Effects of Group-based Agents on Iterated Prisoner's Dilemma Strategies in Evolving Spatial Environments

Gijs de Jong[1] , under supervision of **Neil Yorke-Smith**[1]
[1]TU Delft

Abstract

The feeling of belongingness, to be a member of a group, is rooted in human evolutionary history. Cooperative behaviour within such groups has since been an important research topic. The evolution of cooperation in the iterated prisoner's dilemma (IPD) has been shown to be an effective tool of simulating and analysing this behaviour. However, it is unclear what the effects of group-based agents on IPD strategies in evolving spatial environments are. This paper investigates *how* this cooperation emerges by proposing an evolving spatial model that applies a genetic algorithm to its agents, extended to work with three distinct group types. This genetic algorithm makes use of four genetic operators: cloning, crossover, mutation and inversion. The groups considered are kin, clans and their combination. Cooperation is measured by examining (1) populations levels divided into *nice*, *balanced* and *nasty* groups, and (2) average cooperation levels of both strategies and games played per iteration. Experiments for all group types with two reproduction preferences were conducted. Three distinct conclusions can be drawn from the results. First, strategies evolved through domestic reproduction exhibit more cooperation for in-group opponents and more defection for out-group opponents. Second, strategies evolved through wealthy reproduction exhibit the same increase and decrease of cooperation as domestic reproduction, but to a smaller degree. Third, an evolving spatial environment with group-based agents develops subgroups, defined by similar strategies and restricting group-wide cooperation. Thus, agents with cooperative domestic strategies and defective foreign strategies win, and there is a positive correlation between group size and restriction of group-wide cooperation.

Keywords: prisoner's dilemma, kin, clans, genetic algorithms, evolutionary games, spatial games

1 Introduction

The persistence of altruism and selfishness in nature is a long discussed topic in evolutionary game theory [1]. The prisoner's dilemma (PD) has become a standard in game theory describing cooperation between two individuals [2]. It is a non-zero-sum game, where both players are given the choice to either cooperate (C) or defect (D). The reward of a single defection is defined to be higher than mutual cooperation, and mutual defection is lower than mutual cooperation [3]. In the classic symmetric case, the formal payoff matrix can be described by Table 1 satisfying the following two inequalities [4]:

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

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

		Player 2	
		C	D
Player 1	C	R, R	S, T
	D	T, S	P, P

Table 1: Payoff matrix of a two-player symmetric Prisoner's Dilemma

The terms represent the payoffs each player receives. Payoff R is defined as the *reward* both players receive for mutual cooperation. Term P is the *punishment* payoff rewarded for mutual defection. The *temptation* payoff T is rewarded to a player for solely defecting. Lastly, S is defined as the *sucker* payoff for a sole cooperator [4].

The iterated prisoner's dilemma (IPD) is a well-known extension of the original problem. It describes the same game, played repeatedly in succession [5]. Respective players often possess memory, allowing them to remember the results of previous games. The iterated prisoner's dilemma provides an abstract framework to analyse the evolution of cooperation [6]. Multi-agent simulations utilising this framework are often used to explain cooperation and cohesion (or the opposite) in nature and society [7].

Two important extensions of the IPD have been proposed: a spatial and evolving one. These extensions are derived from analysing physical environments. The aim here is not to provide an accurate representation, but a simple one sufficient to allow for complex simulated results [8].

In the first extension of the IPD, agents operate in a two-dimensional spatial environment. Agents can be represented as moving points in a grid. One iteration of the IPD can be seen as an agent moving and playing a game with a neighbouring opponent, the latter being derived from other model properties [9].

The second extension of the IPD involves an evolving environment. In general, this implies an environment that applies genetic algorithms exhibiting evolutionary processes, such as natural selection or reproduction, on its agents [10]. These processes require agents to maintain an energy value, which describes its fitness or "success". Note that the success of an agent is no different from the success of its respective strategy. Over time, an agent's fitness increases or decreases, depending on the outcome of related games or iterations [11]. Natural selection can then occur by eliminating agents without energy. Producing offspring could require a minimal energy level and deduct a given cost. Genetic algorithms have shown to be successful at discovering effective strategies in given environments [10]. Effective or winning strategies are often measured by their (relative) population size, meaning the total number of agents exhibiting that particular strategy.

This paper will discuss an additional extension of the IPD regarding an environment containing group-based agents. Group-based agents are interconnected through membership in the same groups. Related work involves investigating this extension with restricted group types from group- or agent-based perspectives. Hirshleifer and Rasmusen examined non-spatial populations as a single group with the possibility to (temporarily) remove defectors [12]. This possible expulsion was shown to be effective at promoting cooperation. Julstrom analysed the evolution of cooperation using genetic algorithms, when the population was segregated into a fixed number of clans [13]. Two important assumptions were made: contests (the total number of games) between clan members were longer, and reproduction only happened between same-clan members. Such a non-spatial environment was shown to promote higher levels of cooperation for domestic (same-clan) contests and lower levels for foreign contests. Li and Kendall introduced collective strategies for an evolutionary non-spatial IPD, which made agents cooperate with their group members and defect against any others [14]. These strategies were shown to be effective at maintaining stable populations, favouring collective behaviour. This does not necessarily imply more cooperation, merely cooperation to the same degree.

The groups used in the related work are restricted with regards to their size and composition. The environments are also limited, as they represent a non-spatial IPD. It is not yet known how cooperation (in strategies) evolves when a more relaxed definition of groups is used. Furthermore, how this evolution would develop in an evolving spatial environment is unclear.

The objective of this paper is to fill this gap by answering the main research question: how do group-based agents affect the iterated prisoner's dilemma strategies in evolving spatial environments? Three distinct group types are considered: kin, clans and their combination. Generalised definitions of groups are used in order to lift the restrictions mentioned for the related work. The evolving environment is defined similarly to Julstrom [13], applying a genetic algorithm on its agents.

The hypothesis this paper drafts is that group-based agents exhibiting cooperative domestic strategies and defective foreign strategies win in the iterated prisoner's dilemma in evolving spatial environments. A cooperative strategy is positively inclined towards cooperation, which is derived from two metrics considering the "niceness" and game cooperation respectively. Domestic and foreign strategies are in-group and out-group strategies. An agent is considered winning if its strategies are present in the majority of the population.

Section 2 dives into related work used for the base model in this paper. Section 3 discusses the main contributions made. Next, section 4 gives a detailed model description. Section 5 contains the experimental setup and results of simulations. Section 6 describes ethical and epistemic values of this paper. Section 7 discusses the results with respect to related work. Finally, section 8 contains the conclusion and a recommendation for future work.

2 Related work

This paper considers an extension on the demographic prisoner's dilemma for its base model. This section will define the demographic prisoner's dilemma and what its extension entails.

The demographic prisoner's dilemma (DPD) was proposed by Epstein as a "minimal" evolving spatial version of the IPD originally proposed by Axelrod [15]. Minimal refers to dropping two assumptions many other models relied on. The first assumes agents, or strategies, have access to memory. The second allows for "tags" in an environment, permitting agents to distinguish one from another. The new environment derived from dropping these assumptions consists of indistinguishable agents inheriting either an always cooperate or always defect strategy. This environment is evolving in the sense that agents accumulate energy by playing games, allowing them to procreate or die. Offspring will inherit the strategy of their parent, with a chance of mutation occurring. A spatial variant of the IPD was used to show that cooperation emerges in spatial zones specifically. The spatial model used is very similar to the one proposed by Smaldino [16]. In each round of the DPD spatial model, all agents move in a random available direction, play an iteration of the PD against their neighbours, and produce offspring. Neighbouring patches (for movement and offspring) and agents (opponents) are derived from an agent's vonNeumann neighbourhood. This procedure is explained further in section 4.1.

