Impact of Trust in the Spatial and Evolutionary Iterated Prisoner's Dilemma

Tommaso Tofacchi[1] , ***under supervision of*** **Neil Yorke-Smith**[1] ,
[1]Department of Electrical Engineering, Mathematics and Computer Science, Technische Universiteit Delft, Netherlands

Abstract

The complex dynamics that shape human interactions have been the object of studies for decades. One of the most famous examples to analyse the cooperative or defective nature of people is the Iterated Prisoner's Dilemma (IPD), a versatile dynamic game which has been shown capable of capturing the emergence of cooperative behaviours in environments where people are rationally tempted to defect. Trust plays a crucial role in such social interactions, as it intrinsically influences a human on which behaviour to adopt with a partner: a key aspect that the Prisoner's Dilemma (PD) game fails to register. In this study, two variants of the traditional PD game are proposed in the context of spatial and evolutionary Iterated Prisoner's Dilemma; each variant features the inclusion of trust as a variable that affects the rewards of the two-player game. The goal of this paper is to study the impact that trust has on the efficacy of traditional Prisoner's Dilemma strategies played by agents in simulated environments. The results obtained over different experiments confirm that trust indeed fosters cooperating behaviours among agents, and allows them to more easily populate worlds that feature harsh living conditions.

1 Introduction

Game Theory is the branch of studies that focuses on observing and understanding the dynamics of interaction among economic agents, which make choices according to personal utilities and therefore generate outcomes [1]. Due to the framework's ease of use, Game Theory has often been used to model social interactions, in order to better understand via simulation the behaviours of humans when facing such situations.

Among many, the *Prisoner's Dilemma* (PD) is arguably the most renowned game in literature. The game consists of two players – namely, two prisoners – who must choose whether to cooperate or defect with each other [2]. The players choose autonomously and with no possibility to communicate with their partner, then finally receive a payoff according to their choices. The payoffs are calculated as in Table 1 according to the inequalities 1 and 2, which ensure non-exploitation of mutual defection at alternating turns to obtain maximum rewards. Due to the payoff distribution, rational agents will always choose defection, as it provides the best personal payoff regardless of the opponent's choice; however, both players can receive higher rewards if they both decide to cooperate.

Table 1: Payoff table for the PD game; C_i corresponds to a player i choosing cooperation, D_i for defection. In a (X, Y) pair of rewards, reward X goes to player A and reward Y to player B.

		Player B	
		C_B	D_B
Player A	C_A	R, R	S, T
	D_A	T, S	P, P

$$T > R > P > S \tag{1}$$
$$2R > S + T \tag{2}$$

Such cooperative behaviour is shown to emerge rather easily among the agents as they play repeated games of the Prisoner's Dilemma in succession, i.e., when the PD is *iterated* (IPD) and thus transitions from being static to dynamic. Later studies have confirmed the emergence of cooperative behaviours also in IPD games played in spatial and natural environments [3, 4], where agents face costs of living in a simulated world at every turn and have to pay them with their earned payoffs (the energy).

Despite enabling the study of cooperation among agents, the traditional IPD game fails to address another fundamental component of human interactions that lays behind cooperation itself: *trust*. Whereas cooperation indicates acting together with a partner to obtain mutual rewards, trust digs deeper into the reasons why cooperation happens. In fact, trust is the willingness to take risks under uncertainty [5, 6]; the action of an individual of revealing some own vulnerabilities to a partner, in order to receive a greater payoff according to their actions. Trusting a partner, however, consequentially means also accepting the risk that in case of missed cooperation, the defected individual could be severely damaged – proportionally to how much it trusted the partner.

In 2005, Yamagishi et al. proposed an empirical study on different variants of the traditional Prisoner's Dilemma, which aimed to address the role of trust in the game by separating the actions of trust and cooperation [7]; for such, these variations are grouped under the name of *Prisoner's Dilemma with variable Dependence* (PD/D). Yamagishi et al. conducted their study over mixed groups of American and Japanese people playing non-spatial *iterated PD/D* (IPD/D) games, and reported two relevant findings: firstly, the importance of separating trust from cooperation in order to build trust relationships; secondly, the causation of cooperation towards trust – and not the other way around.

In my research, I utilize the conclusions of Yamagishi et al. to build simulated environments that reproduce the human trust behaviour, thus allowing for investigating into the influence of trust over IPD/D games played in spatial and evolutionary environments by computer agents that employ the most diffused strategies in literature. In particular, the two principal variations of the game proposed by Yamagishi et al. are implemented, so that it can be assessed whether trust can foster or diminish the evolution of cooperation among simulated agents in such realistic environments.

I hypothesize that the harsher the environments, the more the separation of trust boosts cooperation. Additionally, the explicit presence of trust as a variable in the IPD/D games leads to more defecting behaviours when the society is wealthier (thus living conditions are less competitive).

A background of the research topic can be found in Section 2, and Section 2.1 presents a summary of the approaches to trust in the Prisoner's Dilemma game found in literature. The two PD/D variations are explained in Section 2.2, and their implementations are detailed in Section 4.

In Section 3, the works from van Tilburg [4] and Gevers [8] are presented, as they provided benchmark studies to confront the influence of trust in different realistic spatial and evolutionary IPD environments. The strategies used in their two studies and this one are explained in Section 5.

To discuss the hypotheses, a series of experiments in different environments - presented in Section 6 – will be run and measured over the metrics of population's and actions' distribution; the effects of trust on the evolution of cooperation can thus be analyzed, as that is reported in Section 8. In order to transparently include validation of results, replication and reproducibility for the reader, Section 7 addresses the theme of responsible research.

Finally, conclusions and suggestions for future work in the field are provided in Section 9.

2 Background

The Prisoner's Dilemma game has been central in extensive studies over the years. As its iterated version offers an interesting framework to study the evolution of cooperation, different approaches to the game have been developed by researchers depending on the goals of their works.

The primary distinction pertains to the kind of agents that are subject to observation in studies; behaviour analysis of players can be conducted both on humans playing the game, or computer agents which play in simulated environments according to various sets of strategies. In this research, the focus is shifted on the latter category. The Iterated Prisoner's Dilemma game can then be extended – as is the case of study here – over spatial and evolutionary environments [9, 10].

The spatial IPD sees agents play Prisoner's Dilemma rounds in a two- or three-dimensional simulated world, where they can move across to reach other players and partner up. This attempt to recreate spatial playing conditions is often enhanced in order to more closely reproduce natural environments (e.g., by introducing costs for basic agent's needs and actions), thus enabling researchers to investigate how strategies perform in realistic situations.

