Effects of Noise on Cooperation in Harsh Environments

Louis Gevers[a], *under supervision of* Neil Yorke-Smith[a],

[a] *Faculty of Electrical Engineering, Mathematics & Computer Science, Delft University of Technology, Mekelweg 5, 2628 CD Delft, Netherlands.*

Abstract

Interactions in the real world are subject to mistakes and miscommunications. The presence of this noise in interactions challenges cooperation, as one party cannot determine whether the other party did not cooperate on purpose. The Prisoner's Dilemma has commonly been used to study mutual cooperation. Strategies like Tit for Tat that do well in the classic version of the game, perform badly once noise is present. Recent studies that have used the Prisoner's Dilemma to show that harsh environments promote cooperation, currently do not take noise into account. This article therefore uses the Prisoner's Dilemma and Agent Based Modeling and Simulation (ABMS) to describe the relation between the harshness of the environment and noise. From the simulations it follows that the adversity of the environment benefits cooperators and can make cooperation more robust against mistakes. Harsher environments also encourage greater generosity to cope with noise. Yet when uncertainty is high due to higher probability of mistakes or more potential defectors in the environment, contrite behaviours are most successful.

Keywords: Prisoner's Dilemma, noise, harsh environments, generosity, contrition

1. Introduction

An individual's behaviour is considered cooperative if it provides benefits to another individual, and has been naturally selected because of these benefits [1]. The abundant presence of cooperation in the biological world leads to a paradox in evolutionary theory; providing benefits to another would reduce the relative fitness of this behaviour and therefore be selected against [2]. Extensive research has however shown that under the right conditions cooperation can evolve in nature [3, 4, 5]. In 1981, Axelrod and Hamilton formalized the issue of cooperation using the Prisoner's Dilemma Game to study the exact conditions under which cooperative behaviour can evolve [6].

When environmental adversity is high, cooperation in many species counter-intuitively increases [7]. The adversity, or harshness, of an environment increases with the short-term benefits of selfish behaviour. Examples are increased risk of predation and scarcity of resources, where selfish behaviour would be more advantageous short-term. However, in both situations cooperative behaviour in nature has been observed to increase ([8] and [9], respectively).

The same outcome can be observed in spatial Prisoner's Dilemma Games with simulated harshness. Defecting strategies benefit the most at first, but in the long run cooperating strategies recover and later dominate the game [10]. This effect is enhanced when harshness increases; cooperative strategies shrink more at first, and make a more dramatic recovery in the long run.

However, these experiments have been conducted with only strictly cooperating and strictly defecting strategies, assuming that no mistakes were ever made. Other works have explored a larger set of strategies [11], but with the same assumption that no mistakes were made in the decision making. Yet information in real-world interactions is not perfect and errors occur. Strategies that perform very well in a normal Prisoner's Dilemma setting, often fail miserably when occasional mistakes happen [12]. For a more realistic understanding of cooperative behaviours in harsh environments, mistakes (more commonly known as noise) have to be simulated as well.

While various works have studied the effects of noise in spatial versions of the game [13, 14, 15], no literature has been found on how noise affects successful strategies in harsh environments. It has been observed that cooperative, although non-forgiving strategies are more successful in harsh environments [11], but it is unknown if these strategies can remain successful in the presence of noise.

This article explores the relation between noise and harshness of the environment. Using Agent Based Modeling and Simulation (ABMS) [16] on a spatial Prisoner's Dilemma Game with simulated harshness [10], the effects of noise in harsh environments are studied. The results show that the presence of mistakes makes cooperation more difficult, but the harshness of the environment can in some cases help cooperators' robustness against noise. As opposed to previous conclusions that the least generous cooperators are the most successful in harsh environments [11], this article shows how once noise is present highly generous strategies are dramatically more successful than other strategies. Results also show that not only generous but also contrite behaviours are a successful way to account for mistakes in harsh environments.

Section 2 provides relevant background and definitions used in this article. Section 3 describes how the relations between noise and harshness will be studied using agent-based simulations, followed by the exact model description in Section 4. The obtained results using this model are presented in Section 5. Section 6 discusses the reproducibility of the results. Results are interpreted in Section 7 followed by a conclusion and recommendations for future work in Section 8.

2. Background

2.1. Prisoner's Dilemma Game

The Prisoner's Dilemma Game formalizes the problem of achieving mutual cooperation [17]. In the classic version of the game, two players each decide to either cooperate or defect. The individual choice to defect has a higher payoff than to cooperate. But if both parties defect, the payoff is worse than if they both had decided to cooperate. Formally, the Prisoner's Dilemma is represented by matrix with 8 entries, shown in Table 1. The letters represent the payoffs of each player. R stands for *reward*; the payoff when both parties cooperate. T stands for *temptation*; the higher payoff for successfully defecting (i.e. the other player tried to cooperate). S stands for *sucker's payoff*; the payoff when cooperating when the other player defects. P stands for *punishment*; the payoff if both players defect. The Prisoner's Dilemma includes the following inequalities in its definition [17]:

$$S < P < R < T \tag{1}$$

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

These inequalities are what makes the game a dilemma. Condition 1 motivates to defect in order to get the highest temptation T, at the risk of getting punishment P. Although the punishment P is higher than the sucker's payoff S, the player is motivated to cooperate to get the higher reward R by cooperating. This only works if the other player cooperates as well, hence the dilemma. Condition 2 guarantees the only "cooperative solution" is when both players cooperate. Suppose $S + T \geq 2R$, then both players that alternate between cooperating and defecting would receive a reward of $S + T$ for each alternation, i.e. at least $2R$, which would have been obtained if both players had consecutively cooperated.

