

Replication strategies and the evolution of cooperation by exploitation

Markus Brede¹ and Simon Tudge¹

¹University of Southampton, Southampton, SO17 1BJ
markus.brede@soton.ac.uk

Abstract

Introducing the concept of replication strategies this paper studies the evolution of cooperation in populations of agents whose offspring follow a social strategy that is determined by a parent's replication strategy. Importantly, social and replication strategies may differ, thus allowing parents to construct their own social niche, defined by the behaviour of their offspring. We analyse the co-evolution of social and replication strategies in well-mixed and spatial populations. In well-mixed populations, cooperation-supporting equilibria can only exist if the transmission processes of social strategies and replication strategies are completely separate. In space, cooperation can evolve without complete separation of the timescales at which both strategy traits are propagated. Cooperation then evolves through the presence of offspring-exploiting defectors whose presence and spatial arrangement can shield clusters of pure cooperators.

Introduction

Actions that are in the interest of the group but not necessarily to the immediate benefit of the individual are widely observed in the social and biological sciences. Understanding the emergence and sustainability of such altruism or cooperation still poses major challenges to evolutionary game theory and the recent decades have seen very active research in this field. For instance, a recent review article classified five different mechanisms that support altruism (Nowak, 2006). Here, we are mainly interested in one of them: network reciprocity, cf. (Szabó and Fath, 2007; Perc and Szolnoki, 2010) for recent reviews.

Network reciprocity summarizes effects that result from constrained interactions in structured populations in which agents interact with fixed and typically rather small sets of permanent neighbours. In this way interactions between parents and offspring are favoured, i.e. cooperation is supported through positive assortment of strategies. The literature about evolutionary games in structured populations goes back to the seminal paper of Nowak (Nowak and M., 1992) in which spatial games were introduced, observing and describing chaotic patterns of strategies in space. The work was extended in several ways to, e.g., include effects

of noise (Szabó and Toke, 1998) and asynchronicity (Huberman and Glance, 1993) in strategy propagation. Recent research has mainly focused on the evolution of cooperation in population structures modelled by complex networks, finding, e.g. that heterogeneous networks give a strong boost to cooperation (Santos et al., 2006). The latter findings have been extended to evolutionary models on regular graphs in which there is some heterogeneity in agent's abilities to generate payoff. Examples of studies in this direction are (Szolnoki and Szabó, 2007; Perc and Szolnoki, 2008; Brede, 2011a), but also the recent work on teaching and learning (Szolnoki and Perc, 2008; Tanimoto and Yamauchi, 2012). In the latter line of research agents are classified into two groups: (i) teacher agents with an enhanced ability to pass on strategies and (ii) learner agents with reduced abilities to pass on strategies. The co-evolutionary dynamics of fast and slow strategy spread can then generate phases in which cooperation can survive much beyond parameter regimes in which cooperation is supported by network reciprocity alone (Brede, 2013a).

Common to this large bulk of work on cooperation and network reciprocity is the assumption that offspring (in a biological context) or followers (in a social context) adopt exactly the same strategy as parents (or leaders). In fact, one might surmise that this assumption is crucial to allow for positive assortment which enables support for cooperation through network reciprocity. In this paper we introduce a more general framework that aims to challenge this hypothesis and explore its boundaries. We distinguish two traits that describe agent behaviour. The first is the typical social strategy that describes an agent's behaviour in the social dilemma game under consideration. The second is a replication strategy, i.e. a strategy that an agent will pass on as a social strategy to its offspring. In this way every agent is characterised by a tuple (s, e) : a social strategy s and a replication strategy e through which it can determine its offspring's social behaviour. Notably, the social strategy and the replication strategy of an agent can be different: It might be in the interest of an agent to surround itself with offspring (or followers) that are of a different type than itself. Hence,

agents may surround themselves by un-like types, questioning the role of positive assortment by network reciprocity.