The evolutionary IPD, on the other hand, integrates evolutionary mechanics in the game, for instance by allowing agents to reproduce and introducing Darwinian rules that define how their strategies can mutate and adapt to become a more optimal fit for the environment. Further information on both extensions of the IPD game is provided in Section 4, where the models utilized in this research are explained.

2.1 Trust in Previous Research

The role of trust in the Prisoner's Dilemma game has been subject to different studies in the past; the approaches followed to investigate its influence in PD games have been diverse, thus hereby a summary of the most relevant results is reported.

In 2009, Venanzi proposed an interpretation of the human-like trust characteristic in simulated agents systems [11]. It shapes as a series of intuitive and internally similar rating systems for agents, in which the commonly shared idea features *trustworthiness* as an agent's parameter that is calculated via the ratio of cooperations and total actions.

Similarly, Chen et. al in 2011 studied the impact of an agent's reputation score calculated, once again, by considering the ratio between cooperations and total actions [12, 13]; it differentiates from previous approaches by introducing the possibility to receive indirect information from external agents and weighting its influence according to the senders' own reputation.

Trust explicated as such became thus used as the *niceness* indicator for an agent or a strategy, due to its intrinsic correlation with the act of cooperation in the game – as previously explored. All of the mentioned approaches focus on the *passive* trustworthiness of an individual as a reputation indicator that summarizes an agent's actions. Yamagishi et al. [7] relevance in the field comes from their contribution, introduced in Section 1: the authors shift focus to the *active* role that trust plays in human relationships, by directly targeting how it can alter the players' behaviour over the course of repeated games. To achieve so, they modified the calculation of payoffs in the two proposed Prisoner's Dilemma game variants, so that the action of trusting a partner can alter them dynamically. In the next subsection, the two PD/D games are presented to the reader.

2.2 Prisoner's Dilemma with Variable Dependence

The first game proposed by Yamagishi et. al is the *Prisoner's Dilemma with Trust via Matrix Changes* (PD/Dm), in which the payoff matrix of each player increases or decreases according to whether a player decides to trust or not its partner.

If a player decides to trust their partner, then all of the values in the player's payoff matrix are increased (or decreased, in case of negative payoffs) by a set percentage of the initial values. Conversely, in case of mistrust, the payoffs are reduced (or increased) by the same set percentage of the initial values.

The second game is the *Prisoner's Dilemma with Trust via Coin Entrustment* (PD/Dc), an investment game [14] where every player decides how much of its previously earned rewards entrusting its partner with, thus shaping the payoff distribution. When two players are paired, a PD/Dc game is carried out as follows:

1. Players simultaneously entrust their partner with a certain percentage of their reward points earned until that moment.
2. Players simultaneously decide on an independent action to perform – either cooperate or defect.
3. Players receive the appropriate payoffs for their actions as rewards.

In case of cooperation, a player decides to send the received points back to the owner. It thus receives no rewards from this action, and the owner in return receives double the amount of points that it had initially entrusted. In case of defection, an agent retains the amount of points received from the partner. The player is thus rewarded by stealing the points that the partner decided to entrust it with, and the partner does not receive anything back.

Both games share the feature of isolating trust from cooperation, whilst retaining the key PD feature of rewarding players with payoffs according to their choice of cooperating or defecting.

3 Related Work

In this paper, the influence that introducing trust in the spatial and evolutionary Iterated Prisoner's Dilemma can bring to the agents' behaviours is analysed by performance comparison with the previously obtained results in the field. In particular, the data collected via *agent-based modelling and simulation* (ABMS) [15] – thoroughly explained in Section 4 – is confronted with the findings of van Tilburg and Gevers on the performance of IPD strategies run in natural and noisy environments.

In 2018, van Tilburg proposed a model [4] to analyse the behaviour of the ten most diffused strategies in IPD research [16]. In the model, a number of agents are deployed in a two-dimensional grid which simulates a world environment; at every time tick, agents can play a Prisoner's Dilemma game with one of their nearby partners and gain energy points according to their mutual choice of actions (see Table 1). If an agent earns sufficient energy points, it can generate offspring that inherits its key traits in evolutionary fashion (e.g., the strategy to decide which move to play next). Agents can also die by old age or insufficient energy (i.e., when their energy points fall below 0), since at every time tick a *cost of living* is deducted from each player. With this model, van Tilburg showed that cooperative strategies are more likely to survive and reproduce in harsh environments, thus the ones defined by high costs of living, high reproduction costs and low life expectancy.

An extension to the study was proposed in 2020 by Gevers, who focused on the performance of IPD strategies when introducing noise in harsh spatial environments [8]. To simulate the agents' behaviours, Gevers used an approach similar to van Tilburg's and expanded his model by adding a noise parameter and additional, environment-appropriate strategies. The *noise parameter* is a value from 0 to 1 which represents the probability that an agent's action gets inadvertently flipped before it reaches the partner. The added strategies are generous and contrite variants of the *Tit-for-Tat* and *Pavlov* traditional strategies [17]; generosity is described as "leaving a certain amount of another individual's defections unpunished", while contrition is "avoiding to defect as a response to the other individual's defection after own unintended defection" [8, p.3]. Gevers confirmed that van Tilburg's claims over cooperative strategies still hold when noise is involved in the harshness of an environment, therefore generous strategies help in coping with noise; moreover, he demonstrated that in presence of high uncertainty (i.e., when noise values are high) contrite behaviours are the most effective.

To study the impact of trust in such scenarios, this paper reproduces the works from van Tilburg and Gevers, and extends them by implementing the two Prisoner's Dilemma with Variable Dependence game variants by Yamagishi et. al [7] described in Section 2.2. In this way, simulations can be conducted over the same environments (i.e., with the same world settings) for both the traditional Prisoner's Dilemma game and its trust-isolating variants. The results can thus be gathered unambiguously and confronted clearly over the same metrics, enabling to specifically address the influence that trust has in the spatial and evolutionary Iterated Prisoner's Dilemma.

4 Models

To simulate experimental runs and gather data over the influence of trust in the spatial and evolutionary IPD, multiple *agent-based modeling and simulation systems* have been generated on the NetLogo framework. The models are based on van Tilburg's codebase[1], and are four in total: two models for the PD/Dm game and two models for the PD/Dc game. Each of the two models per different game features the specifics and the strategies employed respectively by van Tilburg and Gevers in their works.