2.2. Iterated Prisoner's Dilemma

The version of the game where players play multiple rounds against each other, is the Iterated Prisoner's Dilemma (IPD). The same conditions 1 and 2 introduced in the previous section hold for this version. Behaviour can be recognized and decisions

	C_2	D_2
C_1	R_1, R_2	S_1, T_2
D_1	T_1, S_2	P_1, P_2

Table 1: Matrix representation of the Prisoner's Dilemma. C_x and D_x stand for player x cooperating or defecting, respectively. Entries Y_x represent payoff Y for player x.

adapted, which leads to a much larger set of interesting strategies. For example, if playing against a player that seems to consistently defect, the rational move is to defect as well, even if the original intent was to cooperate.

In Axelrod's and Hamilton's work a computer tournament was organized to study which strategies are successful in the Iterated Prisoner's Dilemma [6]. The best performing strategy was "TIT for TAT". This simple strategy cooperates on the first move, and then copies whatever the player did on the last move. It is therefore a strategy based on cooperation and reciprocity. Other strategies based on cooperation performed significantly better in the tournament than selfish strategies.

2.3. Noise

Noise defines the probability an individual makes a mistake in an IPD game. For example, with noise of 5%, each intended choice has a probability of 5% to result in the opposite choice being carried out. In the classic IPD, successful strategies remain cooperative in the presence of noise, but either posses more generosity or contrition [12]. Generosity is 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.

2.4. Harshness of the Environment

Numerous extensions have been made to the game in order to simulate a more "realistic" environment. Attempts at making an environment more realistic generally consist of adding parameters that mimic real biological settings to influence the game. For example, implementing natural selection and evolution [18], creating a 2 dimensional space in which players can move and play against their neighbours (spatial IPD) [19], allowing players to choose and refuse partners [20], etc.. One particularly interesting model is the one presented by Smaldino et al., a spatial IPD that simulates the harshness of an environment [10].

Their model simulates the adversity of the environment using energy levels and two different parameters [10]. The main idea behind the model is to simulate harshness by removing more energy from the agents. This can be done by increasing the cost of life, i.e. removing more energy from every agent at each turn; or making the sucker's payoff increasingly more negative to punish cooperations with defectors. Adding death and reproduction mechanisms makes the environment competitive, allowing the more successful strategies to survive. The formal model description can be found in Section 4.

3. Agent-Based Modeling and Simulation for Spatial IPD

The advances in computing power over the past decades have made modeling and simulation important means for analysis and experimentation in many different domains [21]. A model is a representation of an original or reference system of interest. A simulation is the operation of this model [22]. The behaviour of the system of interest can then be inferred by experimenting with and analyzing of the operation of the model.

In agent-based modeling and simulation (ABMS), systems are modeled using using autonomous, interacting agents within an environment [16]. This stems from the idea that a global phenomenon can be generated from the actions and interactions between the agents of a system [23]. The strengths of ABMS are its ability to express phenomena that emerge from many interactions, how it provides a natural description of a system, and its flexibility [24]. Given these strengths, ABMS has been used in numerous studies involving complex systems, across a large range of domains such as pandemics, the immune system, and water management ([25], [26], and [27] respectively).

The study of cooperation and competition has seen its fair use of ABMS [28]. Axelrod's original work using the Iterated Prisoner's Dilemma included both mathematical analysis and the use of ABMS in the form of tournaments. The use of such simulations has been thoroughly reviewed by Gotts et al. [29]. Today there seems to be little use to simulation studies for simple Iterated Prisoner's Dilemma games where every two agents have equal probability to meet. In these cases, mathematical analysis is strongly encouraged. The situation is however different for spatial versions of the game. The stochastic nature of such environments makes mathematical analysis particularly difficult, while using ABMS is trivial and has the potential to provide valuable insights.

Given the above reasoning and the used spatial IPD model presented in the next section, ABMS will be used to study the relation between noise and harshness of the environment.

4. Model Description

The ABMS model used to study the relation between noise and harshness in IPD games is based on the work of Smaldino et al. [10]. Their work presents a spatial IPD model that simulates harshness of the environment with a cost of life K and a varying sucker's payoff S. These parameters directly adjust the harshness of the environment, making experimentation over different levels of harshness particularly convenient. Furthermore, this model has been used in other works on cooperation in harsh environments [11, 30]. However, the model only includes always cooperating or defecting strategies, and assumes no mistakes are made. A larger set of strategies with an extension of the model has been explored, but under the same assumption that no mistakes were made and the conclusion that cooperating, but non-forgiving strategies were most successful [11]. In order to study the effects of noise and the use of generosity and contrition in strategies of the IPD, the model of Smaldino et al. is extended. A formal description of this extended model is described in this section.

Agents are initially placed on a $L \times L$ discrete square grid with toroidal boundaries, which serves as the environment. Each simulation starts with N agents, each placed on a random, unique location. Every grid cell can only be occupied by a single agent over the whole simulation, and a total maximum of N^* agents can be present on the entire grid.

Each time step, every agent tries to play a round of the IPD game with another agent. If an agent has not played yet, it will look in its local neighborhood, i.e. the 8 closest cells, for another agent that has not played yet during that time step. If the agent cannot find such a co-player, it moves to a random, unoccupied cell in that same neighborhood. If all cells are occupied, the agent will remain on the same cell and take no action.

Each agent has its own energy level, initially drawn from a uniform distribution between 1 and 50. Energy can be gained or lost by playing a round of the IPD game with another agent. The payoffs of the game are set to $T = 5$, $R = 3$, and $P = 0$; $S \leq 0$ is varied to simulate harshness. Furthermore, a cost of life K is subtracted at the end of each agent's action (for every time step), regardless of whether the agent moved, played, or stayed stationary. If an agent's energy level gets to 0 or less, the agent dies and is removed from the simulation. An agent can attempt to reproduce when its energy is at least 100. If a cell in its neighborhood is unoccupied, for a cost of 50 energy units, a new offspring agent is created on that cell with the same strategy of the original agent and energy level of 50. The total energy is capped at 150, to avoid agents accumulating energy indefinitely.