Another significant difference from Axelrod's original model proposition [10] is a different chain of inequalities used for the payoff matrix. In the DPD model payoffs can also be negative, formally given by:

$$T > R > 0 > P > S \tag{3}$$

Payoffs P and S were chosen to be negative to reinforce worse agents dying. The main reason for using a simplified model was to show that it suffices to generate cooperative zones.

An extension of the DPD was implemented by Radax and Rengs, which features two additional modes of interaction: kin and clans [17]. The definitions for these group types are the same as used in this paper and will be discussed in detail in section 4. While the authors did not publish any work on the extension, they did produce a case study on the original DPD [18]. This case study documented their efforts to replicate the original DPD [19], and attempt to improve the model. All the improvements either clarify missing details in the original paper or improve the usage of random numbers in the model. The extension of groups also contains these improvements.

3 Contribution

This paper reports two modifications of the DPD model. The first modification generalises the strategies and the second modification extends the model by adding support for group-based agents. As explained in section 2, the DPD model only considers two strategies (always cooperate and always defect) [15]. This paper generalises these strategies using a variant of genetic algorithms, of which an early introduction was proposed by Goldberg [20]. The genetic algorithm this paper relies on is similar to the description of Axelrod [10], but extended to work with clans as proposed by Julstrom [21]. This extension is further modified in order to support more group types. Specifically, this paper considers the following group types: kin, clans, and kin and clans combined. The base model discussed in section 2 already supports a restricted version of kin and clans. The main details of the genetic algorithm are given in section 4.3. The extension to the genetic algorithm needed to support additional groups is discussed in section 4.1.

The main reason for these two modifications is that, together, they allow for more insight into the kin- and clan-wise evolution of cooperation. Epstein showed, under specific assumptions, that the assumptions are sufficient to generate cooperative persistence in spatial environments [15]. This paper will not only show that cooperation persists under the given conditions and assumptions, but also show *how* the cooperation persists. The extension of the group type with kin and clans combined is considered, because it allows for examining the cooperation (or the opposite) with respect to each other.

Adding support for generalised strategies through a genetic algorithm means the first assumption of the DPD model is no longer true. The new strategies can be classified as memory-3, as they make use of a three round game history to determine a next move. More details on encoding histories will be given in section 4.1. This violates the first assumption of the DPD model, which assumes agents to be memory free. The second assumption of the DPD still holds, which states the environment is tagless. It is important to note that this entails agents indistinguishable from another regarding their strategy. While this paper proposes a model where agents can recognise the group of an opponent, the respective strategies are still hidden.

4 Model Description

The model this paper proposes is an extension of the DPD model. Agents operate in a spatial grid environment consisting of patches, playing moves from different strategies depending on their opponents. Agents move to find neighbours and produce offspring at a free patch with a suitable mate after reproduction requirements have been met. These requirements consist of a minimum age and wealth. Wealth is accumulated by playing PD games and reduces depending on the cost of living. A flowchart of the main life cycle of agents in this model is depicted in Figure 1. Note that an agent plays a PD game with all of its neighbours. The energy of an agent is equal to its wealth. Metabolism is the cost of living for an agent, which is deducted after every round.

4.1 Model definitions

This section will present the definitions for the underlying concepts of this model.

Group types

As discussed in section 3, this paper considers three distinct group types: kin, clans and their combination. This paper considers generalised definitions for kin and clans, aiming to lift restrictions present in related work. Kin is a group defined by similar ancestors, which grows in size through reproduction. While the list of ancestors of an agent is constant, storing ancestors of both parents (and their parents) allows an agent to be in multiple kin. A clan is a group defined by mutual recognition of a clan leader and grows through accumulating new members. This paper defines clan membership to be atomic and immutable, meaning clans also only grow through reproduction.

Encoding strategies

The same procedure proposed by Bukhari is used to encode strategies [3], but with one distinction following the work of Errity [22]. This procedure considers memory-3 strategies that rely on a three round game history to determine a next move. Strategies are encoded as a bitstring consisting of C's (1's) for cooperate and D's (0's) for defect. A previous game is stored in the history as a 2-letter string consisting of the agent's and opponent's move. For a history of the last three games, this means a total of 64 (2^6) moves can be represented. An agent's move index is determined by looking up the index of its history in a table of all possible game histories. The actual move is simply the value of the strategy at the given index. Axelrod suggested that 6 extra bits could be prepended to the strategy to represent the fictitious pre-game history [10]. However, assuming a random pre-game history is not appropriate for offspring produced in a spatial environment that does not use a replacement technique. Following the method of Errity [22], 7 bits are prepended to the strategy instead. These bits represent the first move, and histories for first and second game respectively. This removes the need for assuming any pre-game history and means strategies are of length 71. The resulting search space (2^{71}) is too big to be searched exhaustively, which a search procedure through genetic algorithms could fix. Table 2 shows the encoding of a potential strategy with corresponding bits or indices.

In order to extend the environment with group-based agents, group-based strategies are needed. An extended variant of the procedure proposed by Julstrom will be used for this [21]. This procedure simply assigns multiple strategies to an agent, each keeping track of their own history. With both kin and clans present, agents need a total of 4 strategies for all group relations: in-kin and in-clan, in-kin and out-clan, out-kin and in-clan, and out-kin and out-clan. It is important to note that game histories are remembered by group relations. A three round history of the in-kin and in-clan strategy does not necessarily represent the moves of a single opponent or group.

Bit	History	Move
0	First move	C/D
1	C	C/D
2	D	C/D
3	CC	C/D
4	CD	C/D
5	DC	C/D
6	DD	C/D
7	CCCCCC	C/D
8	CCCCCD	C/D
9	CCCCDC	C/D
10	CCCCDD	C/D
11	CCCDCC	C/D
⋮	⋮	⋮
69	DDDDDC	C/D
70	DDDDDD	C/D

Table 2: Encoded strategy of 71 bits [3]

Strategy types

In addition to measuring overall cooperation in strategies and games played, examining the population could give more insights. To analyse the population, all strategies (and therefore the population) are divided into three groups: nice, balanced and nasty. This division is done by a computed cardinality based on two metrics considering betrayal and punishment with respective thresholds. Formally, cardinality c of a strategy with betrayal c_b and punishment c_p is of the following form:

$$c = 1 - \frac{c_b + c_p}{2} \tag{4}$$

The reason for favouring these metrics over, for example, a simple count of Cs in a strategy, is that this count does not reflect the true "niceness" of a strategy. Whether a move is "nice" or not depends on the history associated with it. The reason for not including other metrics, such as forgiveness, is that betrayal and punishment metrics combined are already sufficient to give an approximate division of the population.