One might also interpret our framework as a very simple model of social niche construction (Powers et al., 2011). The term social niche construction was recently introduced to describe a situation in which agents can evolve preferences for the social group they interact with. Using the example of preferences for group size it was then demonstrated that co-evolution of such preferences and social strategies can naturally support cooperation. In our context here, by their replication strategy, agents can influence the environment in which they live and thus improve their chances to generate payoff in competitive games. Using the often-studied framework of the prisoner’s dilemma game, we will explore under which circumstances such a simple co-evolutionary model can allow for additional support for cooperative strategies.

Real-world inspiration for the above assumption of differences between social and replication strategies is not hard to come by. For instance, in models of teaching and learning the above framework allows for situations in which teachers can teach strategies different from their own. Arguably, this is a more realistic and general framework than the one considered in previous work. In a biological context one might interpret the model as a simple model of cell differentiation.

The present work thus follows in a line of recent advances in the understanding of the co-evolution of individual-level traits and cooperation (Szolnoki et al., 2009; Powers et al., 2011; Perc and Wang, 2010; Brede, 2011b, 2013b).

The organization of the paper will be as follows. We start with a detailed description of the model framework and then describe and explain results in the section thereafter. The paper concludes by a summary and discussion section that puts our main results into context and discusses implications and future work.

Model

In more detail, we consider the following model of an evolutionary one-off prisoner’s dilemma in space. A set of N agents are associated with the sites of a graph whose links define interactions and directions of strategy propagation. In case of experiments in well-mixed populations this social networks is a complete graph, otherwise we perform experiments on an $L \times L$ square lattice with von Neumann neighbourhoods. Agents are characterized by two strategy traits, a social strategy trait $s \in \{0, 1\}$ and a replication strategy trait $e \in \{0, 1\}$. We use the convention that state “0” corresponds to a strategy of pure defect and state “1” corresponds to pure cooperate. Agents play a prisoner’s dilemma with payoff matrix parametrized in the conventional form

$$\begin{pmatrix} R & S \\ T & P \end{pmatrix} = \begin{pmatrix} 1 & -r \\ 1+r & 0 \end{pmatrix}, \quad (1)$$

such that the parameter r characterizes the toughness of the game setting. In Eq. (1) R stands for the reward for mutual

cooperation, S for the “sucker’s” payoff, T for the temptation to defect, and P for the punishment for mutual defection. A small $r \ll 1$ corresponds to very mild dilemma settings, whereas $r \rightarrow 1$ characterizes very tough dilemmas. Hence, we distinguish four strategies: (i) cooperators who want their offspring to cooperate ($s = 1, e = 1$), (ii) cooperators who want their offspring to defect ($s = 1, e = 0$), (iii) defectors who want their offspring to cooperate ($s = 0, e = 1$) and (iv) defectors who pass on defection to their neighbours ($s = 0, e = 0$). This model may easily be extended by including context-dependent inheritance, i.e. the offspring determining trait would then depend on the social strategy currently played, but we reserve a thorough investigation of this case for future work and concentrate on the simplest setup in this paper.

In the following we will also consider the impact of various timescales in the evolution of social and replication strategies. The spread of both strategies might occur on separate or similar timescales. In the case of joint strategy pass, an agent will adopt the desired social strategy of a parent as well as its replication strategy. In case of disparate pass, an agent might either adopt the parents’ desired social strategy or its replication strategy. To distinguish these cases and to investigate the effects of disjoint strategy pass we introduce a probabilistic framework for the spread of strategies: With probability p_s only the social strategy is imposed, otherwise, with probability p_a only the replication strategy is passed on, and in the remaining cases (i.e. with probability $p_p = 1 - p_s - (1 - p_s)p_a$) the social strategy is imposed and the replication strategy passed on. The timescales of the spread of social and replication strategies are then given by $T_s = 1/(p_s + p_p)$ and $T_a = 1/(p_a + p_p)$.