The harshness of the environment can therefore directly be controlled by the cost of living K, and the sucker's payoff S. A higher K and a lower S lead to harsher environments, with $P < K < R$. If $K > P$, defectors could survive on their own, and if $K > R$ cooperators would not be able to survive and all the agents would go extinct.

In order to simulate mistakes, an extra parameter for noise E is introduced. E is the probability that an action of an agent actually results in the opposite action. For example, for $E = 0.01$, an agent that decides to cooperate (according to its strategy) has a 1% chance of defecting instead.

Strategies are evenly distributed over the N initial agents. In order to study the impact of noise on cooperation directly, the original two strategies (always cooperating and always defecting) will be considered. However, to study whether generosity and contrition can be successful in harsh environments, a second experiment with a larger set of strategies is required.

In 2012, M. Jurišić et al. conducted a review of IPD strategies over the last 40 years. In this review they present a set of 9 default strategies, described in Table 2. Given the wide range of strategy types it provides, this set will be included in the model for competition experiments.

Designation	Description
ALLC	Strategy always plays cooperation
ALLD	Strategy always plays defection
RAND	Strategy has a 50% probability to play cooperation or defection
GRIM	It starts with cooperation, but after the first defection of its opponent continues with defection
TFT	It starts with cooperation and then copies the moves of the opponent
TFTT	As TFT but defects after two consecutive defections
STFT	As TFT but starts with defection
TTFT	As TFT but for each defection retaliates with two defections
Pavlov	Action results are divided into 2 groups, positive actions are T and R and negative actions are P and S - if the result of previous action belonged to the first group, action is repeated and if the result in the second group, then the action was changed, it is also called win-stay, lose shift.

Table 2: Default types of strategies presented in the review of Iterated Prisoner's Dilemma Strategies [31].

In the study of the effects of noise in the classic IPD, 3 strategies were highlighted [12]. These are described in Table 3. *Generous Tit for Tat* (GTFT) and *Generous Pavlov* (GPavlov) are generous variants of the original strategies, while *Contrite Tit for Tat* (CTFT) is the contrite variant of TFT. These strategies performed better than other strategies in the presence of noise in the classical IPD. These strategies will be included in the model with the other strategies to study whether their performance translates to harsh environments as well. Note however that the used model of this article is a spatial IPD.

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

Table 3: Strategies to cope with noise in a classic Iterated Prisoner's Dilemma [12]

Grim has shown that in the presence of noise greater generosity is to be expected in a spatial context than in a classic one [15]. The most successful strategy in his work was an even more generous version of TFT that cooperates 60% of the time it would otherwise defect. To study whether the harshness of the environment (and not only the spatial nature) leads to increased generosity, two extra generous strategies are included in the model. These strategies are *Spatial Generous Tit for Tat* (SGTFT) and *Spatial Generous Pavlov* (SGPavlov) and are described in Table 4.

Designation	Description
SGTFT	As TFT, but has a 60% probability of cooperating when it would otherwise defect.
SGPavlov	As Pavlov, but has a 60% probability of cooperating when it would otherwise defect.

Table 4: Spatial versions of the generous strategies. Greater generosity is to be expected in spatial environments, with a GTFT cooperating 60% of the time it would otherwise defect as best performing strategy [15]

.

5. Simulation Results

The model described in the previous section is implemented using the MASON Multi Agent Simulation Toolkit [32]. For every round, all agents were stepped once in random order. Results were measured at the end of each round, i.e. after all agents have been stepped. Interested readers willing to reproduce the results are encouraged to read Section 6 for more information.

Two experiments were conducted on the model, each with a different purpose. The first experiment is aimed at studying the effects of noise and harshness on cooperation, while the second experiment focuses on the effects on different strategies.

The independent variables for both experiments were the cost of life K, sucker's payoff S, and noise level E. K was varied between 0 and 2.5, and S between -2.5 and 0. Higher (or lower) values would result in an environment that is too harsh for agents to survive. E was varied between 0 (no noise) and 0.1 (10% noise).

The setup and results of both experiments are described in Section 5.1 and 5.2, respectively.

5.1. Effects of noise and harshness on cooperation

In order to create a perfect balance between cooperating and defecting agents, only always cooperating (ALLC) and always defecting (ALLD) strategies were used in the model for this experiment. These strategies were equally distributed: half of the initial agents were cooperating and the other half defecting. This is the same setup as the original model [10].

In this experiment the average fraction of cooperators was measured for different levels of noise and harshness. The average fraction of cooperators was computed as follows: For round i, let C_i be the number of cooperators (ALLC agents) in that round, and D_i the number of defectors (ALLD agents). The average fraction of cooperators F is the fraction of the sum of the number of cooperators over all rounds over the sum of all agents over all rounds as shown in Equation 3:

$$F = \frac{\sum\limits_i C_i}{\sum\limits_i C_i + \sum\limits_i D_i} \tag{3}$$

This metric of cooperation differs from the one used in the original model [10]. In the original model, the average cooperation frequency represented the levels of cooperation. This frequency is similar to the average fraction of cooperators F, but the number of cooperations is counted instead. It is however difficult to interpret results with this metric in a noisy environment. Remember that noise affects both cooperators and defectors. This means that every cooperation has a given probability to result in defection, while every defection has the same probability to result in cooperation. Therefore, in an environment with a majority of cooperators there will be more accidental defections than accidental cooperations, and vice versa. Cooperation could seem to decrease or increase, while being a simple side effect of noise. Take an extreme example: an environment with only cooperators. Under no noise, 100% of the actions are cooperations. But under 10% noise, only 90% of the actions are due to accidental defections. The number of cooperators has not changed, but the average cooperation frequency did. This justifies the use of average fraction of cooperators F over the average cooperation frequency in noisy environments.