The first metric is the betray-cardinality, which describes the factor of betrayal in a strategy. Betrayal is defined as defecting before your opponent does, so defecting against cooperation [5]. Ta-

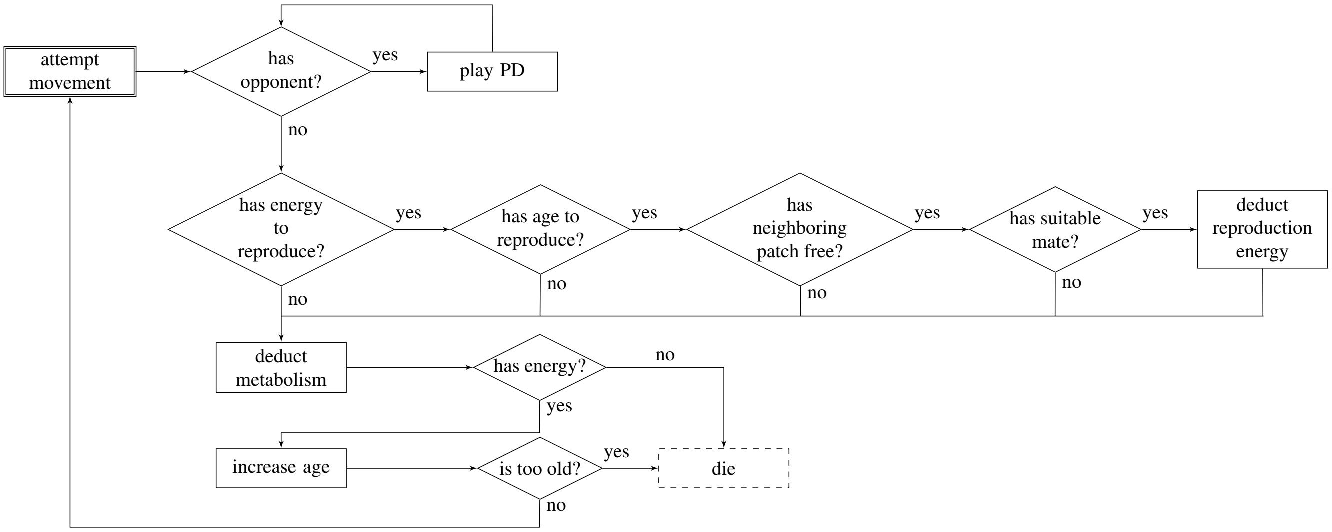

Figure 1: Flowchart of agent life cycle with first action indicated by a double line border

ble 2 shows all indices that allow for betrayal are in the sequence $(0, 1, 3, 5, ...)$, or in the form $\frac{(i^2+i)}{(i-1)^2}$ for the ith index respectively.

The second metric is the punish-cardinality, describing the factor of punishment in a strategy. Punishment, or retaliation, is defined as defecting always and only after your opponent has [5]. From Table 2, these indices are in the sequence $(2, 4, 8, 10, ...)$ or in the form $\frac{2}{(i-1)^2}$ for the ith index.

The thresholds used to determine what group a strategy belongs to for an average cardinality c are defined such that $nasty : c < 0.45$, $balanced : c \geq 0.45 \land c \leq 0.55$ and $nice : c > 0.55$. The interval of balanced strategies is small, because the large majority of strategies would otherwise fall into it due to the average of metrics and the nature of randomly initialised agents.

Neighbourhood types

In order to determine the neighbours of an agent, two different neighbourhood types are considered. The first is the von Neumann neighbourhood, which considers neighbours to be only those agents directly adjacent to the patch of the respective agent [23]. This means an agent can have a maximum of four direct neighbours. The second neighbourhood type is the Moore neighbourhood. This considers all agents in surrounding patches as neighbours [24]. An agent can detect a total of 8 neighbours following this neighbourhood type.

It is important to note that neighbouring locations derived from the neighbourhood type are not only used for finding agents. All interactions depend on the neighbourhood type, such as finding a patch to move to or a free patch to produce offspring.

4.2 Reproduction

Reproduction is used as a tool to evolve successful agents and, over time, eliminate less successful agents. The requirements for reproduction are given by a minimum age and wealth. Once these are met, an agent also has to find a free neighbouring patch and a suitable mate to reproduce with. A selection procedure, meaning a method of determining which agents participate in reproduction, is often used in other work [3,13]. However, by enforcing the given requirements for reproduction, this selection becomes obsolete.

A free patch is derived from the selected neighborhood type (see section 4.1). The best suitable mate is found according to whether the agent prefers *domestic* reproduction. If true, the agent finds a mate with the highest wealth that it is the closest related to. The second metric is dominant, meaning an agent prefers a mate who is closely related and has a lower wealth over the opposite. The order of closest relations according to the type of group(s) used is given in Table 3 below. If an agent prefers *wealthy* reproduction, the best suitable mate is simply the neighbouring agent with the highest wealth.

Producing new offspring has a predetermined cost which is deducted from both parents. Note that a replacement technique is not applied here, meaning parents continue to participate in IPD games. This is contrary to non-spatial evolutionary simulations where produced offspring replace their parents [13, 25]. The reason for this is to allow successful agents to produce more offspring. Depending on the group(s) used, the ancestors (regarding kin) and clan will also have to be determined. Deriving new strategies from parents is always done using the genetic algorithm discussed in section 4.3.

Group(s)	Ordering (closest to furthest)
kin	in-kin, out-kin
clan	in-clan, out-clan
kin and clan	in-kin and in-clan, in-kin and out-clan, out-kin and in-clan, out-kin and out-clan

Table 3: Relations ordering

Remembering ancestors

As a single offspring is produced by two parents, both of their ancestors will have to be remembered. The length of the list of ancestors is limited to emulate restricted memory of agents. In the case that the limit is reached, oldest members are forgotten first. To ensure fair division of ancestors from the two parents, each agent remembers them separately, constrained by half the size of the total ancestors. The separate ancestors are derived from the union of half of the respective parent's ancestors. In order to make room for the new ancestor, namely the parent itself, the oldest member of the ancestors of parent two from the respective parent is removed. Parent two is always chosen over a random parent, as the order of parents is already random. An example of new ancestors of maximum length of 8 for offspring o from parents p_1 and p_2 is shown in Figure 2.

$$\begin{array}{llcccc}
p_1: & ancestors_{p_1}: & \boldsymbol{a} & \boldsymbol{b} & c & d \\
 & ancestors_{p_2}: & \boldsymbol{e} & f & g & h \\
 & & & & & \\
p_2: & ancestors_{p_1}: & \boldsymbol{i} & \boldsymbol{j} & k & l \\
 & ancestors_{p_2}: & \boldsymbol{m} & n & o & p \\
\hline
o: & ancestors_{p_1}: & p_1 & a & b & e \\
 & ancestors_{p_2}: & p_2 & i & j & m
\end{array}$$

Figure 2: Example new ancestors with maximum length of 8 for offspring o created by parents p_1 and p_2

Determining new clan

Since agents cannot belong to more than one clan, determining a new clan for offspring has to be done differently. The new clan is trivially defined as the clan of the first parent. This is fair in the sense that the order of parents in reproduction is random, which will ensure the same results multiple runs.

4.3 Genetic Algorithm

To derive new strategies from parents, this paper considers four genetic operators affecting reproduction occurring chronologically: cloning, crossover, mutation and inversion. These are, depending on the group(s) present, applied separately, but on all strategies of the two parents. Applying them separately means that a certain mutation in the first strategy of offspring is not necessarily also present in the second strategy. Four operators are considered as they are sufficient to explore the search space of strategies and avoid premature convergence. An agent's strategies can be seen as its chromosomes, consisting of genes for moves.

Clone

The clone operator simply copies one of the chromosomes from the two parents. The respective chromosome is determined by a randomly generated direction value.