All the models share key features on how they represent the spatial and evolutionary IPD game. An $L \times L$ discrete grid is populated with a variable number of agents, each occupying a single, non-overlapping cell and employing a specific strategy; a thorough list of the utilised strategies is provided in Section 5. Agents possess an energy level, which is drawn from a 1 to 50 uniform distribution at simulation setup and signals the agent's death when it falls below 0. Time is measured discretely in ticks; at every tick, a *cost of living* amount of energy is subtracted from every agents' energy

[1] A personal acknowledgement to van Tilburg, who permitted the use and share of his code in my work.

points. During the tick, agents can thus perform a series of actions described as follows:

- Pair up with a Moore neighbour (another agent that occupies one of their 8 adjacent cells), and play a Prisoner's Dilemma game together to earn energy points.
- Move to an adjacent cell, if no neighbour to play a game with is found.
- Reproduce by the cost of 50 energy points and, on one of the free adjacent cells, hatch an agent that has 50 energy points and the same strategy as the parent.
- Evolve their strategy to match one of the better performing neighbours. For such, an *evolution threshold* is indicated, so that if the energy level of an agent + the evolution threshold is $\leq$ than the neighbour's energy level, the agent adopts the neighbour's strategy. If *evolution threshold* $= 0$, then evolution is not applied.

In the following subsections, the differences in models for the PD/Dm and PD/Dc variants are provided.

4.1 PD/Dm

In the Prisoner's Dilemma with Trust via Matrix Changes, the core difference with the traditional Prisoner's Dilemma game resides in the fact that the payoff matrix changes according to the trust placed by a player in its partner.

To model this in an ABMS system, a trust mechanism has been introduced according to the conclusions of Yamagishi et al. As noted by the authors, "cooperation leads to trust, not the other way around" [7, p.303]: for such, agents store a *trust array* in which every cell contains a number that summarizes the results of past cooperation-defection received by a specific, previously-played player. For every time that a partner cooperates, the number – initially 0 – is incremented by 1; for every defection, it is instead decreased by 1.

An *increase trust factor*, a number between 0 and 1, indicates the percent increase or decrease in payoffs (relative to the initial payoff matrix) that an agent experiences per action when playing a PD/Dm game with a partner. In order to calculate the appropriate variation in the payoff matrix, the partner's corresponding cell is fetched from the *trust array* and multiplied by the *increase trust* factor (see Formula 3).

As an example: if *increase trust* $= 0.1$ and an agent partners up with another agent who in previous encounters cooperated 5 times and defected 2, its corresponding cell in the *trust array* will contain the value 3, and the agent's payoff matrix will increase (or decrease, in case of negative payoffs) by $0.1 \times 3 = 30\%$ with respect to the initial payoff matrix's values.

$$\mathrm{P_{increase}} = \text{increase-T} \times \text{T-array}_i[j] \times \mathrm{P_{initial}} \tag{3}$$

The payoff increase in agent i's payoff matrix when playing a PD/Dm game with agent j. P = *Payoffs*, T = *Trust*.

4.2 PD/Dc

In the Prisoner's Dilemma with Trust via Coin Entrustment, changes in the dynamics of the game are more substantial: there is no payoff matrix involved in the decision of the payoffs as a game's outcome.

To replicate the trust dynamic in this scenario (explained in Section 2.2), at hatch every agent is assigned a *base trust* value, which is a number from 0 to 1 that indicates an agent's inclination towards trusting an opponent. If an agent is the child of another agent, its *base trust* value is randomly picked in a 30% interval ($\pm 15\%$) around the parent's *base trust*. Every agent stores a *trust array* in the same fashion as in Section 4.1; an *increase trust* value from 0 to 1 is present as well, albeit with a slightly different use. In fact, in the PD/Dc game agents decide on the amount of energy to entrust their opponent with according to Formula 4.

$$\mathrm{E_{to\ entrust}} = \mathrm{E}_i \times (\mathrm{T_{base}} + \text{T-array}_i[j] \times \text{increase-T}) \tag{4}$$

The calculation of energy that agent i entrusts agent j with when playing a PD/Dc game. E = *Energy*, T = *Trust*.

4.3 Van Tilburg's and Gevers' Differences

The codebase for the simulation suite on NetLogo was provided by van Tilburg, thus it is entirely consistent with the implementation in his research; the codebase was then extended by implementing the two PD/D variants of the game.

Gevers' model was re-implemented in NetLogo by extending van Tilburg's: a *noise* parameter was added coherently to his studies, thus as a number from 0 to 1 indicating the probability that an agent's action gets flipped due to miscommunication in a noisy environment.

5 Strategies

Van Tilburg's and Gevers' approaches deal with two different subjects in the field of the spatial IPD, therefore the authors analysed different sets of agents' strategies.

To ensure correctness and fidelity to the original studies when comparing them to the results obtained with trust-based variants of the PD games, every strategy has been provided in the corresponding NetLogo models.

5.1 Van Tilburg's Strategies

Van Tilburg studied the performance of the 9 strategies classified as *default type* by Jurišić et al. [16]; the authors conducted an analysis of the strategies proposed in the three major IPD tournaments held from 2004 to 2005 and listed the 9 *default types* as follows (see Table 2). As the Pavlov strategy can be played with either cooperation or defection as the initial behaviour, the resulting total number of strategies for van Tilburg is 10.

5.2 Gevers' Strategies

Gevers also studied the performance of the same 9 *default type* strategies. However, Gevers dealt with harsher environments than van Tilburg's, characterized by noise that could alter the choice of an agent's action. He thus considered 5 additional strategies to the *default types*, which correspond to generous and contrite variants of the Tit-for-Tat (TFT) and Pavlov strategies [17]. Gevers' added strategies are described in Table 3.

Table 2: Strategies used in Van Tilburg's research [4].

Designation	Description
ALLC	A strategy that always plays cooperation.
ALLD	A strategy that always plays defection.
RAND	A strategy has a 50% probability of playing either cooperation or defection.
GRIM	A strategy that starts with cooperation; after its opponent's first defection, it continues with defection.
TFT	A strategy that starts with cooperation, then replicates the moves of the opponent.
STFT	As TFT, but the strategy starts with defection.
TFTT	As TFT, but the strategy defects after two consecutive defections.
TTFT	As TFT, but each opponent's defection is retaliated with two defections.
Pavlov	A strategy that considers action results as divided into 2 groups: positive actions (T and R), and negative actions (P and S). If the result of the previous action belonged to the first group, the action is repeated; if the result was in the second group, then the action is changed. It is also called *win-stay, lose shift*.

Table 3: Strategies used in Gevers' research [8].