Agents were placed on a 50×50 grid ($L = 50$) and ran for 10^5 rounds. The initial population size N was set to 10% of the grid size ($0.1 \cdot L^2$) with a maximum

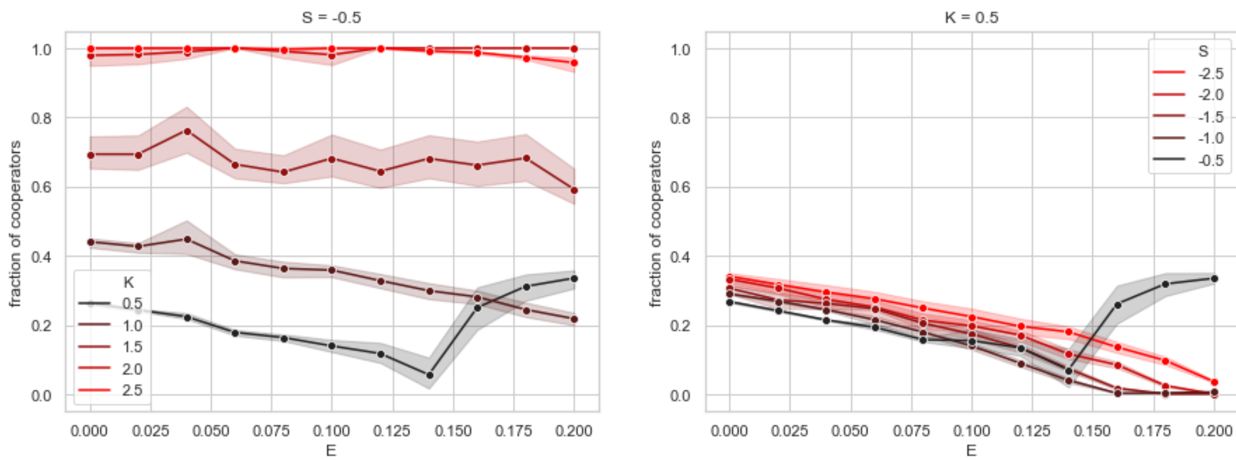

Figure 1: Influence of noise on the average fraction of cooperators for different levels of harshness. On the left figure the sucker's payoff S is kept fixed at a non harsh value, and cost of life K is varied. On the right figure K is kept fixed at a low value and S is varied. Redder lines represent harsher environments.

size N^* of 50% of the grid size $(0.5 \cdot L^2)$. Originally, each combination of independent variables (cost of life K, sucker's payoff S, and noise E) was ran 10 times, and it became clear that varying S or K led to different types of results. In a second batch, either S or K were held constant, and each combination ran 20 times. The results of this second batch are summarized in Figure 1. Results for other combinations were similar, except for very high S and K where the environment was too harsh for agents to survive.

The grid size and number of rounds is chosen to be big enough to provide meaningful results. Originally the experiment was ran on a 100×100 grid with 10^6 runs to replicate the results of the original model [10]. However, these simulations were time-consuming and were not realistic for this research's time frame. Smaller grid sizes were explored and similar results were found down to $L = 40$. For smaller grid sizes populations would occasionally go extinct for harsh environments. $L = 50$ was therefore chosen to be safe. For this grid size the average cooperation frequency seemed to stabilize well before 10^5 rounds. A couple of experiments were ran for 10^6 rounds with the same results.

When the environment becomes harsher (by increasing K or decreasing S), cooperation seems to increase. This aligns with the results in the work of Smaldino et al. [10]. Furthermore, increasing harshness with the cost of life K seems to have a bigger impact than varying the sucker's payoff S. This might be due to the fact that cost of life affects both cooperators and defectors, while the sucker's payoff only hurts players that cooperate with defectors.

The rates at which cooperation decreases seems the same for all levels of harhsness, with the exception of a high cost of life ($K \geq 1.5$) or an environment that is not harsh ($S = -0.5$ and $K = 0.5$). The fact that the rates at which cooperation decreases seem to be the same in general leads to believe that the harshness of the environment does not have a direct impact on the effects of noise. But the side effect of increased harshness, namely increased cooperation, could explain the deviations from this rate. In an environment without noise ($E = 0.0$), if the harshness is in favor of cooperators (fraction of cooperators ¿ 0.5), cooperation seems to remain around the same levels when noise increases.

Interestingly, when the environment is not harsh ($S = -0.5$ and $K = 0.5$) cooperation suddenly increases for higher noise ($E \geq 0.16$). This can be explained by two factors. First, noise can act as a "defense mechanism" for cooperators by giving them