Hence, our evolutionary simulations consist of an asynchronous process iterating the following steps:

- Seed all agents with randomly chosen initial social and replication strategies.
- Randomly pick a focus agent, say i , and choose a reference agent j from one of its four von Neumann neighbours at random.
- Evaluate game interactions of the focus agent with its neighbours to determine its accumulated payoff $\pi_i^{(\text{game})}$ and follow the same procedure to calculate the accumulated payoff $\pi_j^{(\text{game})}$ of the reference agent from interactions with its neighbours.
- After evaluating game payoffs, a cost c is deducted from payoffs of agents who attempt to spread a strategy different from their social strategy, i.e.

$$\pi_i = \pi_i^{(\text{game})} - c(1 - \delta_{s_i e_i}), \quad (2)$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. A cost $c > 0$ accounts for the fact that imposing social strategies different

from your own might involve a costly effort to ‘convince’ the opponent. Unless otherwise stated experiments are carried out with $c = 0$ and the influence of a non-zero cost is only evaluated at the end of the paper.

- In a next step, a focus i agent will adopt the strategy of the reference agent j with a likelihood that depends on the difference in payoffs, i.e.

$$P(j \rightarrow i) = \frac{\exp(\pi_j/\kappa)}{\exp(\pi_j/\kappa) + \exp(\pi_i/\kappa)}. \quad (3)$$

In the above equation the parameter κ introduces noise in the replication process, the larger κ , the larger the chance for inferior strategies to spread. In all following simulations we set the noise level to a relatively large value of $\kappa = 1$. This choice is motivated by reasons of computational feasibility, because the evolutionary dynamics becomes very slow for low levels of noise when the timescales of cluster expansion are dominated by the timescales of change of local configurations of $s = 0, e = 1$ defectors surrounded by cooperators at the boundaries of clusters of pure cooperators/defectors which can become entrenched for a very long time (see also the results section).

- Strategy spread (with the probability $P(j \rightarrow i)$ defined above) occurs in the following way. With probability p_p the reference agent will impose his desired social strategy and will also transfer his replication strategy ($s_i = e_j$ and $e_i = e_j$). Otherwise, with probability p_s only the social strategy is imposed ($s_i = e_j$) and in the remainder of cases, i.e. with probability $p_a = \frac{1-p_p-p_s}{1-p_s}$ only the replication strategy is passed on from j to i ($e_i = e_j$). The timescales for joint or disjoint spread of the traits (parametrized via p_p) and distinction of timescales for the spread of social and replication strategies (parametrized via p_s) prove crucial parameters to understand the dynamics of social evolution in this context.
- The payoff evaluation and strategy updating steps are then repeated for a sufficiently large number of steps until a quasistationary state has been reached. Then, average concentrations of all strategy concentrations are sampled from another TN iterations (note that in this paper time is always measured in units of full lattice sweeps).

In the following we employ computer simulations of systems composed of in between 10^4 and 1.6×10^5 agents to construct phase diagrams of parameter regions in which the evolutionary dynamics can allow cooperation to survive.

Results

Well mixed populations

Before discussing spatial simulations it is worthwhile analysing the case without network reciprocity, i.e. a well-mixed population in which individuals meet at random and




Figure 1: Dependence of the concentrations of defect and cooperate strategies on dilemma toughness for a well-mixed population of size $N = 40000$ and noise level $\kappa = 0.01$. For $p_p = 1$ cooperation can always survive, but for $p_p < 1$ defection wins out for $r > 0$ (and since they all overlap circles at $n = 0$ are omitted for $r > 0.01$).

strategies spread according to Eq. (3) on the basis of payoff gathered from interactions with the entire population. For simplicity, we will not distinguish timescales given by p_s and p_a and assume $p_s/p_a = 1$ in the following. It is then straightforward to describe the evolutionary dynamics of strategy concentrations n_i by a set of rate equations in the form:

$$\dot{n}_i = \sum_{k,l} n_k n_l a_{kl}^{(i)}, \quad (4)$$

where the indices label the four possible strategies 00, 10, 01, and 11 and the matrices $a^{(i)}$ contain information about conversions between strategies according to the rules set out in the previous section. It is worth noting that $n_{00} + n_{10} + n_{01} + n_{11} = 1$, i.e. there are only three relevant degrees of freedom.