Crossover

The crossover operator is defined as a 2-point crossover occurring with probability P_c. This combines the genes from two parent chromosomes by randomly selecting two crossover points. The direction value generated for the clone operator is used to determine what parts of which chromosome are copied with respect to the crossover points. An example of a 2-point crossover with one-chromosome strategies is given in Figure 3. Note that if the probability P_c of this operator performing is zero, simple natural selection will occur through cloning "successful" chromosomes.

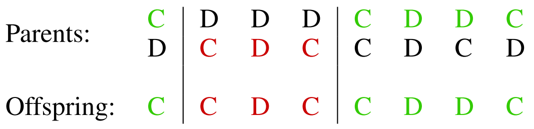

Parents:	C	D	D	D	C	D	D	C
	D	C	D	C	C	D	C	D
Offspring:	C	C	D	C	C	D	D	C

Figure 3: Example 2-point crossover in given direction

Mutation

The mutation operator is defined as a bit-flip of new genes copied from parent chromosomes with probability P_m. This is performed with low probability ($P_m \approx 0.01$), as this provides an effective way to explore solutions without preventing stability over time.

Inversion

The inversion operator is defined as inverting the entire chromosome of new offspring with probability P_i. This is performed with very low probability ($P_i \approx 0.0001$) and serves as a way to both explore more drastic solutions and strategy resistance. The latter use prevents sub-optimal solutions from dominating during earlier stages of a simulation.

5 Experimental Setup and Results

The proposed model was implemented in NetLogo [26], and the results of simulations were evaluated over 5 runs. Running simulations five times already showed statistical significant results, see section 6.2. Every round (or tick), all agents execute one iteration of their main life cycle as depicted in Figure 1. Values for the given plots were computed at the end of every round. In all simulations, a threshold of 2000 ticks was sufficient for low variance. This variance was measured by grouping the strategies (and therefore population) into nice, balanced and nasty groups and examining their ratio over the last 50 ticks. Because mutation and inversion influence strategies, the variance in population ratios can never be fully eliminated. Simulations running for 4000 ticks were also examined to ensure true convergence of the final results.

A grid of 100 × 100 patches was selected for the experiments, as this showed a good balance between simulation time and total number of possible agents. A population density of $\frac{1}{4}$ was used, as it proved to be minimal for sustaining life through reproduction. A relative low density was selected, because it allows more room for search space exploration in strategies. For the given configuration of the grid, this means 2500 agents were initialized at the start. The first agents are assigned random positions, strategies, an empty list of ancestors, and their own distinct clan. A random age drawn from a uniform distribution and initial wealth equal to the reproduction cost are also set. The default set of model parameters kept constant throughout all simulations are depicted in Table 4.

Description	Value
Payoff matrix	$R = 5$, $P = -5$, $T = 6$, $S = -6$
Probabilities genetic operators	$P_c = 0.95$, $P_m = 0.01$, $P_i = 0.0001$
Maximum age	20
Minimum reproduction wealth	10
Minimum reproduction age	0
Reproduction cost	6
Cost of living	0

Table 4: Default set of model parameters

To properly answer the main question posed by this paper, the evolution of cooperation is evaluated separately for different group types. In order to put these results into perspective, base results of a simulation without groups are first presented. For each of the given group types, two simulations with two different agent sets are considered. The first contains agents preferring domestic reproduction, and the second preferring wealthy reproduction. Details of specific group-related parameters are given in the subsections below.

Two metrics are used to measure cooperation. The first examines population levels divided into nice, balanced and nasty groups, as discussed in section 4.1. The population levels are normalized with respect to the maximum population size. The second measures average cooperation levels of the games played per round. The average cardinality or "niceness" of strategies used for the first metric is also given. The reason for adding levels of game cooperation is that, since agents rarely exercise every move in their strategy, it is not appropriate to rely on strategy cooperation alone. Using the two metrics together gives a better and sufficient representation of the actual presence of cooperation.

5.1 No groups

For a simulation without groups, balanced strategies seem to dominate in Figure 4. However, cooperative strategies strategies depicted are still preferred over defective ones in Figure 5. It is important to note that the initial majority of balanced strategies due to random initialisation does not cause them to dominate in the long term. They dominate because they are simply more effective. Evolving a population initially consisting of unconditional cooperate strategies, converged towards the same population ratios and cooperation levels.

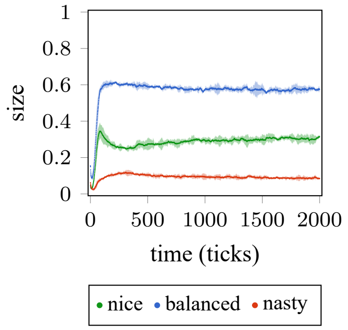

Figure 4: Levels of nice, balanced and nasty populations with no groups in vonNeumann neighbourhood

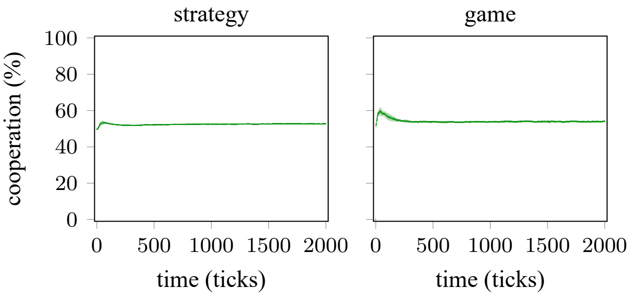

Figure 5: Cooperation with no groups in vonNeumann neighbourhood

5.2 Kin

As discussed in section 4.2, the number of ancestors an agent can remember is variable. To investigate the impact of changing this memory, two experiments for domestic reproduction were conducted with maximum ancestors of 4 and 40 respectively. An increase with a factor of 10 was used to ensure a difference in results would emerge if a causation exists.

Domestic reproduction

A clear difference in cooperation levels emerges depending on the size of ancestors, see Figures 6 and 8. Game cooperation levels are significantly higher for a lower number of ancestors (Figure 7). This is counter-intuitive, as both contests and the total number of games for in-kin are lower (by 36%) due to smaller kin. Longer contests lead to better evolved strategies, which would result in more cooperation for in-kin games, as shown in Figure 9.

The reason for increased cooperation with smaller kin is that group-strategies are not defined by their group, and can differ amongst its members. Kin defined by many ancestors is functionally divided into smaller groups with distinct in-kin strategies. Over time, this restricts the rise of cooperation present for a smaller number of ancestors. The increase of defection with 4 ancestors can be similarly explained. Due to higher cooperation levels within kin, agents can afford, meaning without dying, to take more risk for the temptation payoff. Since in-kin cooperation with 40 ancestors is lower, the risk agents can take for defection is also lower.

The difference in strategy and game cooperation for 4 ancestors can be explained by the repeated patterns that occur in contests. For example: a three round in-kin history containing 2 cooperate moves is much more likely than other histories. Over time, this causes those respective moves to reinforce the probabilities of the corresponding histories occurring. The same is also the case for out-kin strategies, but due to less games also reinforced less.

