238–243, 1978.
[64] J. J. Argo, K. White, and D. W. Dahl, "Social comparison theory and deception in the interpersonal exchange of consumption information," Journal of Consumer Research, vol. 33, no. 1, pp. 99–108, 2006.
[65] B. M. Zagorsky, J. G. Reiter, K. Chatterjee, and M. A. Nowak, "Forgiver triumphs in alternating prisoner's dilemma," PLoS One, vol. 8, no. 12, 2013.

MATTHIEU NADINI received B.Sc. degree in Physics from University of Modena and Reggio Emilia in 2015 (cum laude), M2R in Physics of Complex Systems from Paris Diderot University in 2017, and M.Sc. in Physics of Complex Systems from Polytechnic University of Turin in 2017 (cum laude). Presently, he is a Ph.D. student in the Dynamical Systems Laboratory at New York University Tandon School of Engineering. His research interests include network science, epidemiology, computational social science, and human behavior.

PEERAYOS PONGSACHAI is currently finishing B.Sc. degree in Mechanical Engineering with a minor in Electrical Engineering at New York University Tandon School of Engineering. He worked as a Research Assistant at the Dynamical Systems Laboratory in Summer of 2019.

CHIARA SPINELLO received B.Sc. degree in Biology and M.Sc. degree in Neurobiology from "Sapienza" University of Rome in 2015 and 2017, respectively. She is currently a full time Assistant Researcher and a Laboratory Manager of the Dynamical Systems Laboratory at New York University Tandon School of Engineering. She is the author of ten publications in the field of behavioral neuroscience.

DANIEL ALBERTO BURBANO-L. (M'12) received B.Sc. degree in Electronic Engineering and M.Sc. degree in Industrial Automation from the National University of Colombia, Manizales, Colombia, in 2010 and 2012, respectively, and Ph.D. degree in Control Engineering from the University of Naples Federico II, Italy. From 2017 till 2018, he was a Post-Doctoral Researcher at Northwestern University, USA, working on network topology identification. He is currently an Assistant Professor/Faculty Fellow at New York University. His research is focused on the analysis and control of complex systems and networks with particular attention to problems in network estimation and control, animal behavior, consensus, and synchronization.

MAURIZIO PORFIRI (SM'12–F'19) received M.Sc. and Ph.D. degrees in Engineering Mechanics from Virginia Tech, in 2000 and 2006; a "Laurea" in Electrical Engineering (with honours) and a Ph.D. in Theoretical and Applied Mechanics from "Sapienza" University of Rome and the University of Toulon (dual degree program between Italy and France), in 2001 and 2005, respectively. He is currently an Institute Professor at New York University Tandon School of Engineering, with appointments in the Departments of Mechanical and Aerospace Engineering, Biomedical Engineering, and Civil and Urban Engineering. Dr. Porfiri is a Fellow of the American Society of Mechanical Engineers (ASME) and the Institute of Electrical and Electronic Engineers (IEEE). He has served in the Editorial Board of ASME Journal of Dynamics Systems, Measurements and Control, ASME Journal of Vibrations and Acoustics, IEEE Control Systems Letters, IEEE Transactions on Circuits and Systems I, and Mechatronics. Dr. Porfiri is the Director of the Dynamical Systems Laboratory and he is engaged in conducting and supervising research on dynamical systems theory, multiphysics modeling, and underwater robotics. He is the author of more than 300 journal publications and the recipient of the National Science Foundation CAREER award. He has been included in the "Brilliant 10" list of Popular Science in 2010 and his research featured in all the major media outlets, including CNN, NPR, Scientific American, and Discovery Channel. Other significant recognitions include invitations to the Frontiers of Engineering Symposium and the Japan-America Frontiers of Engineering Symposium organized by National Academy of Engineering; the Outstanding Young Alumnus award by the college of Engineering of Virginia Tech; the ASME Gary Anderson Early Achievement Award; the ASME DSCD Young Investigator Award; and the ASME C.D. Mote, Jr. Early Career Award.