For the transition matrices one finds:

$$a^{(00)} = \begin{pmatrix} 0 & 0 & -1/2 & -p \\ \beta P & \beta/2 & 0 & 0 \\ \beta/2 & 0 & \beta p & 0 \\ \alpha P & \alpha/2 & 0 & 0 \end{pmatrix}, \quad (5)$$

where we introduce the shortcuts $\beta = 1 - p_p/2$, $\alpha = 1 - p_p$, $P = 1/(1 + \exp(-\Delta\pi/\kappa))$, and $p = 1/(1 + \exp(\Delta\pi/\kappa))$ to simplify notation with payoff difference between defectors and cooperators

$$\Delta\pi = n_C(1+r) - (-(1-n_C)r + (1-n_C)) \quad (6)$$

$$= r, \quad (7)$$

where $n_C = n_{11} + n_{10}$ is the concentrations of agents with social strategy cooperate.

For instance, if an agent with strategy 00 meets an agent with strategy 01, the agent following 00 will adapt its strategy with likelihood $1/2$ (since both strategies achieved the same payoff). If strategy 00 learns from 01, there are three cases that need to be distinguished. (i) with probability $1 - p_p$ 00 learns the social strategy that 01 wishes to impose (i.e. 1) and 01's replication strategy (i.e. 1) and hence converts to strategy 11. (ii) with probability $p_p/2$ 00 only learns the social strategy 01 wishes to impose, i.e. 00 converts to 10 and (iii) 00 may only learn 01's replication strategy, i.e. 00 converts to 01. In all three cases 00 converts to a strategy different from 00, hence the entry $a_{13}^{(00)} = -1/2$. Similarly, if 10 encounters 00 the chance that 10 will learn from 00 is given by P . Either learning only the social strategy 00 wishes to impose (probability $p_p/2$) or learning both the social strategy 00 wishes to impose and 00's replication strategy (probability $1 - p_p$) will convert 10 to 00. Hence the entry $a_{21}^{(00)} = \beta/2$.

Analogous equations for the remaining three matrices $a^{(10)}$, $a^{(01)}$, and $a^{(11)}$ can be derived, i.e.

$$a^{(10)} = \begin{pmatrix} 0 & 0 & p_p/4 & pp_p/2 \\ -\beta P & -\beta/2 & -\beta P & -\beta/2 \\ 0 & 0 & 0 & 0 \\ Pp_p/2 & p_p/4 & 0 & 0 \end{pmatrix}, \quad (8)$$

$$a^{(01)} = \begin{pmatrix} 0 & 0 & p_p/4 & pp_p/2 \\ 0 & 0 & 0 & 0 \\ -\beta/2 & -\beta p & -\beta/2 & -\beta p \\ Pp_p/2 & p_p/4 & 0 & 0 \end{pmatrix}, \quad (9)$$

$$a^{(11)} = \begin{pmatrix} 0 & 0 & \alpha/2 & p\alpha \\ 0 & 0 & p\beta & \beta/2 \\ 0 & 0 & \beta/2 & \beta p \\ -P & -1/2 & 0 & 0 \end{pmatrix}. \quad (10)$$

The systems of equations (4) is a system of three non-linear equations. Even though an analytical analysis of stationary states might be possible, numerical integration of (4) provides enough insight for the present purposes. Figure 1 gives the dependence of stationary strategy concentrations obtained by numerical integration of (4) on the dilemma strength for two scenarios of strategy pass for $\kappa = 0.01$ (note that noise levels should be measured per interaction, i.e. a very small value in the well-mixed case with all-to-all interactions corresponds to larger noise values on sparse grids).

The first, illustrated by square symbols in Fig. 1 corresponds to completely asynchronous strategy pass, i.e. $p_p = 1$. In this case for all $r > 0$ the population is split into roughly two thirds defectors (equal halves of which carry both replication strategies) and one third cooperators (with again equal halves carrying both replication strategies). In

contrast, for any $p_p \neq 1$ (round symbols) the social strategy cooperate is found to die out, i.e. $n_{10} + n_{11} = 0$, and the two social defect strategies share the population in equal proportions.