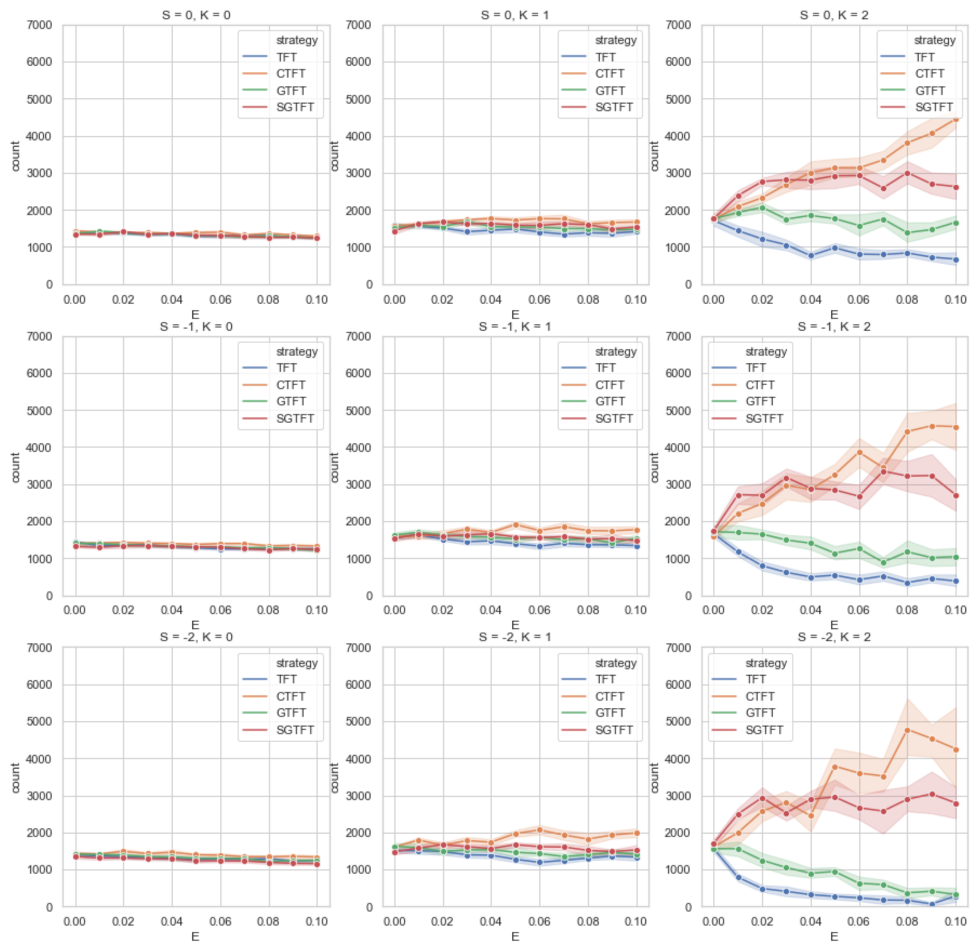

Figure 2: Influence of noise on variants of the TFT strategy in a tournament setting for different levels of harshness.

the probability to defect against a defector (and a probability of a defector cooperating with a cooperator). Second, when adversity of the environment is low accidentally defecting with another cooperator is less critical.

5.2. Effects of noise on strategies in harsh environments

This experiment took form as tournaments with six different sets of strategies under different levels of noise and harshness. The first set was the set of default strategies, presented in Table 2. The second, third, and fourth set were the same as the first set, with the exception of TFT being replaced with its noise tolerant counterparts, CTFT, GTFT, and SGTFT (see Table 3 and Table 4). In the two last sets the Pavlov strategy was replaced with the noise tolerant GPavlov (Table 3) and SGPavlov (Table 4).

Tournaments took place on a 100×100 grid with an initial population size N at 10% of the grid's capacity ($0.1 \cdot L^2$). The strategies of the chosen set were distributed evenly over the initial agents. The maximum capacity N^* was set to the grid size (L^2) and simulations ran until the grid was full, which would happen after ca. 300 to 5000 rounds, depending on

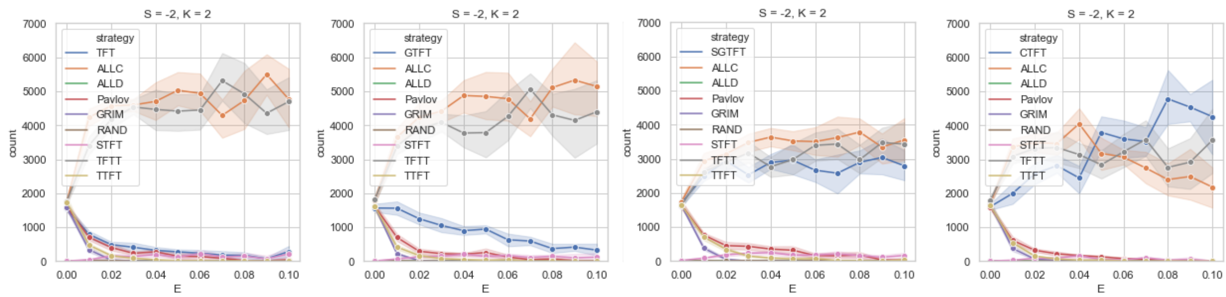

(a) TFT (b) GTFT (c) SGTFT (d) CTFT

Figure 3: Performance of TFT strategy and its variants (in blue) compared to other strategies in a tournament setting in a very harsh environment.

the harshness and noise of the environment (harsher and noisier leads to more rounds).

The number of agents with a given strategy at the end of the simulation served as performance measure. Successful strategies could accumulate more energy, reproduce quicker, and therefore grow their population size faster than their peers. Due to the random initial position of agents, these numbers can slightly vary per tournament. Every tournament is therefore repeated 10 times (for every set of strategies and independent variables) and the average count of these runs is used as final metric.

The performance of the TFT strategy compared to its noise tolerant variants GTFT, SGTFT, and CTFT for different levels of harshness is shown in Figure 2. Similarly to the results of the first experiment shown in Figure 1, the effects of increasing cost of life K are much greater compared to decreasing the sucker's payoff S. When the environment is not harsh ($K = 0$) there does not seem to be much difference in performance between strategies. Even when noise is high, the 4 strategies perform the same.

The impact of noise on the different strategies becomes clear when the cost of life K is increased. In harsher environments, higher noise seems to decrease the performance of TFT. GTFT, the generous version that cooperates 10% of the time it would otherwise defect, also sees a decrease in performance, though at a smaller rate. This could be due to the spatial nature of the model, as the more generous SGTFT is performing better when noise increases. The contrite version of TFT, CTFT, seems to thrive in harsh environments with high noise ($E > 0.04$).

There are 3 trends that can be observed when comparing the performance of TFT, GTFT, SGTFT, and CTFT with other strategies in the tournament (Figure 3a, 3c, and 3d respectively). First, the very generous strategies ALLC (always cooperate), TFTT (TFT but only defect after two consecutive defections), and SGTFT (TFT but cooperate 60% of the time it would otherwise defect) seem to dominate other strategies once noise is present. Second, the less generous GTFT's performance seems less affected by noise than other strategies, but does not seem to benefit from it similarly to TFTT, ALLC and SGTFT. Finally, for higher noise ($E > 0.05$) CTFT seems to outperform all the other strategies.

Unlike GTFT and SGTFT, the generous versions of Pavlov (GPavlov and SGPavlov) do not seem to see the same benefits when noise is present. Figure 4 compares the performance between the three strategies in a harsh environment. Pavlov and GPavlov seem to perform the same regardless of the level of noise. SGPavlov is more successful under noise compared to the other two strategies, but without any performance increase like SGTFT. The same results were observed for different levels of harshness.