Designation	Description
GTFT	As TFT, but it has a 10% probability of cooperating when it would otherwise defect.
CTFT	As TFT, but it has 3 states: *contrite*, *content* and *provoked*. The strategy begins with cooperation and stays there until there is an unilateral defection. If the agent was the victim while *content*, it becomes *provoked* and defects until a cooperation from the other player causes it to return *content*. If the agent was a defector while *content*, it becomes *contrite* and cooperates. When *contrite*, it only becomes content after successfully cooperating.
GPavlov	As Pavlov, but it has a 10% probability of cooperating when it would otherwise defect.
SGTFT	As TFT, but it has a 60% probability of cooperating when it would otherwise defect.
SGPavlov	As Pavlov, but it has a 60% probability of cooperating when it would otherwise defect.

6 Results

To investigate the influence that trust can have in the spatial and evolutioanry IPD, a set of simulations with different environments was analyzed. More specifically, in this section the various scenarios are organized as:

- IPD/Dm simulations
  - In natural and evolutionary environments (based on van Tilburg's study).
  - In noisy environments (based on Gevers' study).
- IPD/Dc simulations
  - In natural and evolutionary environments.

All the experiments have been run with $L = 100$ (thus 100×100 grid size), an initial *population size* of 10% of the available grid cells (thus 1000 agents), with strategies uniformly distributed over the players. Every experiment features its own set of *cost of living* (or K), *increase trust*, *payoff matrix values* and *noise* (where applicable). For every experiment, the number of executed runs and their criterion of termination will be reported; when referring to *until stabilization*, it is to be intended as "every strategy's agent-count stayed in a margin of 0.1% of the total population for 50 rounds".

6.1 IPD/Dm Simulations

The experiments involving the PD/Dm variation of the Prisoner's Dilemma game are reported here. For the experiments in natural and evolutionary environments, as they are based on van Tilburg's studies, the strategies from his article are used (see Section 5.1); for the experiments in noisy environments, as they are based on Gevers' studies, the best performing noise-specific strategies found in his article are used, alongside the remaining traditional ones (thus ALLC, ALLD, RAND, GRIM, STFT, TFTT, TTFT, CTFT and SGPavlov).

Natural and Evolutionary Environments

For all the PD/Dm experiments conducted in natural and evolutionary environments, the payoff values are set as $T = 4, R = 2, P = 0$ and $S = -1$.

The first simulation compares the percentage of strategies among agents over the first 350 ticks in van Tilburg's research to the results obtained by substituting the PD game with its PD/Dm variant, under the same circumstances and over 20

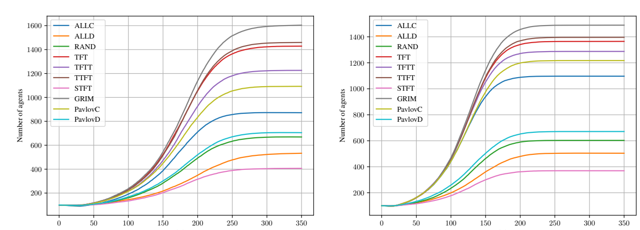

Figure 1: Agents count per strategy over the first 350 ticks.

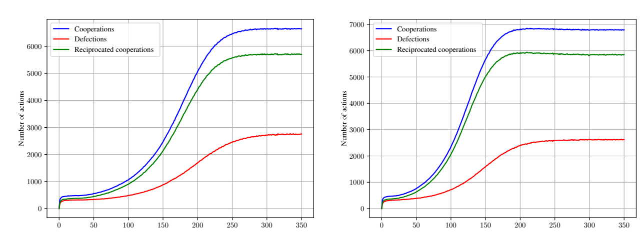

Figure 2: Actions count over the first 350 ticks.

runs (see Figure 1). An added metric is included, and that is the number of actions happening in the simulation at every tick (see Figure 2). For this experiment, *cost of living* $= 0.25$, the *evolution threshold* is set at 0, and *increase trust* $= 0.1$ in the PD/Dm variant.

The second experiment compares the variation in numbers of actions that agents perform at the last tick of runs when the PD/Dm game is played with different values of *increase trust* (see Figure 3). Specifically, for the *increase trust* parameter, 6 values from 0.0 (corresponding to the traditional Prisoner's Dilemma game) to 0.5 in 0.1 increases are considered, as well as both scenarios with *evolution threshold* set at 0 and 20. The *cost of living* is set at 0.25. Once again, the obtained results are averaged over 20 runs, each of which terminated at population stabilization.

The third experiment focuses on analyzing the impact of varying the *increase trust* parameter over different *cost of living* values. For such, the time ticks that on a 10-run average every simulation took to achieve a fully populated grid are reported in Table 4 for *evolution threshold* $= 0$, and in Table 5 for *evolution threshold* $= 20$. *N.A.* indicates that the runs were not able to reach 100% grid population in those simulations.

Noisy Environments

The experiments ran in noisy environments focus on investigating the impact of trust over collaboration in the setting proposed by Gevers; for this reason, the *evolution threshold* has been set to 0 and never changed during the analysis. Moreover, the same payoffs from Gevers' study have been used, thus $T = 5$, $R = 3$, $P = 0$ and $S = -2$. To differentiate with previous experiments in non-noisy environments, the *noise* parameter has been set to 0.1.

The first experiment compares the variation in numbers of actions that agents perform at the last tick of runs, when the PD/Dm game is played with varying *increase trust* at different values of *cost of living* (see Table 6 for cooperations, Table 7 for reciprocated cooperation, Table 8 for defections). Runs are ended when population reaches stabilization, or after 2000 ticks in case of impossibility to stabilize; 10 runs have been executed per every parameter combination.

In the second experiment, the graphs of actions over ticks are presented to the reader, with the aim to investigate the evolution of cooperation/defection over different *cost of living* values with *increase trust* $= 0$ and 0.3 (see Figures 4 and 5). The last tick before stabilization/tick limit (here set at 2500, with 10 runs per value combination) indicates the maximum scale over the *x-axis*, thus the differences over graphs.

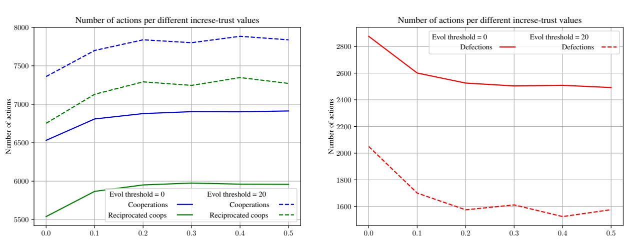

Figure 3: Actions count at termination over various *increase trust*.