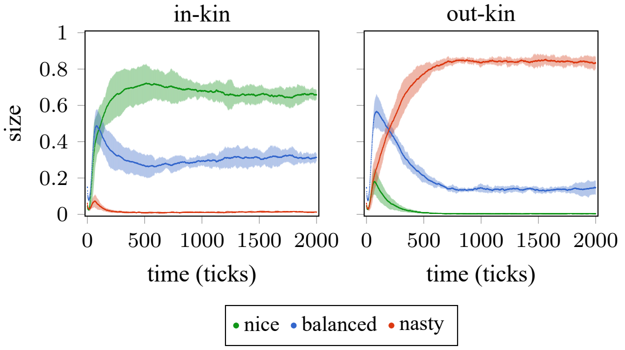

Figure 6: Levels of nice, balanced and nasty populations preferring domestic reproduction with kinship with max 4 ancestors in vonNeumann neighbourhood

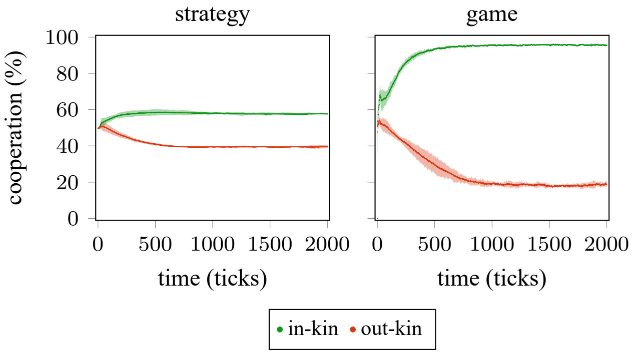

Figure 7: Cooperation preferring domestic reproduction with kinship with max 4 ancestors in vonNeumann neighbourhood

Wealthy reproduction

The difference in results between domestic and wealthy reproduction in Figure 10 are small, but still statistically significant (see section 6.2). Wealthy reproduction results in slightly less cooperation and defection for in-kin and out-games respectively, see Figure 11. The small difference is likely due to the restricted neighbourhoods: the wealthiest neighbouring agent is often also in the same kin.

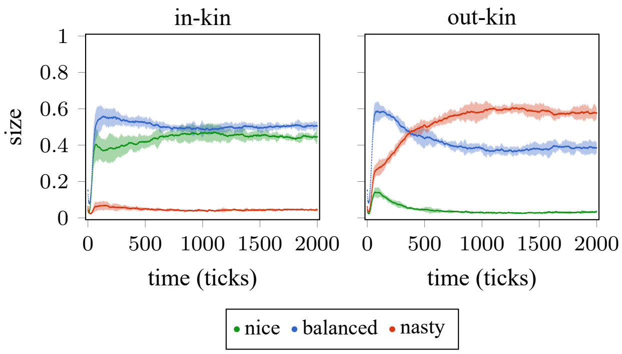

Figure 8: Levels of nice, balanced and nasty populations preferring domestic reproduction with kinship with max 40 ancestors in vonNeumann neighbourhood

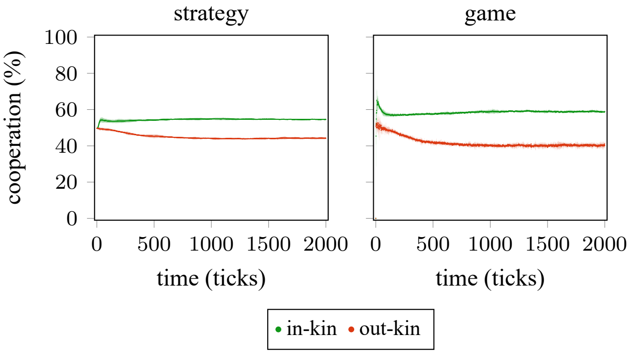

Figure 9: Cooperation preferring domestic reproduction with kinship with max 40 ancestors in vonNeumann neighbourhood

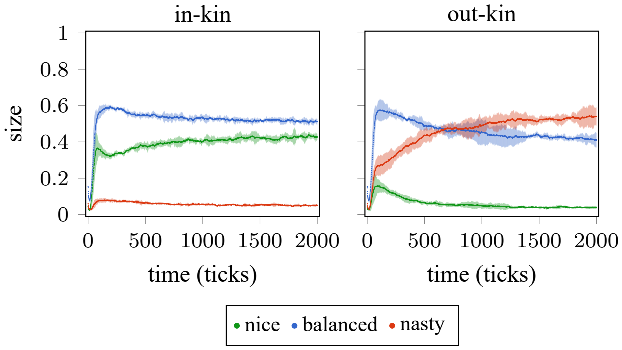

Figure 10: Levels of nice, balanced and nasty populations preferring wealthy reproduction with kinship with max 40 ancestors in vonNeumann neighbourhood

5.3 Clan

Since clan size can not be restricted by a parameter, cooperation will simply be evaluated for agents preferring domestic and wealthy reproduction.

Domestic reproduction

Agents preferring domestic reproduction cooperate more often in clans, and are likely to defect for non-clan members (Figures 12 and 13). Defection for out-group games is higher than cooperation for in-group games. This is due to the "punishable" effect of individual cooperation, directly affecting survival.

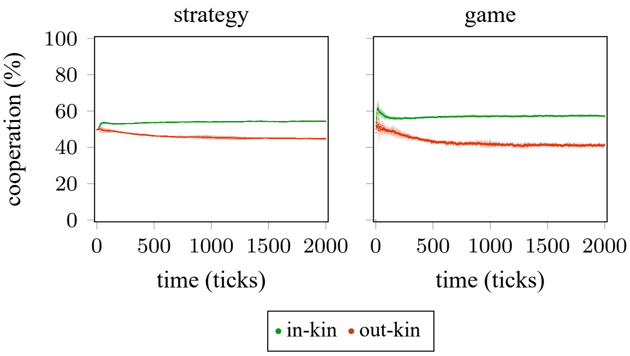

Figure 11: Cooperation preferring wealthy reproduction with kinship with max 40 ancestors in vonNeumann neighbourhood

Levels of game cooperation are slightly higher than strategy cooperation. This can be explained similarly to kin simulations: contests and histories fall into patterns only exercising certain strategy indices.

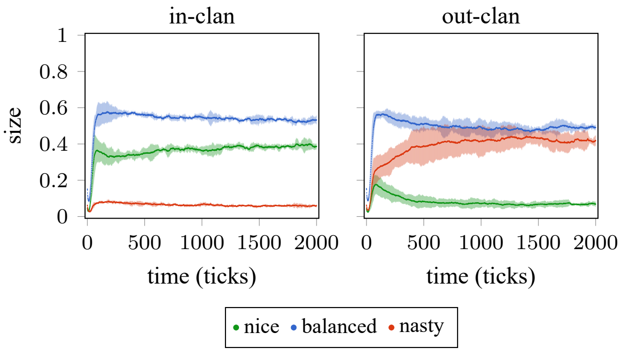

Figure 12: Levels of nice, balanced and nasty populations preferring domestic reproduction with clans in vonNeumann neighbourhood

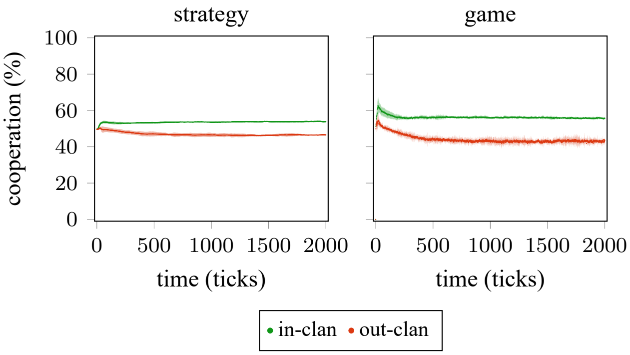

Figure 13: Cooperation preferring domestic reproduction with clans in vonNeumann neighbourhood

Wealthy reproduction

Population levels are different for agents preferring wealthy reproduction, as shown in Figure 14. Balanced strategies seem to dominate for both in- and out-clan games, with nice and nasty populations next. As opposed to domestic reproduction, out-clan defection is not higher than in-clan cooperation (Figure 15). Similarly to Figure 11, the expected results should not deviate much from clans with wealthy reproduction. Atomic clan membership would increase out-clan games over simulations with large kin, as agents can be members of multiple kin. However, large clans present in spatial zones reduce the probability of finding out-clan opponents significantly.