It is easy to understand why this is the case. Strategy $s = 0$ and $e = 1$ can earn the same payoff as pure defectors with $s = 0$ and $e = 0$. However, in a well-mixed population it cannot profit from generating offspring who cooperate, because cooperation can be exploited by the entire population of defectors. Hence, agents with $s = 0$, $e = 1$ can generate the same number of 'offspring' as $s = 0$, $e = 0$; but their descendants die out without conferring an advantage on the parents. The situation is different if the spread of social strategy and replication strategy are completely separated: In this case the population of social cooperators is always reinforced by an inflow from the pool of social defectors with replication strategy defect (who earn equal payoff as pure defectors) and can also not be suppressed through interactions with pure defectors, because there is always a one-half chance that the cooperative trait survives due to separate strategy pass.

Spatially distributed populations

As we have seen in the previous section on well-mixed populations, replication strategies that differ from social strategies can only support cooperation if the spread of social and replication strategies is completely separated. The reason for this is that social niche construction cannot operate in well-mixed populations: offspring that plays the social strategy cooperate can be exploited by the entire population and does not bestow any specific benefit to the parent who gave birth to it. Rather, the effect for strategies with replication strategy $e = 1$ is negative: Their offspring will replicate less well than the parent because it can be exploited by the entire population of defectors. One would anticipate that this situation can be different in viscous populations. In the latter case, parents can accrue specific individual fitness benefits by surrounding themselves by cooperators. It appears reasonable to surmise that the consequential increase in reproductive fitness for parents might compensate for the loss in fitness of offspring, thus enabling cooperative strategy traits to survive. We will explore this scenario for spatial games in some detail below.

Figure 2 illustrates simulation results for the evolution of the four strategies in two typical settings in which replication of the two components of a strategy, social strategy s and replication strategy e , are to some extent disjoint ($p_p = 0.6$). The figures also give the frequency f_c of mutually cooperative interactions. In the first setting with lower dilemma toughness (top panel), cooperation can grow to dominance. In the second with somewhat larger dilemma toughness (bottom panel), an equilibrium state in which all four strategies coexist is reached. Spatial arrangements of the strategies that correspond to such a mixed state are illustrated in Fig.




Figure 2: (Average) evolution of social and reproduction strategies for a prisoner's dilemma and average fraction of mutually cooperative interactions f_c . (a) With $r = 0.18$ and $p_p = 0.6$ when cooperation grows to dominance and (b) with $r = 0.22$ when an equilibrium between the strategies is reached (on a 200×200 torus with $\kappa = 1$).

3.

These first experiments which we show in Fig. 2 illustrate two important points: (i) As hypothesised above, when including opportunities for social niche construction via replication strategies, cooperation can survive in spatial arrangements, even if strategy pass is not completely disjoint. (ii) Disjoint transmission of social and replication strategies can allow for the dominance of cooperation in regimes of dilemma games far beyond regimes normally supported by network reciprocity (i.e. for a typical spatial game with von Neumann neighbourhoods with $\kappa = 0.1$ the extinction threshold for cooperation is around $r_c = 0.021$ (Hauert and Szabo, 2004) and even somewhat smaller with $r_c \approx 0.017$ for $\kappa = 1$).

The typical spatial arrangements in Fig. 3 also provide an intuitive understanding why replication strategies can support cooperation in spatial settings. The figure shows the




Figure 3: Example configuration of an equilibrium arrangement of the four strategies (for $r = 0.22, p_p = 0.6, \kappa = 1$). Colors are (s, e) : red (D,D), light red (D,C), green (C,C) and blue (C,D). (70*70) C,D and D,C only occur at the boundaries of larger C,C and D,D clusters.

presence of large homogeneous clusters of pure defectors ($s = 0, e = 0$, dark red) and pure cooperators ($s = 1, e = 1$, green). Strategies with $s \neq e$ only occur at the boundaries of these clusters. A cursory glance at Fig. 3 which is confirmed by the results shown in the bottom panel of Fig. 2 also suggests that the strategy $s = 0, e = 1$ (blue) is far more prominent than strategy $s = 1, e = 0$. The reason is that a social defect strategy can earn larger payoffs than a social cooperate strategy.