Table 4: Ticks to full grid with *evolution threshold* $= 0$.

	Increase trust					
Cost of living	0.0	0.1	0.2	0.3	0.4	0.5
0.00	289	237	220	215	208	210
0.25	398	278	243	254	241	221
0.50	458	305	307	290	269	279
0.75	585	393	334	322	330	335
1.00	748	466	409	442	393	390
1.50	2002	621	534	547	721	552
2.00	N.A.	905	812	755	641	651
2.50	N.A.	1690	1054	1023	987	948
3.00	N.A.	N.A.	N.A.	3529	2614	2463
3.50	N.A.	N.A.	N.A.	N.A.	N.A.	N.A.
4.00	N.A.	N.A.	N.A.	N.A.	N.A.	N.A.

Table 5: Ticks to full grid with *evolution threshold* $= 20$.

	Increase trust					
Cost of living	0.0	0.1	0.2	0.3	0.4	0.5
0.00	283	232	220	212	216	207
0.25	352	281	262	247	244	230
0.50	448	317	288	287	292	276
0.75	548	383	325	317	343	300
1.00	758	447	396	352	357	350
1.50	1899	581	522	470	470	475
2.00	N.A.	908	737	639	662	636
2.50	N.A.	1624	1155	1084	904	948
3.00	N.A.	N.A.	N.A.	3614	2850	2488
3.50	N.A.	N.A.	N.A.	N.A.	N.A.	N.A.
4.00	N.A.	N.A.	N.A.	N.A.	N.A.	N.A.

Table 6: Cooperations at the end of simulations.

	Increase trust					
Cost of living	0.0	0.1	0.2	0.3	0.4	0.5
0	3529	3765	3859	4061	4138	4152
1	4668	4280	5286	5593	5696	5919
2	2820	3929	4209	4970	5982	5953
3	40	3900	4215	4408	4624	5742

Table 7: Reciprocated cooperations at the end of simulations.

	Increase trust					
Cost of living	0.0	0.1	0.2	0.3	0.4	0.5
0	2581	2832	2933	3110	3174	3174
1	3662	3350	4084	4367	4426	4619
2	2516	3435	3598	4167	4963	4903
3	33	3519	3774	3912	4056	4967

Table 8: Defections at the end of simulations.

	Increase trust					
Cost of living	0.0	0.1	0.2	0.3	0.4	0.5
0	2469	2298	2255	2324	2327	2370
1	2302	2137	2809	2821	2924	2990
2	506	910	1195	1671	2153	2289
3	15	656	775	908	1069	1501

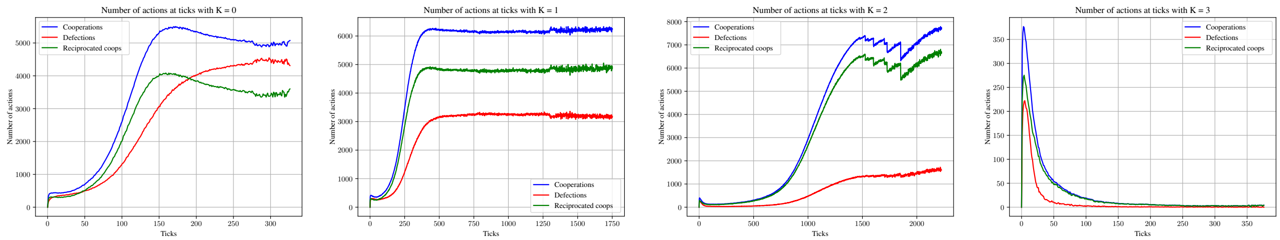

Figure 4: Actions over ticks for *cost of living* $= 0, 1, 2, 3$ at *increase trust* $= 0$.

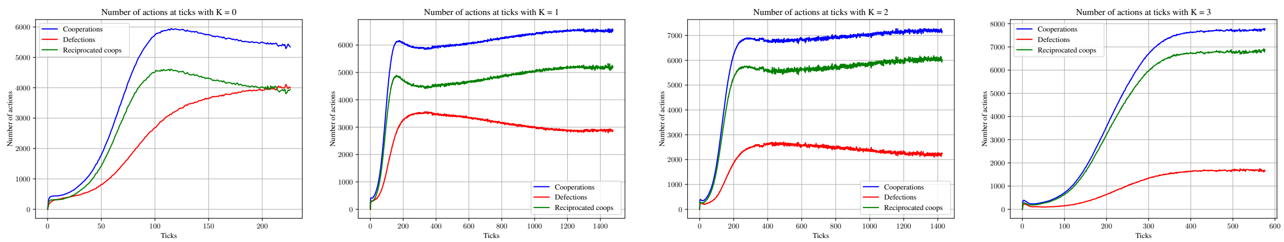

Figure 5: Actions over ticks for *cost of living* $= 0, 1, 2, 3$ at *increase trust* $= 0.3$.

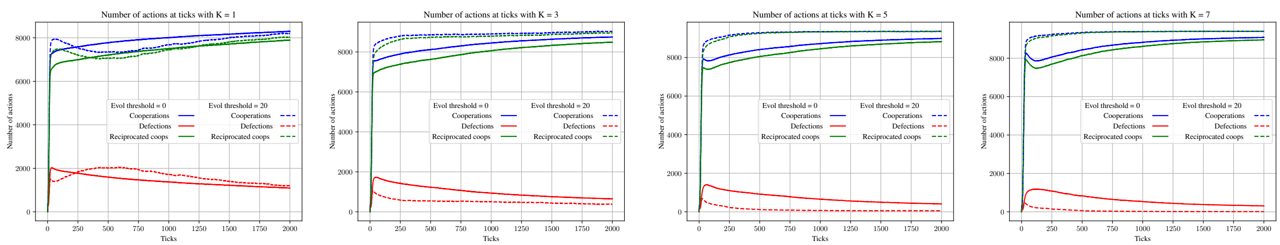

Figure 6: Actions over ticks for *cost of living* $= 1, 3, 5, 7$.

