

CrossMark
click for updates

Research

Cite this article: McAvoy A, Hauert C. 2015

Structural symmetry in evolutionary games.

J. R. Soc. Interface **12**: 20150420.<http://dx.doi.org/10.1098/rsif.2015.0420>

Received: 11 May 2015

Accepted: 9 September 2015

Subject Areas:

computational biology, biomathematics

Keywords:

game theory, finite population, evolution

Author for correspondence:

Alex McAvoy

e-mail: alexmcavoy@gmail.com

Structural symmetry in evolutionary games

Alex McAvoy and Christoph Hauert

Department of Mathematics, University of British Columbia, Vancouver, Canada V6T 1Z2

In evolutionary game theory, an important measure of a mutant trait (strategy) is its ability to invade and take over an otherwise-monomorphic population. Typically, one quantifies the success of a mutant strategy via the probability that a randomly occurring mutant will fixate in the population. However, in a structured population, this fixation probability may depend on where the mutant arises. Moreover, the fixation probability is just one quantity by which one can measure the success of a mutant; fixation *time*, for instance, is another. We define a notion of homogeneity for evolutionary games that captures what it means for two single-mutant states, i.e. two configurations of a single mutant in an otherwise-monomorphic population, to be ‘evolutionarily equivalent’ in the sense that *all* measures of evolutionary success are the same for both configurations. Using asymmetric games, we argue that the term ‘homogeneous’ should apply to the evolutionary process as a whole rather than to just the population structure. For evolutionary matrix games in graph-structured populations, we give precise conditions under which the resulting process is homogeneous. Finally, we show that asymmetric matrix games can be reduced to symmetric games if the population structure possesses a sufficient degree of symmetry.

1. Introduction

One of the most basic models of evolution in finite populations is the Moran process [1]. In the Moran process, a population consisting of two types, a mutant type and a wild-type, is continually updated via a birth–death process until only one type remains. The mutant and wild-types are distinguished by only their reproductive fitness, which is assumed to be an intrinsic property of a player. A mutant type has fitness $r > 0$ relative to the wild-type (whose fitness relative to itself is 1), and in each step of the process an individual is selected for reproduction with probability proportional to fitness. Reproduction is clonal, and the offspring of a reproducing individual replaces another member of the population who is chosen for death uniformly at random. Eventually, this population will end up in one of the monomorphic absorbing states: all mutant type or all wild-type. In this context, a fundamental metric of the success of the mutant type is its ability to invade and replace a population of wild-type individuals [2].

In a population of size N , the probability that a single mutant in a wild-type population will fixate in the Moran process is

$$\rho = \frac{1 - r^{-1}}{1 - r^{-N}}. \quad (1.1)$$

In this version of a birth–death process, the members of the population are distinguished by only their types; in particular, there is no notion of spatial arrangement, i.e. the population is *well mixed*. Lieberman *et al.* [3] extend the classical Moran process to graph-structured populations, which are populations with links between the players that indicate who is a neighbour of whom. In this structured version of the Moran process, reproduction happens with probability proportional to fitness, but the offspring of a reproducing individual can replace only a *neighbour* of the parent. Since individuals are now distinguished by both their types (mutant or wild) *and* locations within the population, a natural question is whether or not the fixation probability of a single-mutant

type depends on where this mutant appears in the population. Lieberman *et al.* [3] show that this fixation probability is independent of the location of the mutant if everyone has the same number of neighbours, i.e. the graph is *regular* [4]. In fact, remarkably, the fixation probability of a single mutant on a regular graph is the same as that of equation (1.1)—an observation first made in a special case by Maruyama [5]. This result, known as the Isothermal Theorem, is independent of the number of neighbours the players have (i.e. the *degree* of the graph).