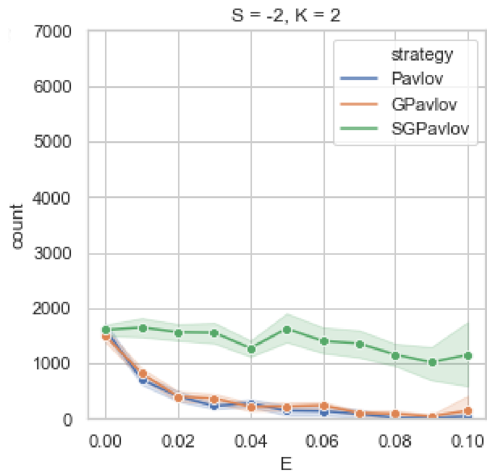

Figure 4: Influence of noise on Pavlov and generous variant strategy in a tournament setting with harsh environment.

6. Responsible Research

This section provides more information on the reproducibility of the results in order to address the reproducibility crisis. The complete model description is given in Section 4. The experiment procedures and used parameters are given in Section 5. These should provide interested readers enough information to reproduce the obtained results. Due to the stochastic nature of ABMS, results are subject to slight variations. The complete code used to produce the results of this article together with the raw data are therefore made available on a public repository[1]. The seeds of the random number generator used to produce the data are included in the raw data files in order to reproduce the exact same results. The code uses the MASON Multi Agent Simulation Toolkit [32], which is well documented online[2].

[1]`https://gitlab.com/louis.gevers/noise-harsh-ipd`

[2]`https://cs.gmu.edu/~eclab/projects/mason/manual.pdf`

7. Discussion

7.1. Summary of results

In the presence of strict cooperators and defectors noise seems to affect cooperation on most levels of harshness. The more noise is present, the more difficult cooperation becomes. However, harsher environments lead to increased cooperation. In environments where harshness results in a majority of cooperators, cooperation seems more robust against noise.

When the environment is populated with a more varied set of behaviours, the relation between noise and harshness becomes more apparent. When adversity is low, the success of different behaviours does not seem to be affected by noise. However, when the adversity is increased, most noise tolerant behaviours prevail. Generous behaviours that forgive opponent defections by continuing to cooperate are dramatically more successful in harsh conditions when noise is present. Contrite behaviours that focus on fixing one's own mistakes rather than forgiving the opponent mistakes finds similar success to generous behaviour in harsh environments. While slightly outperformed when noise is low, their success directly increases with noise. Contrite behaviours are therefore more successful when communication is particularly unreliable.

Not all classic noise tolerant behaviours have this success however. The Pavlov strategy fails just as bad as a classic TFT, and so does the more generous GPavlov. As a result of greater generosity SGPavlov is more successful, yet it is still way outperformed by SGTFT. Behaviours based on reciprocity with great generosity (even unconditional cooperation) or contrition are therefore preferred.

7.2. Interpretation

Both experiments show that noise must not be neglected when studying cooperation. Cooperation is directly affected by noise, on all levels of harshness of the environment. In non harsh environments, cooperation will decrease in the presence of noise. While in harsher environments it becomes more crucial to actively cope with noise. The success of generous or contrite behaviours can be truly appreciated when communication is unreliable and the adversity of the environment is high.

7.2.1. Diversity of behaviours

The results of both experiments seem to differ from each other. It follows from the first experiment that the presence of noise makes cooperation more difficult, while in the second experiment generous and cooperative strategies thrive under noisy conditions when environmental adversity is high. Furthermore, the most generous strategy, ALLC, which was used in both experiments performed differently in each. In the first experiment its performance decreased in presence of noise, or remained stable when the harshness of the environment was high. Meanwhile in the second experiment its performance greatly increased with noise in harsh environments.

This difference in cooperation could potentially be explained by the work of Chong and Yao, which describes how the diversity of behaviours affects the success of generous strategies under noise [33]. In their work they study the effects of noise on the IPD game with intermediate choices in a co-evolutionary environment. Instead of having different sets of fixed strategies, behaviour is modeled using neural networks. They draw two important conclusions based on their results. First, higher levels of noise are detrimental to cooperation. Second, in the presence of noise the diversity of behaviours helps the evolution of cooperation. While their method is vastly different from this article, their conclusion could explain how noise discourages cooperation when only ALLC and ALLD strategies are present (experiment 1) and encourages cooperative and generous strategies in the presence of a larger set strategies (experiment 2).

7.2.2. Generosity

The success of generosity in noisy environments has been widely supported. Numerous studies [12, 15, 34, 35] conclude that adding generosity is a successful way for cooperation to succeed in presence of noise. The results of this article indicate that the importance of generosity is emphasized when the harshness of the environment increases.

Unlike the classic IPD where moderate generosity (GTFT) works best [12], the results of this article indicate that greater generosity (SGTFT, TFTT, and ALLC) leads to more success. Before jumping to any conclusions, note that the model is a spatial IPD and not a classic IPD. Grim has shown that in the presence of noise greater generosity is to be expected in a spatial context than in a classic one [15]. While the most successful strategy in that work is SGTFT (cooperates 60% of the time it would otherwise defect), its generosity is still inferior to ALLC (cooperates 100% of the time). This combined with the fact that ALLC performed better than SGTFT in harsher environments leads to the conclusion that in the presence of noise, harsher environments promote greater generosity.

This contradicts the previous conclusions of the success of less generous strategies such as GRIM in harsh environments [11], which confirms the importance of studying the Prisoner's Dilemma with noise. Real-world interactions are error-prone, and the presence of such mistakes have a great impact on the success of different strategies. This leads to believe that great generosity is more likely to thrive in real environments with harsh conditions than non forgiving behaviours.