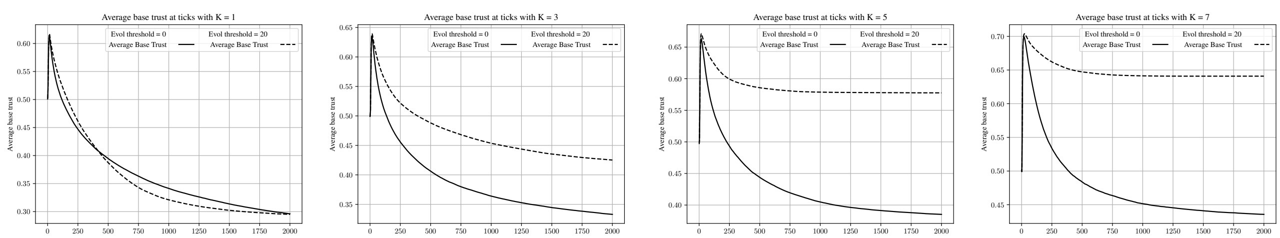

Figure 7: Average *base trust* values over ticks for *cost of living* $= 1, 3, 5, 7$.

	Strategy									
Cost of living	ALLD	TFTT	ALLC	PavlovC	TFT	TTFT	GRIM	STFT	PavlovD	RAND
1	17	1702	455	1458	2071	1667	1621	416	266	328
3	9	1331	448	1404	2444	2033	1828	199	100	205
5	4	1064	474	1374	2403	2278	2118	96	88	99
7	5	886	444	1354	2673	2106	2292	85	91	63

Table 9: Number of agents per strategy at end of simulations with *evolution threshold* $= 0$.

	Strategy									
Cost of living	ALLD	TFTT	ALLC	PavlovC	TFT	TTFT	GRIM	STFT	PavlovD	RAND
1	714	221	131	742	3063	2730	2398	0	0	0
3	135	568	345	1058	2782	2577	2534	0	0	0
5	0	1113	768	1530	2196	2240	2152	0	0	0
7	0	1305	1050	1565	1946	1913	2221	0	0	0

Table 10: Number of agents per strategy at end of simulations with *evolution threshold* $= 20$.

6.2 IPD/Dc Simulations

This subsection presents the simulations that feature the PD/Dc variant of the Prisoner's Dilemma game. For the different nature of the game (more diverse than PD/Dm vs PD), no direct comparison with previous studies is possible. However, the settings that define the environments hereby used resemble the natural and evolutionary scenarios of van Tilburg. Due to this, the strategies featured in these experiments are the ones used by van Tilburg in his studies (see Section 5.1).

Natural and Evolutionary Environments

The first experiment presents the graphs of actions over ticks with the aim to investigate the evolution of cooperation over four different *cost of living* values $= 1$, 3, 5 and 7 (see Figure 6), as well as the graphs of the agents' average *base trust* values over ticks (see Figure 7). To study the impact of evolution, the *evolution threshold* is set at 0 and 20. Results are obtained as the average of 10 runs per parameter combination, each of which terminating after 2000 ticks.

In the second experiment, we observe the strategies' development among agents at the end of simulations with *cost of living* values $= 1$, 3, 5 and 7 (see Tables 9 and 10). To study the impact of evolution, the *evolution threshold* is set at 0 and 20. Results are obtained as the average of 10 runs per parameter, each of which terminating after 2000 ticks.

7 Responsible Research

In this section, insight is provided over the reproducibility of the results obtained in the research.

The codebase used in the research and the results obtained by running the simulations are available online upon request for validation, replication, reproduction or extension. The reader can download the models and run them in NetLogo, either to recreate the results or to analyse different scenarios by modifying the parameters introduced in Section 4.

As described in Section 6, every experimental case has been run 10 or 20 times – according to the type of simulation –, in order to provide consistent averaged results. Due to the stochastic nature of the ABMS systems involved (it being due to the grid setup or some strategies' own randomness), it cannot be expected to replicate consistently the same exact results at every iteration of the simulations. However, the *coefficient of variation* (or CV) metric have been used for every experiment to analyze the results' consistency over averaged runs (see Formula 5).

$$CV = \frac{\text{standard deviation}}{|\text{mean}|} \tag{5}$$

In the IPD/Dm simulations run in natural and evolutionary environments, the *CV* mean value for the collected data in each experiment is between 0.03 to 0.12.

In IPD/Dm simulations run with added *noise*, max *CV* can rise up to 0.65 in the first experiment: this is due to both the added randomness of the *noise* parameter and the inclusion of generous/contrite variants of traditional strategies, that rely on probabilities. Low-number results (e.g., the ones obtained in high *costs of living* scenarios) also negatively impact the score, as being subject to higher relative variation impacts the *CV* value, but have a low influence on the outcome of the simulation if the rest of the reported data is considerably larger. The second experiment registers a mean *CV* value of 0.03.

In the final IPD/Dc simulations, the mean $CV = 0.37$ for both the evolutionary and non-evolutionary cases in the first experiment; in the second experiment, the value can grow up to 0.72 for the evolutionary scenario (and 0.64 for the non-evolutionary one). This for similar numerical reasons as in the noisy simulations, whilst the generally higher mean values can depend on the different nature of the PD/Dc game.

In every scenario, the mean *CV* value is well below 1, which is considered to be a good indicator of low variance in collected data, especially for stochastic simulations.

8 Discussion

In the following section, an overall interpretation of every simulation case is provided to the reader; every experiment is summarized in its results, highlighting the key findings and observations. The reported behaviours are then related to the hypotheses advanced in Section 1, to finally assess the knowledge that can be deduced from studying the impact of human trust in the agent-simulated spatial and evolutionary Iterated Prisoner's Dilemma.

8.1 Summary of Results

The results are analyzed per category of experiments, as listed in Section 6.

IPD/Dm Simulations in Natural and Evolutionary Environments

The first experiment focuses on observing how the different strategies spread across the grid's population with and without the trusting mechanic (in Figure 1); the number of actions over time ticks is also reported, to study the evolution of cooperation (in Figure 2). The strategies distribution remains overall consistent over both the with and without trust scenarios, as the only difference noticed in the with-trust simulation is the wider gap that separates the best 6 strategies from the worst-performing 4. When analyzing the ticks that the agents need to stabilize, however, it is reported that the simulation with trust takes $\approx$ 230 ticks to reach full-grid population, whilst the non-trust simulation takes ≈ 300 ticks: a 23% performance improvement. The same behaviour can be seen in the actions graph, where the amount of both total and reciprocated cooperations is also slightly larger in the trusted setting than the non-trusted one (+142 and +147 respectively).

The second experiment deals with the impact that different *increase trust* values have on the number of actions happening when simulation reaches stabilization (in Figure 3). The scenarios analyzed include both values of 0 and 20 for the *evolution threshold* parameter, to observe how evolution is affected by trust; these values were chosen as they are the critical points to notice a consistent impact of evolution, as observed by van Tilburg. When *evolution threshold* $= 0$, the resulting action curves show a distinct increase in cooperation (therefore decrease in defection) with trust increasing from 0 to 0.3, then a stabilized behaviour for the remaining values; when *evolution threshold* $= 20$, the same pattern repeats, albeit with more fluctuation for values from 0.3 to 0.5. The effects of evolution in mitigating defecting behaviours studied by van Tilburg finds confirmation in these results.