Let us now consider the effect of $s = 1, e = 0$ and $s = 0, e = 1$ on the clusters of pure cooperators and defectors. When replicating in the direction of pure cooperators $s = 1, e = 0$ either reproduces itself or (assuming disjoint strategy pass) it produces a defector with $s = 0, e = 1$. However, since $s = 1, e = 0$ earns the same payoff as pure cooperation $s = 1, e = 1$ can only invade clusters of cooperators via neutral drift. By the same token it only rarely gets a chance to replicate when competing against defection, and if so, it cannot reproduce itself (since any pure defector n would either only be influenced in its social strategy via the replication strategy, i.e. $s_n = e = 0$ or would additionally imitate the replication strategy $e_n = e = 0$ which correspond to its own strategy anyway). Hence, $s = 1, e = 0$ impedes the spread of pure cooperation into clusters of pure defectors, but also, by transitioning into $s = 0, e = 1$, delays invasions of defection into clusters of pure cooperators.

What about the spread of the strategy $s = 0, e = 1$? The propagation of $s = 0, e = 1$ is more relevant at the boundary of clusters, since, following social defect, it will typi-




Figure 4: (a) Dependence of the frequency of mutually cooperative interactions f_c on the dilemma strength for various values of p_p on a 200×200 torus. It becomes apparent that cooperation finds more and more support, the more frequent uncorrelated replication events become. (b) Dependence of the concentrations of the various strategies on the dilemma strength for $p_p = 0.4$. There are two phases dominated by pure cooperators or pure defectors (small and large r and an in-between phase in which all strategies can coexist). (c) In contrast, for $p_p = 0$ only pure cooperators and pure defectors survive and the phase diagram from the standard PD without replication strategies is reproduced.




Figure 5: Dependence of critical thresholds for the extinction of strategies on the probability for disjoint trait propagation p_p .

cally harvest a larger payoff than $s = 1, e = 0$. On the one hand, when interacting with pure cooperation, it will always surround itself by pure cooperation. On the other hand, when interacting with a pure defector, it will either generate a $s = 1, e = 0$ defector or cause a transition of the neighbour to pure cooperation. Hence, even though $s = 0, e = 1$ exploits cooperators, it also shields clusters of cooperators from the invasion of defection and promotes the spread of the pure cooperate strategy.

When considering the role of all four strategies at the boundaries of compact clusters of pure cooperators and pure defectors it is also important to recognize that the strategy $s = 0, e = 1$ will typically generate the largest payoff (because being on average surrounded by more cooperators than pure defect) and thus replicate most often. Even though

being thus most successful in terms of replication, it can only recreate itself indirectly – its offspring will never follow the same strategy.

The mechanism which supports cooperation in the simulations shown above principally works as follows. Offspring-exploiting defectors $s = 0, e = 1$ are the most successful strategy, but cannot recreate directly, and, as a result, serve as support for cooperation. It is evident that in case of joint strategy propagation a 'checkerboard pattern' of $s = 0, e = 1$ interspersed with $s = 1, e = 1$ would be evolutionarily stable. However, since $s = 0, e = 1$ cannot recreate itself and is only generated at boundaries of defectors and cooperators when disjoint strategy pass is allowed, without the presence of random strategy invasions or mutations such a pattern cannot evolve from random initial conditions (cf. Fig. 4 right panel). Moreover, this checkerboard pattern is not stable in the face of even small degrees of disjoint strategy spread (measured by p_p). If p_p is sufficiently large, offspring-exploiting defectors support pure cooperation in two ways: (i) by shielding clusters of pure cooperators from the invasion of defection, and by (ii) serving as a source of pure cooperators, as a consequence of their own success in replication.

Figure 5 extends our earlier simulation experiments by giving the dependence of the frequency of cooperative interactions and strategy concentrations on the dilemma toughness. A clear dependence of the support for cooperation on the frequency of joint strategy pass p_p is evident and is further supported by the full phase diagram that illustrates the dependence of coexistence regimes and regimes in which single strategies dominate on p_p . As already indicated by the dependencies in Fig. 4 the coexistence regimes are typically rather small and regimes in which either pure cooperation or pure defection take over the entire population dominate the diagram of Fig. 5. Coexistence is only found in a small




Figure 6: Dependence of critical thresholds for the extinction of strategies on timescales for the spread of the social strategy (p_s) and for the replication strategy (p_a) for fixed $p_p = 0.6$. For $r > r_0$ only pure defection survives, for $r < r_1$ only pure cooperation survives and for $r_0 > r > r_1$ all four strategies can coexist. The faster social strategies spread relative to replication strategies, the more support for cooperation. Also the coexistence region becomes larger the slower the spread of the social strategy.

borderline region between the regimes of pure strategy dominance.