7.2.3. Contrition

Under higher noise, contrition (CTFT) seems to outperform generosity. The same phenomenon can be observed in the classic IPD [12]. The advantage of CTFT is that it works well in overtaking environments with defectors, while generous strategies rely on the presence of other cooperating strategies to succeed [36]. The presence of less generous or defecting strategies such as ALLD and STFT in this article's tournaments might explain why CTFT outperforms the generous variants.

While contrition seems to work better than generosity in heterogeneous environments, the existing literature about its success is less convincing. Smith and Price introduced the concept of an evolutionarily stable strategy (ESS), a strategy that would be stable under natural selection [37]. Different studies provide various conditions under which CTFT can be evolutionarily stable [12, 38], however in general contrition is not considered to be an ESS [39]. This leads to believe that while contrite behaviours might work well in simulations, generosity is more likely to emerge in the real world.

7.3. Current limitations

It is worth noting that the Prisoner's Dilemma Game is a simple model that attempts to explain complex behaviours. While this game is a powerful model for reciprocal altruism [40], it is limited and cannot provide an accurate explanation for every phenomenon of altruism in nature. In his work, Connor presents 3 different models that have been used to explain altruistic behaviours in animals more accurately than the Prisoner's Dilemma [41]. The Prisoner's Dilemma is therefore not the only way to model cooperation, and other models should be considered before drawing conclusions for specific real-life phenomena.

This article is based on a specific model of the Prisoner's Dilemma. The model in question is based on the work of Smaldino et al. [10]. Another study has extended this model with more realistic features (evolution and aging) [11], yet both the extension and the model used in this article suffer from the same limitation: using a fixed set of strategies.

Each experiment used a fixed set of well-defined strategies. The main benefit of using such strategies lies in their simplicity. They are easy to encode and setting up an experiment is straightforward. And since these strategies are simple and distinct, interpreting results gives a broad view about different behaviours. Using well-defined strategies is therefore suitable when studying the general effects on different behaviours.

There are limitations to using such fixed strategies. While broader questions can be answered, the obtained results lack depth. To illustrate this three examples from this article are given. First example, when noise is present harsher environments promote more generosity. With the current results however it is unclear how much more generosity is to be expected with harshness. To answer this question an approach with continuous levels of generosity is preferred, as in Grim's study of generosity in spatial IPD [15]. Second example, it follows from the results of this article and existing literature [33, 36] that the presence of some strategies influences the performance of other strategies. Which strategies are influenced and how they are influenced remains unclear. Third example, there exists more complex strategies based on different concepts than generosity and contrition, such as constructing a model of the other individual's behaviour [42], that could potentially cope with noise.

8. Conclusions and Future Work

This article describes the relation between noise and harshness of the environment. While harsh environments benefit cooperators, noise makes cooperation more difficult. Hence the interest in studying the relation between the two. When the adversity of the environment is too weak to be beneficial for cooperators, the effects of noise are greater and cooperation decreases. Counterintuitively, harsher environments can help cooperators to be more robust against mistakes. Furthermore, harsh environments promote greater generosity. Strategies based on reciprocity that forgive the great majority (if not all) of the opponents defections are most successful in harsh environments when communication is unreliable. For higher noise, contrition is more successful than generosity. The strategy *Contrite Tit for Tat* that focuses on fixing its own mistakes rather than forgiving the opponents mistakes outperforms even the most generous strategies in harsh environments.

The results of this work are based on the model of Smaldino et al. [10], using the classic strategies presented in the review of strategies by M. Jurišić et al. [31] and noise tolerant strategies summarized in Wu and Axelrod's work [12]. The model can be extended in the future to study the impact of other environmental factors in the presence of noise. For example, an extension of the model could be used to study the effects of evolution and ageing on different strategies when noise is present [11]. Environmental factors also include the (non) presence of other behaviours. Further work could study the presence of which behaviours catalyse or restrict generosity, and explore more complex behaviours by using different models than fixed strategies.

References

[1] S. A. West, A. S. Griffin, and A. Gardner, "Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection," *Journal of evolutionary biology*, vol. 20, no. 2, pp. 415–432, 2007.

[2] O. Harman, *The price of altruism: George Price and the search for the origins of kindness*. WW Norton & Company, 2011.

[3] D. S. Wilson and L. A. Dugatkin, "Group selection and assortative interactions," *The American Naturalist*, vol. 149, no. 2, pp. 336–351, 1997.

[4] S. A. West, A. S. Griffin, and A. Gardner, "Evolutionary explanations for cooperation," *Current biology*, vol. 17, no. 16, pp. R661–R672, 2007.

[5] J. A. Fletcher and M. Doebeli, "A simple and general explanation for the evolution of altruism," *Proceedings of the Royal Society B: Biological Sciences*, vol. 276, no. 1654, pp. 13–19, 2009.

[6] R. Axelrod and W. D. Hamilton, "The evolution of cooperation," *science*, vol. 211, no. 4489, pp. 1390–1396, 1981.

[7] P. Andras, J. Lazarus, and G. Roberts, "Environmental adversity and uncertainty favour cooperation," *BMC Evolutionary Biology*, vol. 7, no. 1, p. 240, 2007.

[8] I. Krams, A. Bērziņš, T. Krama, D. Wheatcroft, K. Igaune, and M. J. Rantala, "The increased risk of predation enhances cooperation," *Proceedings of the Royal Society B: Biological Sciences*, vol. 277, no. 1681, pp. 513–518, 2010.

[9] R. M. Callaway, R. W. Brooker, P. Choler, Z. Kikvidze, C. J. Lortie, R. Michalet, L. Paolini, F. I. Pugnaire, B. Newingham, E. T. Aschehoug, *et al.*, "Positive interactions among alpine plants increase with stress," *Nature*, vol. 417, no. 6891, pp. 844–848, 2002.

[10] P. E. Smaldino, J. C. Schank, and R. McElreath, "Increased costs of cooperation help cooperators in the long run," *The American Naturalist*, vol. 181, no. 4, pp. 451–463, 2013.