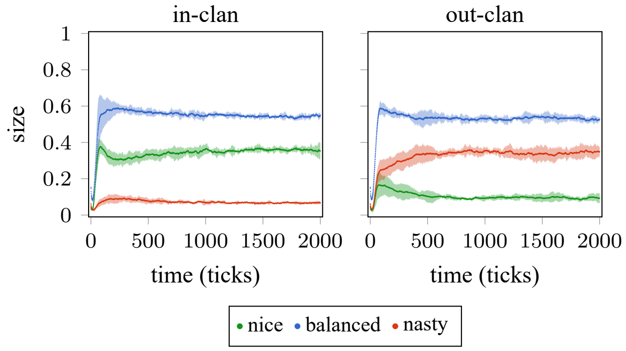

Figure 14: Levels of nice, balanced and nasty populations preferring wealthy reproduction with clans in vonNeumann neighbourhood

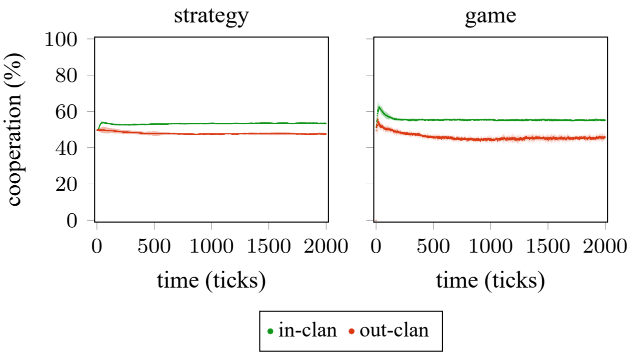

Figure 15: Cooperation preferring wealthy reproduction with clans in vonNeumann neighbourhood

5.4 Kin and Clan

Since kin and clans combined allow for more insight into their relative cooperation levels, simulations operating in Moore neighbourhoods have also been examined. A constant number of 40 ancestors is used, so the difference between in-clan and in-kin games is smaller.

Domestic reproduction

As depicted in Figures 16 and 17, expected similar cooperation levels derived from separate kin- and clan simulations develop. Agents are nicer to those they are closer related to and defect more often to strangers. One interesting aspect is that agents defect more often to out-kin and in-clan opponents than out-kin and out-clan opponents. The reason for this is the restriction of vision of agents due to the vonNeumann neighbourhood used. With clans being very large and spatially present, in-clan games make up the large majority (92%) of the games played.

The Moore neighbourhood offers a solution for this by increasing the vision of agents. Figures 18 and 19 show that cooperation levels are much better defined. While the number of out-kin and out-clan games is still small, the defection levels are much closer to out-kin and in-clan games.

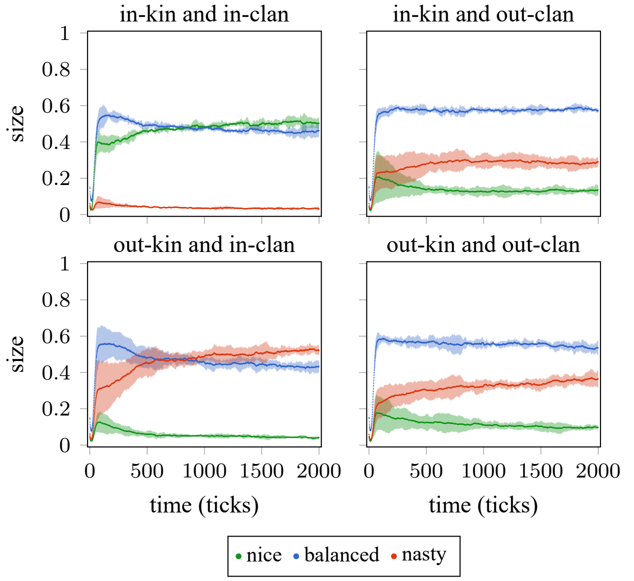

Figure 16: Levels of nice, balanced and nasty populations preferring domestic reproduction with kinship with max 40 ancestors and clans in vonNeumann neighbourhood

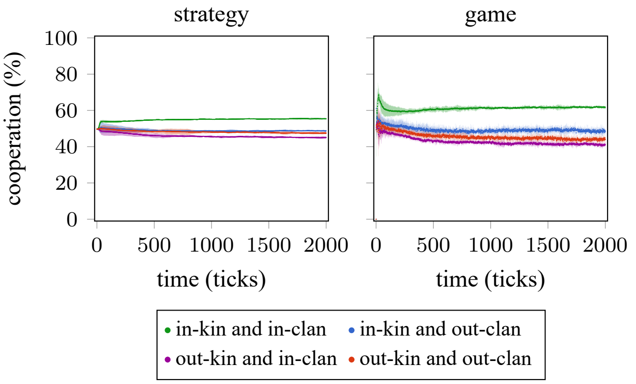

Figure 17: Cooperation preferring preferring domestic reproduction with kinship with max 40 ancestors and clans in vonNeumann neighbourhood

Wealthy reproduction

Agents preferring wealthy reproduction shown in Figures 20 and 21 show restricted cooperation levels, similar to the results for wealthy reproduction for kin- and clan-only simulations. When kin and clans are combined, the probability of finding mates that do not belong to the same group is smaller. This reinforces the restriction of developing cooperation levels.

5.5 Other experiments

The experiments above had constant values for the grid size, a set of model parameters, and the cooperation was measured in discrete time. Experiments with different values and measured in continuous time have also been conducted to ensure keeping them constant is justified. Justified here means that they affect groups to the same degree.

Lowering all payoffs meant the reproduction cost and minimum wealth also had to be reduced to ensure survival. This resulted in the same defection levels, but less cooperation. Lowering the payoff for

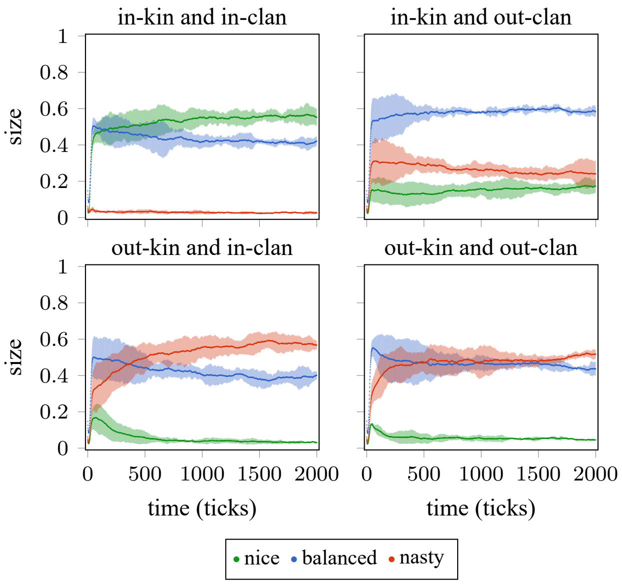

Figure 18: Levels of nice, balanced and nasty populations preferring domestic reproduction with kinship with max 40 ancestors and clans in Moore neighbourhood

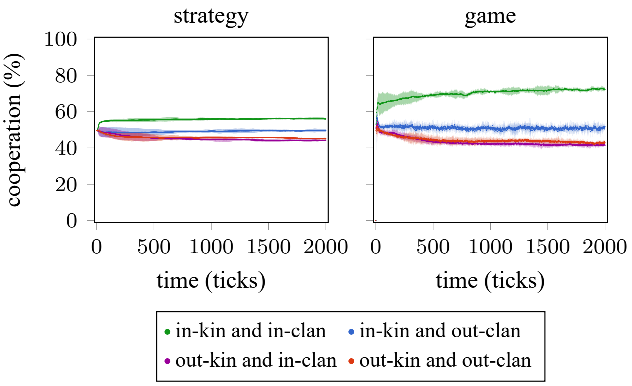

Figure 19: Cooperation preferring domestic reproduction with kinship and clans in Moore neighbourhood

mutual cooperation and decreasing the negative payoff for mutual defection caused the same effect.