The third experiment compares the time in ticks that simulations with different combinations of *increase trust* and *cost of living* values take to reach 100% grid occupation (see Tables 4 and 5). Again, values of 0 and 20 are considered for the *evolution threshold*. It appears clear that introducing higher degrees of trust in the experiments significantly improves efficiency in populating the grid (thus in agents' reproduction, therefore in agents' capability of earning higher rewards). Moreover, trust allows agents to achieve a fully-populated grid at increasingly larger *costs of living*, where the impact of trust is the greatest.

IPD/Dm Simulations in Noisy Environments

The first experiment proposes an analysis on the number of actions observed at the last tick of the simulations in a noisy environment, thus when *noise* $= 0.1$ (see Tables 6, 7, and 8). The trend of trust's influence is the same as the one noticed in van Tilburg-based experiments for less harsh environments (thus with lower *costs of living*); however, it is interesting to see how in such environments the gap between the number of cooperations and defections is drastically lower than what previously noticed. For the harsher environment at *cost of living* $= 2$, however, the increase in defections significantly outgrows the increase in reciprocated cooperations and total cooperations (at *increase trust* $= 0.5$, $+352\%$ defections and just $+95\%$ in reciprocated cooperations, $+111\%$ in cooperations). The ratio then normalizes again in favour of cooperations for *cost of living* $= 3$.

The second experiment focuses on the evolution of cooperation over time for the first experiment's settings where *increase trust* $= 0$ and 0.3, chosen as a representative value to investigate the influence of trust (see Figures 4 and 5). The graphs show in a different manner the effects that trust has on mitigating defection, as well on increasing the survival and reproduction abilities of agents in harsh conditions (see the graphs for *cost of living* $= 3$).

IPD/Dc Simulations in Natural and Evolutionary Environments

The first experiment provides the action graphs for the IPD/Dc game played in environments with *cost of living* $= 1$, 3, 5 and 7 (see Figure 6), at the two *evolution threshold* of 0 and 20. Moreover, the graphs showing the agents' average base trust values over time ticks are reported in Figure 7. The graphs for *evolution threshold* $= 20$ reveal an interesting behaviour at *cost of living* $= 1$: the defecting actions appear to close the gap with cooperations, before reverting to the regularly observed behaviour of decline. This is not observed for larger values of *cost of living*; however, for *evolution threshold* $= 0$ the same attempt at "exploiting the cooperators" is noticed in an inverse pattern, thus increasing as the *costs of living* increase. The graphs concerning the average base trust level of the agents open to different, insightful observations. Firstly, the values for *evolution threshold* $= 0$ drop significantly below the values registered for *evolution threshold* $= 20$ when the *costs of living* increase, as the respective curve gradient descends more rapidly towards 0. Secondly, although the peaks reached by the average base levels of trust remain unchanged in the graphs, it is interesting to notice how their floor values increase as *costs of living* get larger.

The second experiment is a confrontation of strategies diffusion after 2500 ticks, with different values of *cost of living* and *evolution threshold* $= 0$, 20 (see Tables 9 and 10). Here, the difference between the *evolution thresholds* is significant. For *evolution threshold* $= 20$, we notice a complete annihilation of the RAND, PavlovD and STFT strategies regardless of the *cost of living*; for *evolution threshold* $= 0$, that is not the case, as all strategies survive in any scenario. The ALLD strategy performs particularly well for *cost of living* $= 1$ and 2 in the evolutionary scenario, whilst it is not able to reproduce effectively in the non-evolutionary one. Strategies as TFT and TTFT see the largest discrepancies with respect to different *evolution thresholds*: in the non-evolutionary environment they significantly grow in number as *costs of living* increase, but for *evolution threshold* $= 20$ the opposite happens.

8.2 Hypotheses Discussion

In Section 1, two hypotheses over the influence of trust in the evolutionary and spatial IPD were advanced. Here, the same hypotheses are discussed in light of the findings elicited from the results. The two hypotheses develop as follows:

1. The harsher the environments, the more the separation of trust boosts cooperation.
2. The separation of trust leads to more defecting behaviours when the society is wealthier (thus living conditions are less competitive).

To discuss the hypotheses, it is important to verify if trust indeed contributes to boosting cooperation in the overall picture, and eventually how it does so.

In the IPD/Dm simulations in natural and evolutionary environments, trust fosters cooperating behaviours in different manners:

- In the first experiment, defections are slightly reduced and cooperations benefit from that. Moreover, the time needed to reach a stable population is noticeably reduced.
- In the second experiment, higher values of *increase trust* lead to more cooperative agents' behaviours.
- In the third experiment, higher values of *increase trust* allow agents to populate the grid far more efficiently and have a significantly larger impact in harsher environments. The actions' graphs are not included in this article for space constraints, but they conform to the behaviours registered in the first and second experiments.

In the IPD/Dm simulations in noisy environments, trust fosters cooperating behaviours similarly to what seen previously:

- In the first experiment, higher values of *increase trust* favour the emergence of cooperative behaviours. It appears also that defective behaviours grow at a larger rate than cooperations as the *costs of living* increase, although it is due to the fact that defectors can survive more easily by exploiting a larger presence of cooperating neighbours.

- In the second experiment, a boost to cooperation provided by the introduction of trust is evident. In particular, its effects are more evident in the reduced time ticks that are needed to reach a stable development pattern as the *costs of living* increase.

Finally, in the IPD/Dc simulations in natural and evolutionary environments, it is observed that trust fosters cooperating behaviours as follows:

- In the first experiment, trust increases cooperative actions and decreases defections (compared to the traditional PD scenario, but also the PD/Dm variant). Moreover, as the living conditions get harsher, the surviving agents show greater levels of *base trust* values, thus adapting their needs of trusting a partner according to the increasing *costs of living*.
- In the second experiment, trust appears to behave differently in case of *evolution threshold* = 0 or 20 for what concerns the population distribution. In particular, in the first case the population distribution gap between the best performing strategies and the rest becomes more accentuated as the *cost of living* increases; in the second case, conversely, that gap is reduced among the top 6 out of 10 strategies, with the remaining 4 that are reduced to extinction.