The Moran process is *frequency-independent* in the sense that the fitness of an individual is determined by type and is not influenced by the rest of the population. However, the Moran model can be easily extended to account for *frequency-dependent* fitness. A standard way in which to model frequency-dependence is through evolutionary games [6–8]. In the classical set-up, each player in the population has one of two strategies, *A* and *B*, and receives an aggregate pay-off from interacting with the rest of the population. This aggregate pay-off is usually calculated from a sequence of pairwise interactions whose pay-offs are described by a pay-off matrix of the form

$$\begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \left(\begin{array}{cc} a & b \\ c & d \end{array} \right). \end{array} \quad (1.2)$$

Each player's aggregate pay-off is then translated into fitness and the strategies in the population are updated based on these fitness values. Since a player's pay-off depends on the strategies of the other players in the population, so does that player's fitness. Traditionally, this population is assumed to be infinite, in which case the dynamics of the evolutionary game are governed deterministically by the replicator equation of Taylor & Jonker [6]. More recently, evolutionary games have been considered in finite populations [8,9], where the dynamics are no longer deterministic but rather stochastic. In order to restrict who interacts with whom in the population, these populations can also be given structure. Popular types of structured populations are graphs [3,10,11], sets [12] and demes [13,14].

We focus here on evolutionary games in graph-structured populations that proceed in discrete time steps. Such processes define discrete-time Markov chains, either with or without absorbing states (depending on mutation rates). Typically, in evolutionary game theory, one starts with a population of players and repeatedly updates the population based on some update rule such as birth–death [8], death–birth [10,15], imitation [16], pairwise comparison [17,18] or Wright–Fisher [19,20]. These update rules can be split into two classes: cultural and genetic (see [21]). Cultural update rules involve strategy imitation, whereas genetic update rules involve reproduction and inheritance. Without mutations, an update rule may be seen as giving a probability distribution over a number of strategy-acquisition scenarios: a player inherits a new strategy through imitation (cultural rules) or is born with a strategy determined by the parent(s) (genetic rules). Mutation rates disrupt these scenarios by placing a small probability of a player taking on a novel strategy. The way in which strategy-mutation rates are incorporated into an evolutionary process depends on both the class of the update rule and the specifics of the update rule itself. In a general sense, we say that strategy mutations are *homogeneous* if they depend on neither the players themselves nor the locations of the players. This notion of homogeneous strategy

mutations is analogous to that of a symmetric game, which is a game for which the pay-offs depend on the strategies played but are independent of the identities and locations of the players.

The Isothermal Theorem seems to indicate that populations structured by regular graphs possess a significant degree of homogeneity, meaning that different locations within the population appear to be equivalent for the purposes of evolutionary dynamics. However, it is important to note that (i) fixation probability is just one metric of evolutionary success and (ii) the Moran process is only one example of an evolutionary process. For example, in addition to the probability of fixation, one could look at the *absorption time*, which is the average number of steps until one of the monomorphic absorbing states is reached. Moreover, one could consider frequency-dependent processes, possibly with different update rules, in which fitness is no longer an intrinsic property of an individual but is also influenced by the other members of the population. We show that the Isothermal Theorem does not extend to arbitrary frequency-dependent processes such as evolutionary games. Furthermore, we show that this theorem does not apply to fixation *times*; that is, even for the Moran process on a regular graph, the average number of updates until a monomorphic absorbing state is reached can depend on the initial placement of the mutant.

Given that the Isothermal Theorem does not extend to other processes defined on regular graphs, the next natural question is what is the meaning of a spatially homogeneous population in evolutionary game theory? In fact, we argue using asymmetric games [21] that the term 'homogeneous' should apply to an evolutionary process *as a whole* rather than to just the population structure. Even for populations that appear to be spatially homogeneous, such as populations on complete graphs, non-uniform distribution of resources within the population can result in heterogeneity of the overall process. Similarly, for symmetric games, heterogeneity can be introduced into the dynamics of an evolutionary process through strategy mutations. Therefore, a notion of homogeneity of an evolutionary game should take into account *at least* (i) population structure, (ii) pay-offs, and (iii) strategy mutations.