Increasing the cost of living increased noise drastically at the beginning of simulations, as agents need to cooperate to survive. As population density increases, levels of cooperation decrease and converge to the same results presented in this paper. The cost of living was also increased to verify non negative payoffs produced the same results.

Decreasing P_c down to zero lowered cooperation and increased defection slightly. Increasing values for P_m and P_i caused more noise, and above a certain threshold ($P_c > 0.2$ and $P_i > 0.0005$) prevented stabilisation of strategies.

Lowering the maximum age increased the harshness of the environment and challenged survival. Increasing or removing the maximum led to the same results, but increased convergence time up to a factor of 10.

Lowering reproduction requirements and cost caused more noise and showed inconsistent results over multiple runs if low enough. Increasing these values significantly increased the harshness of the environment, and caused extinction if too high.

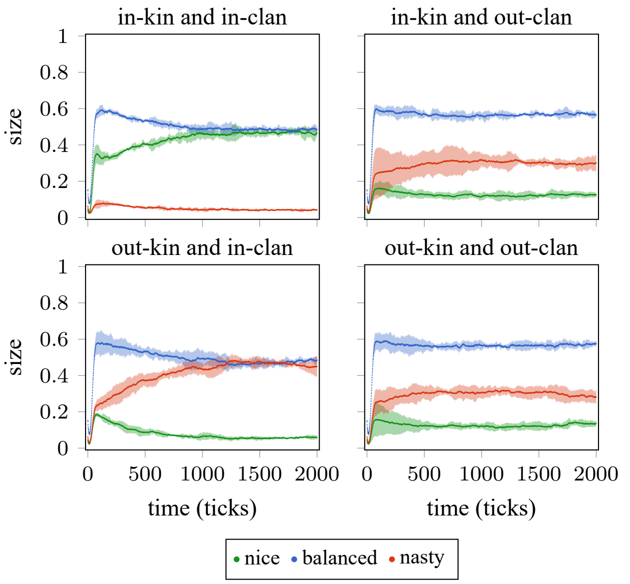

Figure 20: Levels of nice, balanced and nasty populations preferring wealthy reproduction with kinship and clans in vonNeumann neighbourhood

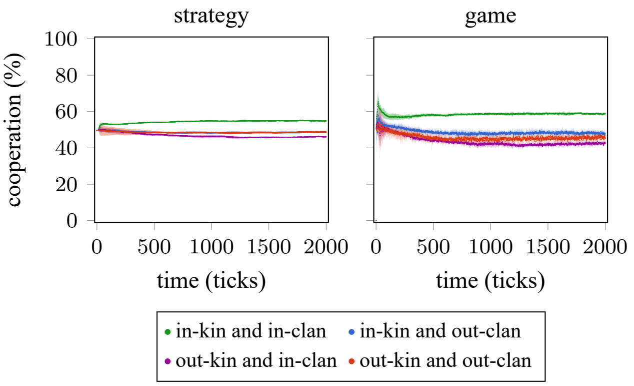

Figure 21: Cooperation preferring wealthy reproduction with kinship and clans in vonNeumann neighbourhood

Simulations running in continuous time (as implemented in NetLogo) produced the same results, but with a slightly higher total execution time.

6 Responsible Research

This section will reflect on the ethical and epistemic aspects of the results produced by this paper.

6.1 Ethics

Choosing agent-based modeling and simulation as a tool to examine the evolution of cooperation is an effective way to avoid premature need for living (human) subjects. Such subjects will also never prefer domestic reproduction for kin as agents did. The reason for analysing agents with this preference is that, in reality, “niceness” is not hard-wired in genes. Domestic reproduction can be seen as the vertical cultural transmission from parents to offspring [27]. Mutation and interaction with strangers can be seen as horizontal cultural transmission affecting strategies [28]. These assumptions are made to show a simple model is sufficient to show the persistence of cooperation in groups. Following experiments could then consider real subjects.

6.2 Reproducibility

The results in this paper can be easily reproduced by implementing the proposed model and running it for discussed parameter sets, averaged over 5 runs. As discussed in section 5, running the simulations 5 times sufficiently indicated the results. A standard t-test assuming unequal variances has been conducted to verify this. Final values of both population and cooperation levels for different simulations were evaluated. The hypothesized mean difference was assumed to be zero and the significance level was set at 0.05. P-values derived from comparing different groups were very small ($< 10^{-3}$). P-values derived from comparing different reproduction preferences for the same group type were bigger, but still under the given threshold (< 0.02). The latter p-values are higher due to the differences in results for the same group type often being smaller.

7 Discussion

The aim of this paper is to investigate how cooperation evolves when group-based agents operate in an evolving spatial environment. This section draws conclusions from the results presented in section 5 with respect to the DPD model [15] and genetic algorithms for clans [21]. The drawn conclusions will be put into broader context by interpreting them based on the real world.

For all group types, cooperation levels seem to be higher for closely related opponents, and lower for the opposite. Nicer in-group strategies and less nice out-group strategies are also preferred. The related parameters used in the simulations in section 5 are very similar to those presented in the DPD model [15]. This work only considers unconditional cooperate and defect strategies without groups, but shows the same persistence of cooperation. The base results with no groups presented in section 5.1 show that balanced strategies defeat cooperative ones, but this is due to them being more effective. Cooperation levels in related work considering a non-spatial model are also higher, but to a larger degree [13]. The reason for this is that the spatial model proposed by this paper restricts agents' opponents to those in neighbouring patches, which prevents uniform evolution of strategies. The non-spatial model contrasts with regard to the defection levels of foreign contests. Whereas this paper shows defection increases with longer contests, the non-spatial model shows the opposite: more cooperation for longer foreign contests.

The main restriction, or feature, of a spatial environment is that agents operate within a given neighbourhood. As discussed in section 5, this neighbourhood often limits the evolution of cooperation. It is important to note that this limitation is not necessarily bad. Simulations running in Moore neighbourhood in section 5.4 show that strategies are better defined, which indicates that increasing this vision further will ultimately lead to the same results for a non-spatial environment. Namely, a non-spatial environment is no different from a spatial one where agents have unlimited vision.

Due to agents operating in restricted neighbourhoods, their environment develops subgroups. These subgroups are spatial zones defined by similar strategies. When a group of neighbours often interact with each other, their strategies evolve together. Through reproduction, such smaller groups develop similar strategies which further reinforce the subgroup. In larger groups, these derived subgroups are likely the reason cooperation between all group members is restricted. This restriction of cooperation results in less defection for out-group opponents, as a degree of cooperation is needed to survive. In non-spatial models from related work, the decrease of defection for foreign contests in larger groups is also present [13].