[11] J. van Tilburg and N. Yorke-Smith, "Strategies for the iterated prisoner's dilemma in a natural environment," 2018.

[12] J. Wu and R. Axelrod, "How to cope with noise in the iterated prisoner's dilemma," *Journal of Conflict resolution*, vol. 39, no. 1, pp. 183–189, 1995.

[13] A. Szolnoki, J. Vukov, and G. Szabó, "Selection of noise level in strategy adoption for spatial social dilemmas," *Physical Review E*, vol. 80, no. 5, p. 056112, 2009.

[14] Z.-X. Wu and P. Holme, "Effects of strategy-migration direction and noise in the evolutionary spatial prisoner's dilemma," *Physical Review E*, vol. 80, no. 2, p. 026108, 2009.

[15] P. Grim, "Spatialization and greater generosity in the stochastic prisoner's dilemma," *BioSystems*, vol. 37, no. 1-2, pp. 3–17, 1996.

[16] C. M. Macal and M. J. North, "Agent-based modeling and simulation," in *Proceedings of the 2009 Winter Simulation Conference (WSC)*, pp. 86–98, IEEE, 2009.

[17] 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.

[18] R. Axelrod, "The evolution of strategies in the iterated prisoner's dilemma," *The dynamics of norms*, pp. 1–16, 1987.

[19] M. A. Nowak and R. M. May, "The spatial dilemmas of evolution," *International Journal of bifurcation and chaos*, vol. 3, no. 01, pp. 35–78, 1993.

[20] E. Stanley, D. Ashlock, and L. Tesfatsion, "Iterated prisoner's dilemma with choice and refusal of partners," *Artificial Life*, vol. 17, no. 3, pp. 131–175, 1992.

[21] B. P. Zeigler, T. G. Kim, and H. Praehofer, *Theory of modeling and simulation*. Academic press, 2000.

[22] A. Maria, "Introduction to modeling and simulation," in *Proceedings of the 29th conference on Winter simulation*, pp. 7–13, 1997.

[23] F. Klügl and A. L. Bazzan, "Agent-based modeling and simulation," *Ai Magazine*, vol. 33, no. 3, pp. 29–29, 2012.

[24] E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," *Proceedings of the national academy of sciences*, vol. 99, no. suppl 3, pp. 7280–7287, 2002.

[25] J. M. Epstein, "Modelling to contain pandemics," *Nature*, vol. 460, no. 7256, pp. 687–687, 2009.

[26] V. A. Folcik, G. C. An, and C. G. Orosz, "The basic immune simulator: an agent-based model to study the interactions between innate and adaptive immunity," *Theoretical Biology and Medical Modelling*, vol. 4, no. 1, p. 39, 2007.

[27] E. Z. Berglund, "Using agent-based modeling for water resources planning and management," *Journal of Water Resources Planning and Management*, vol. 141, no. 11, p. 04015025, 2015.

[28] E. Elliott and L. D. Kiel, "Exploring cooperation and competition using agent-based modeling," *Proceedings of the National Academy of Sciences*, vol. 99, no. suppl 3, pp. 7193–7194, 2002.

[29] N. M. Gotts, J. G. Polhill, and A. N. R. Law, "Agent-based simulation in the study of social dilemmas," *Artificial Intelligence Review*, vol. 19, no. 1, pp. 3–92, 2003.

[30] P. E. Smaldino, "Cooperation in harsh environments and the emergence of spatial patterns," *Chaos, Solitons & Fractals*, vol. 56, pp. 6–12, 2013.

[31] M. Jurišić, D. Kermek, and M. Konecki, "A review of iterated prisoner's dilemma strategies," in *2012 Proceedings of the 35th International Convention MIPRO*, pp. 1093–1097, IEEE, 2012.

[32] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, "Mason: A multiagent simulation environment," *SIMULATION*, vol. 81, no. 7, pp. 517–527, 2005.

[33] S. Y. CHONG and X. YAO, "Behavioral diversity, choices and noise in the iterated prisoner's dilemma," *IEEE transactions on evolutionary computation*, vol. 9, no. 6, pp. 540–551, 2005.

[34] P. A. Van Lange, J. W. Ouwerkerk, and M. J. Tazelaar, "How to overcome the detrimental effects of noise in social interaction: the benefits of generosity.," *Journal of personality and social psychology*, vol. 82, no. 5, p. 768, 2002.

[35] J. Bendor, R. M. Kramer, and S. Stout, "When in doubt... cooperation in a noisy prisoner's dilemma," *Journal of conflict resolution*, vol. 35, no. 4, pp. 691–719, 1991.

[36] M. C. Boerlijst, M. A. Nowak, and K. Sigmund, "The logic of contrition," *Journal of Theoretical Biology*, vol. 185, no. 3, pp. 281–293, 1997.

[37] J. M. Smith and G. R. Price, "The logic of animal conflict," *Nature*, vol. 246, no. 5427, pp. 15–18, 1973.

[38] R. Boyd, "Mistakes allow evolutionary stability in the repeated prisoner's dilemma game," *Journal of theoretical Biology*, vol. 136, no. 1, pp. 47–56, 1989.

[39] C. Hilbe, "Contrition does not ensure cooperation in the iterated prisoner's dilemma," *International Journal of Bifurcation and Chaos*, vol. 19, no. 11, pp. 3877–3885, 2009.

[40] R. Boyd, "Is the repeated prisoner's dilemma a good model of reciprocal altruism?," *Ethology and Sociobiology*, vol. 9, no. 2, pp. 211 – 222, 1988.

[41] R. C. Connor, “Altruism among non-relatives: alternatives to the ‘prisoner’s dilemma’,” *Trends in Ecology Evolution*, vol. 10, no. 2, pp. 84 – 86, 1995.

[42] T.-C. Au and D. Nau, “Accident or intention: that is the question (in the noisy iterated prisoner’s dilemma),” in *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pp. 561–568, 2006.