It is also of interest to investigate the dependence of the support for cooperation on the relative timescales for strategy propagation. To explore this question, we set up experiments with a fixed frequency of disjoint strategy pass and vary the relative frequencies with which reference agents only impose their replication strategy as the desired social strategy of neighbours (i.e. with probability p_s) and the frequency with which neighbours only learn the replication strategy of a reference (i.e. with probability p_a). A typical phase diagram that summarizes our simulation experiments is given in Fig. 6.

Two observations stand out. First, the support for cooperation grows the more prominent the imposition of social strategies relative to the learning of replication strategies. This observation is consistent with our argument about the role of offspring-exploiting defectors. The more frequent the spread of social strategies, the more often they will generate pure cooperators. Second, the regime in which all four strategies can coexist becomes the larger the more frequently agents pick up the replication strategies of their references. This second finding is also intuitively clear from the same argument. The more often replication strategies are learnt, the more often $s = 0, e = 1$ transitions into $s = 1, e = 0$, thus boosting the concentrations of other strategies.

A last point worth investigating is the role of a cost for strategy imposition. To investigate this point we presume




Figure 7: Dependence of critical thresholds for the extinction of strategies on costs for the propagation of unequal strategy traits $s \neq e$ for $p_p = 0.5$.

that the behaviour in the standard game, i.e. imposing agents' own social strategy on neighbours, is free. In contrast, imposing a strategy different from an agent's social strategy needs "convincing", i.e. it comes at some cost c , cf. Eq. (2). The experiments carried out in this way allow us to test the stability of the standard framework and answer the question "Would differences between social and replication strategies evolve if teaching is costly?". To explore this question we fix the frequency p_p of disjoint trait transmission and assume that both traits spread at equal rates (i.e. $p_s = p_a$). Figure 7 summarizes these experiments by presenting a phase diagram for the dependence of extinction thresholds on cost assumptions. Clearly, imposing a cost for producing offspring with social strategies different from an agent's social strategy reduces support for cooperation. Such behaviour would naturally be expected, since imposing costs penalizes the "mixed" strategies $s = 0, e = 1$ and $s = 1, e = 0$ and our previous argument relied on the presence of the first of those to support cooperation via the exploitation of offspring. Nevertheless, even imposing costs that are very substantial compared to game payoffs, cooperation can exist in regimes far beyond the support it would find from the network reciprocity of the spatial grid (with an extinction threshold of $r_c \approx 0.017$, see Fig. 4).

From the data presented in Fig. 7 it is also noteworthy that costs reduce support for the coexistence regime. Costs penalize strategies with $s \neq e$ and hence mixed phases in which all four strategies can co-exist are increasingly suppressed the larger the costs imposed.

Conclusions

In this paper we have introduced a simple way to explore social niche construction (Powers et al., 2011) on spatial networks. In our framework, on top of a social strategy every

agent is endowed with a second trait, a replication strategy, which allows the agent to determine the social strategy of its offspring. We then explored the co-evolution of social and replication strategies, subject to various assumptions about the timescales of spread of both strategy components.

Analyzing the dynamics of the co-evolution in the prisoner's dilemma, we have established that cooperation can only be supported in well-mixed populations if social and replication strategies are never both passed on from parent to offspring. In a social context this corresponds to the rather unrealistic assumption that the timescales of learning the respective traits are completely separated. In a biological context, this assumption translates into assumptions about the traits being located on uncoupled separate genes. As demonstrated by our exploration of the spatial prisoner's dilemma, the presence of a structured population can mitigate this strict condition. We have shown that in spatial settings cooperation can find very strong support, even if the simultaneous passing on of social and replication strategies is rather unlikely. The main driver of the support for cooperation is the prevalence of offspring-exploiting defectors which can generate the largest payoffs in the game. Offspring-exploiting defectors are found to be in a similar role as payoff-distinguished agents in (Perc and Szolnoki, 2008; Brede, 2011a): by virtue of their enhanced ability to pass on strategies they assume a "leadership" role (Zimmermann and Eguíluz, 2005). Different from previous models like (Zimmermann and Eguíluz, 2005; Perc and Szolnoki, 2008; Brede, 2011a), however, such agents with $s \neq e$ never replicate identically and thus offspring-exploiting cooperators facilitate the spread of cooperation by surrounding themselves with cooperators.