With respect to the first hypothesis, it can be stated that trust fosters cooperation more in harsher environments. Although the second experiment with the IPD/Dm played in noisy environments seems to demonstrate that the impact of trust in cooperation increase is greater at lower *cost of living* values, all of the other IPD/Dm and IPD/Dc experiments show the opposite. In particular, it is interesting to notice how the effects of trust on cooperation affect significantly the time that it takes for agents to reproduce and occupy larger areas of the grid, while at the same time allowing to populate worlds with harsher living conditions. This exact behaviour is also reported in the noisy experiments, that provide even harsher scenarios due to the added probability of actions miscommunication.

The second hypothesis does not appear to generally hold in the IPD/Dm simulations: trust seems to mitigate the diffusion of defectors in any scenario. With trust, however, the defectors grow in number when the environments are harsher: this shows the opportunistic behaviour hinted in the hypothesis, as defectors are able to proliferate more when surrounded by more cooperators to exploit. Moreover, in the IPD/Dc simulations (particularly, in the first experiment) an interesting behaviour can be noted for *evolution threshold* = 0: in all but one instance (*cost of living* = 3), the initial pattern of cooperations outgrowing in numbers the defections is briefly inverted for a few ticks span, as a result of the defectors' attempt to exploit the cooperators when sufficiently wealthy. This tendency, however, is rapidly inverted, as afterwards defectors begin to decline and cooperators rise again.

9 Conclusions and Future Work

This research aims to investigate how the introduction of trust in the spatial and evolutionary Iterated Prisoner's Dilemma game affects the performance of simulated agents that employ the most studied strategies in literature. To achieve so, the trusting mechanics have been implemented in ABMS systems according to the findings of Yamagishi et al. [7], who investigated the correlation between trust and cooperation among humans playing two iterated variations of the traditional Prisoner's Dilemma game. The two previous works by van Tilburg [4] and Gevers [8] provide interesting benchmarking scenarios for natural and noisy environments, thus have been used as a reference to confront the variation in performance of the different strategies when introducing trust.

The simulations indeed show that the human feature of trust helps in facilitating the emergence of collaboration among agents. Players who can decide whether trusting or not their partner – and consequentially modify their payoffs – are able to achieve higher rewards more rapidly, meaning that they can reproduce more efficiently and populate a given spatial environment in less time. Additionally, the harsher the living conditions of an environment, the more trust is useful to agents to survive, and can allow them to colonize environments where, otherwise, the *cost of living* would prevent them to do so.

The effects of allowing agents to evolve their strategies and adopt the one of a better performing neighbour are comparable to the observations of van Tilburg in his paper: they foster cooperation, thus amplifying some effects of the introduction of trust. However, when the evolution of strategies is applied to the PD/Dc game, a curious outcome is noticed: as the *cost of living* increases, evolution balances the distribution of the top-6 strategies among the agents – an occurrence that does not happen when the evolution of strategies is not applied.

The noisy environments provided by Gevers provide insight to notice the effects of trust in even harsher environments, where actions can be subject to miscommunication errors. As a result, defectors are more active in these scenarios, thus the ability of trust to mitigate defective behaviours over time can be appreciated also under more difficult circumstances.

In this research, only sets of strategies proposed by [16] and [17] are analyzed to report changes in cooperation and defection. An extension of the study that includes genetic strategies (such as the ones proposed by [18] and [19]) could be of interest, in particular if focusing more on the kind of strategies that develop optimally in presence of trust, instead of the evolution of cooperation through actions. Finally, trust is only one of the human features that play a role when interacting with a partner; in order to more realistically study human behaviour through the games of Game Theory in simulated environments, a different set of influencing features could be modelled and analyzed.

References

[1] Don Ross. Game Theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2019 edition, 2019.

[2] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.

[3] Michael Doebeli and Christoph Hauert. Models of cooperation based on the prisoner's dilemma and the snowdrift game. *Ecology Letters*, 8(7):748–766, 2005.

[4] Jasper van Tilburg. *Performance of Strategies for the Iterated Prisoner's Dilemma in a Natural Environment*, 2018.

[5] The Anh Han, Cedric Perret, and Simon T. Powers. When to (or not to) trust intelligent machines: Insights from an evolutionary game theory analysis of trust in repeated games. *Cognitive Systems Research*, 68:111–124, 2021.

[6] Niklas Luhmann, Christian Morgner, and Michael King. *Trust and power*. Polity Press, 2018.

[7] Toshio Yamagishi, Satoshi Kanazawa, Rie Mashima, and Shigeru Terai. Separating trust from cooperation in a dynamic relationship: Prisoner's dilemma with variable dependence. *Rationality and Society*, 17(3):275–308, 2005.

[8] Louis Gevers. *Effects of Noise on Cooperation in Harsh Environments*, 2020.

[9] Paul E. Smaldino. Cooperation in harsh environments and the emergence of spatial patterns. *Chaos, Solitons & Fractals*, 56:6–12, 2013. Collective Behavior and Evolutionary Games.

[10] Martin A. Nowak and Robert M. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, Oct 1992.

[11] Matteo Venanzi. *Trust, Kinship and Locality in the Iterated Prisoner's*, 2009.

[12] Bo Chen, Bin Zhang, and Weidong Zhu. Combined trust model based on evidence theory in iterated prisoner's dilemma game. *Intern. J. Syst. Sci.*, 42(1):63–80, January 2011.

[13] Bo Chen, Bin Zhang, and Hua qing Wu. Misreporting behaviour in iterated prisoner's dilemma game with combined trust strategy. *International Journal of Systems Science*, 46(1):31–43, 2015.

[14] Joyce Berg, John Dickhaut, and Kevin McCabe. Trust, Reciprocity, and Social History. *Games and Economic Behavior*, 10(1):122–142, July 1995.

[15] Charles M. Macal and Michael J. North. Agent-based modeling and simulation. In *Proceedings of the 2009 Winter Simulation Conference (WSC)*, pages 86–98, 2009.

[16] Marko Jurišić, Dragutin Kermek, and Mladen Konecki. A review of iterated prisoner's dilemma strategies. In *2012 Proceedings of the 35th International Convention MIPRO*, pages 1093–1097, 2012.

[17] Jianzhong Wu and Robert Axelrod. How to cope with noise in the iterated prisoner's dilemma. *The Journal of Conflict Resolution*, 39(1):183–189, 1995.

[18] Andrew Errity. *Evolving Strategies for the Prisoner's Dilemma*, 2003.

[19] Adnan Haider. *Using Genetic Algorithms to develop strategies for the Prisoner's Dilemma*, 2005.