Looking at the "real" world, such subgroups are present everywhere. The concept belongingness, or the human need to be a member of a group, is an important aspect of human nature [29]. Today's societies are divided into several layers of groups, some with higher cohesiveness (or cooperation) than others. Even societies themselves can be seen as large groups consisting of a complex network of subgroups.

8 Conclusions and Future Work

In this paper, the effects of group-based agents on iterated prisoner's dilemma strategies in evolving spatial environments were examined. Two modifications of the DPD model are proposed. The first generalises strategies using a genetic algorithm that applies four genetic operators. The second extends it with kin and clans. Both von-Neumann and Moore neighbourhoods are considered in the spatial environment.

The main hypothesis this paper drafted was that group-based agents exhibiting cooperative domestic strategies, meaning for in-group opponents, and defective foreign strategies, meaning for out-group opponents, win in the iterated prisoner's dilemma in evolving spatial environments. Three distinct conclusions can be drawn from the experiments conducted to test this hypothesis:

1. Strategies evolved through domestic reproduction exhibit more cooperation for in-group opponents and more defection for out-group opponents.
2. Strategies evolved through wealthy reproduction exhibit the same increase and decrease of cooperation as domestic reproduction, but to a smaller degree.
3. An evolving spatial environment with group-based agents develops subgroups, defined by similar strategies and restricting group-wide cooperation.

The main conclusion drawn by combining these findings together is that agents with cooperative domestic strategies and defective foreign strategies win, and there is a positive correlation between group size and restriction of group-wide cooperation. Larger groups develop more subgroups and its group-wide cooperation is therefore more restricted. Smaller groups consist of fewer subgroups, thus exhibiting significantly higher in-group cooperation and out-group defection. This answers the main question and confirms the hypothesis proposed by this paper. It is important to note that this conclusion resolves around the group types used in this paper, which are present in spatial zones. In this case, the group size alone is dominant in restricting group-wide cooperation.

Future work could resolve around considering more complex strategies. There are two ways of increasing their complexity. The first considers a larger number of fixed strategies, which can then be more specific. An example could be agents considering different kin relations based on a distance metric. As mentioned in section 4.2, the ancestors for two parents are stored separately. A distance metric could then determine whether an opponent is a member of (both) the ancestors, and be used to adjust an agent's moves accordingly. This allows for a closer analysis of relations within kin. A second way of increasing the complexity of strategies is to make them kin- or clan-specific. An example could be clans considering separate strategies for other clans. The establishment of trust between individual clans can then be analysed. An important focus for increasing the complexity of agents should be ensuring significant interaction for each of their strategies, as insufficient usage will prevent consistent convergence towards optimal strategies. An easy solution would be to both increase the total number of agents through a larger grid size, and the size of neighbourhoods. The latter can be achieved by either increasing the range of vision of agents or using different (irregular) neighbourhoods.

References

[1] A. Carter, "Evolution and the problem of altruism," *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, vol. 123, no. 3, pp. 213–230, 2005.

[2] A. Rapoport and A. M. Chammah, *Prisoner's Dilemma*. University of Michigan Press, Ann Arbor, 1965.

[3] S. Bukhari, "Using genetic algorithms to develop strategies for the prisoners dilemma," Dec. 2005.

[4] S. Kuhn, "Prisoner's dilemma," in *The Stanford Encyclopedia of Philosophy* (E. N. Zalta, ed.), Metaphysics Research Lab, Stanford University, winter 2019 ed., 2019.

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

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

[7] R. Chiong and M. Kirley, "Effects of iterated interactions in multiplayer spatial evolutionary games," *IEEE Transactions on Evolutionary Computation*, vol. 16, no. 4, pp. 537–555, 2012.

[8] R. Axelrod, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press, 1997.

[9] J. Grujić, C. Gracia-Lázaro, M. Milinski, D. Semmann, A. Traulsen, J. A. Cuesta, Y. Moreno, and A. Sánchez, "A comparative analysis of spatial prisoner's dilemma experiments: Conditional cooperation and payoff irrelevance," *Scientific Reports*, vol. 4, p. 4615, Apr. 2014.

[10] R. Axelrod, *Genetic Algorithms and Simulated Annealing*, ch. The Evolution of Strategies in the Iterated Prisoner's Dilemma, pp. 32–41. Morgan Kaufman, Los Altos, CA, 1987.

[11] M. Nowak and R. May, "Evolutionary games and spatial chaos," *Nature*, vol. 359, pp. 826–829, Oct. 1992.

[12] D. Hirshleifer and E. Rasmusen, "Cooperation in a repeated prisoners' dilemma with ostracism," *Journal of Economic Behavior Organization*, vol. 12, no. 1, pp. 87 – 106, 1989.

[13] B. A. Julstrom, "Clans and cooperation in the iterated prisoner's dilemma," in *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*, GECCO '12, (New York, NY, USA), p. 1463–1464, Association for Computing Machinery, 2012.

[14] J. Li and G. Kendall, "A strategy with novel evolutionary features for the iterated prisoner's dilemma," *Evolutionary computation*, vol. 17, pp. 257–74, Feb. 2009.

[15] J. Epstein, "Zones of cooperation in demographic prisoner's dilemma," *Complexity*, vol. 4, pp. 36–48, Jan. 1998.

[16] P. E. Smaldino and J. C. Schank, "Movement patterns, social dynamics, and the evolution of cooperation," *Theoretical Population Biology*, vol. 82, no. 1, pp. 48 – 58, 2012.

[17] W. Radax and B. Rengs, "Demographic prisoner's dilemma - clan/kinship extension." Visited on 04/06/2020.

[18] W. Radax and B. Rengs, "Prospects and pitfalls of statistical testing: Insights from replicating the demographic prisoner's dilemma," *Journal of Artificial Societies and Social Simulation*, vol. 13, Oct. 2010.

[19] W. Radax and B. Rengs, "iterated pd demographic 107 artificial." Visited on 04/06/2020.

[20] D. E. Goldberg and J. H. Holland, "Genetic algorithms and machine learning," *Machine Learning*, vol. 3, pp. 95–99, Oct. 1988.

[21] B. A. Julstrom, "The development of cooperation in genetic algorithms for the iterated prisoner's dilemma," Mar. 2000.

[22] A. Errity, "Evolving strategies for prisoner's dilemma," 2003.

[23] E. W. Weisstein, "von neumann neighborhood. From MathWorld—A Wolfram Web Resource." Visited on 01/06/2020.

[24] E. W. Weisstein, "Moore neighborhood. From MathWorld—A Wolfram Web Resource." Visited on 01/06/2020.

[25] D. B. Fogel, "Evolving behaviors in the iterated prisoner's dilemma," *Evol. Comput.*, vol. 1, p. 77–97, Mar. 1993.

[26] U. Wilensky, "Netlogo," http://ccl.northwestern.edu/netlogo/, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL, 1999.

[27] L. L. Cavalli-Sforza and M. W. Feldman, *Cultural transmission and evolution: A quantitative approach*. Princeton University Press, 1981.

[28] J. Duffy *Southern Economic Journal*, vol. 64, no. 3, pp. 791–794, 1998.

[29] K.-A. Allen, *The Psychology of Belonging*. June 2020.