We have also presented a number of further experiments that corroborate the robustness of the above finding. Support for cooperation is robust to changes of the timescales of strategy spread over several orders of magnitude and also the inclusion of substantial costs for imposing social strategies different from an agent's own social strategy do not alter outcomes in a qualitative way.

One may wonder if the framework in which we introduced replication strategies in this paper is too restrictive. In other words: Would our main findings be robust if replication strategies were context dependent, i.e. influenced by the social strategy of agents which replicate such that an agent in the role of a social cooperator may wish to impose a different strategy on its neighbours than when being in the role of a social defector? We reserve a more comprehensive analysis of this more general setting for future work.

Acknowledgements

The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

References

- Brede, M. (2011a). The evolution of cooperation on correlated payoff landscapes. *Artificial Life*, 17:365–373.
- Brede, M. (2011b). Playing against the fittest: A simple strategy that promotes the emergence of cooperation. *EPL*, 94:30003.
- Brede, M. (2013a). Costly advertising and the evolution of cooperation. *PLoS ONE*, available online: <http://dx.plos.org/10.1371/journal.pone.0067056>.
- Brede, M. (2013b). Short versus long term benefits and the evolution of cooperation in the prisoner's dilemma game. *PLoS ONE*, 8(2):e56016.
- Hauert, C. and Szabo, G. (2004). Game theory and physics. *Am. J. Phys.*, 73:405–414.
- Huberman, B. A. and Glance, N. S. (1993). Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci. USA*, 90:7716–7718.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314:1560–1563.
- Nowak, M. A. and M., M. R. (1992). Evolutionary games and spatial chaos. *Nature*, 359:826–829.
- Perc, M. and Szolnoki, A. (2008). Social diversity and promotion of cooperation in the spatial prisoner's dilemma game. *Phys. Rev. E*, 77:011904.
- Perc, M. and Szolnoki, A. (2010). Coevolutionary games - a mini review. *BioSystems*, 99:109–125.
- Perc, M. and Wang, Z. (2010). Heterogenous aspirations promote cooperation in the prisoner's dilemma game. *PLoS ONE*, 5:e15117.
- Powers, S. T., Penn, A. S., and Watson, R. A. (2011). The concurrent evolution of cooperation and the population structures that support it. *Evolution*, 65:1527–1543.
- Santos, F. C., Rodrigues, J., and Pacheco, J. (2006). Graph topology plays a determinant role in the evolution of cooperation. *Proc. R. Soc. B*, 273:51–55.
- Szabó, G. and Fath, G. (2007). Evolutionary games on graphs. *Phys. Rep.*, 446:97–216.
- Szabó, G. and Toke, C. (1998). Evolutionary prisoner's dilemma game on a square lattice. *Phys. Rev. E*, 58:69–73.
- Szolnoki, A. and Perc, M. (2008). Coevolution of teaching activity promotes cooperation. *New J. Phys.*, 10:043036.
- Szolnoki, A. and Szabó, G. (2007). Cooperation enhanced by inhomogeneous activity of teaching for evolutionary prisoner's dilemma game. *EPL*, 77:30004.
- Szolnoki, A., Vukov, J., and Szabo, G. (2009). Selection of noise level in strategy adoption for spatial social dilemmas. *Phys. Rev. E*, 80:056112.
- Tanimoto, J. and Brede, M. and Yamauchi, A. (2012). Network reciprocity by coexisting learning and teaching strategies. *Phys. Rev. E*, 85:032101.
- Zimmermann, M. and Eguíluz, V. M. (2005). Cooperation, social networks, and the emergence of leadership in a prisoner's dilemma game with adaptive local interactions. *Phys. Rev. E*, 72:056118.