If the strategy-mutation rates are miniscule, then the population spends most of its time in monomorphic states. With small mutation rates, one can define an embedded Markov chain on the monomorphic states and use this chain to study the success of each strategy [22,23]. That is, when a mutation occurs, the population is assumed to return to a monomorphic state before another mutant arises. Thus, the states of interest are the monomorphic states and the states consisting of a single mutant in an otherwise-monomorphic population. We say that an evolutionary game is *homogeneous* if any two states consisting of a single mutant (*A*-player) in a wild-type population (*B*-players) are mathematically equivalent. We make precise what we mean by 'mathematically equivalent' in §2, but, informally, this equivalence means that any two such states are the same up to relabelling. In particular, *all* metrics, such as fixation probability, absorption time, etc., are the same for any two states consisting of a single *A*-mutant in a *B*-population. We show that an evolutionary game in a graph-structured population is homogeneous if the graph is vertex-transitive ('looks the same' from each vertex), the pay-offs are symmetric and the strategy mutations are homogeneous. This result holds for any update rule and selection intensity.

Finally, we explore the effects of population structure on asymmetric evolutionary games. In the weak selection limit,

we show that asymmetric matrix games with homogeneous strategy mutations can be reduced to symmetric games if the population structure is arc-transitive ('looks the same' from each edge in the graph). This result is a finite-population analogue of the main result of McAvoy & Hauert [21], which states that a similar reduction to symmetric games is possible in sufficiently large populations. Thus, we establish that this reduction applies to any population size if the graph possesses a sufficiently high degree of symmetry. Our explorations, both for symmetric and asymmetric games, clearly demonstrate the effects of population structure, pay-offs and strategy mutations on symmetries in evolutionary games.

2. Markov chains and evolutionary equivalence

2.1. General Markov chains

The evolutionary processes we consider here define discrete-time Markov chains on finite state spaces. The notions of symmetry and evolutionary equivalence that we aim to introduce for evolutionary processes can actually be stated quite succinctly at the level of the Markov chain. We first work with general Markov chains, and later we apply these ideas to evolutionary games.

Definition 2.1 (Symmetry of states). Suppose that $X = \{X_n\}_{n \geq 0}$ is a Markov chain on a (finite) state space, \mathcal{S} , with transition matrix \mathbf{T} . An automorphism of X is a bijection $\phi: \mathcal{S} \rightarrow \mathcal{S}$ such that $\mathbf{T}_{\phi(s), \phi(s')} = \mathbf{T}_{s, s'}$ for each $s, s' \in \mathcal{S}$. Two states $s, s' \in \mathcal{S}$ are said to be symmetric if there exists $\phi \in \text{Aut}(X)$ such that $\phi(s) = s'$.

Definition 2.1 says that the states of the chain can be relabelled in such a way that the transition probabilities are preserved. This relabelling may affect the long-run distribution of the chain since it need not fix absorbing states, so we make one further refinement in order to ensure that if two states are symmetric, then they behave in the same way:

Definition 2.2 (Evolutionary equivalence). States s and s' are *evolutionarily equivalent* if there exists an automorphism of the Markov chain, $\phi \in \text{Aut}(X)$, such that

- (i) $\phi(s) = s'$;
- (ii) if μ is a stationary distribution of X , then $\phi(\mu) = \mu$.

For a Markov chain with absorbing states, the notions of symmetry and evolutionary equivalence of states need not coincide (see example B.16). However, if the Markov chain has a unique stationary distribution (as would be the case if it were irreducible), then symmetry implies evolutionary equivalence:

Proposition 2.3. *If X has a unique stationary distribution, then two states are symmetric if and only if they are evolutionarily equivalent.*

We show in appendix B that a Markov chain symmetry preserves the set of stationary distributions (lemma B.7), so if there is a unique stationary distribution, then condition (ii) of definition 2.2 is satisfied automatically by any

symmetry. Proposition 2.3 is then an immediate consequence of this result.

If s and s' are evolutionarily equivalent, then it is clear for absorbing processes that the probability that s fixates in absorbing state \bar{s} is the same as the probability that s' fixates in state \bar{s} (and similarly for fixation times). If the process has a unique stationary distribution, then the symmetry of s and s' implies that this distribution puts the same mass on s and s' . These properties follow at once from the fact that the states s and s' are equivalent up to relabelling.

2.2. Markov chains defined by evolutionary games

Our focus is on evolutionary games on fixed population structures. If S is a finite set of strategies (or 'actions') available to each player, and if the population size is N , then the state space of the Markov chain defined by an evolutionary game in such a population is $\mathcal{S} = S^N$. For evolutionary games without random strategy mutations, the absorbing states of the chain are the monomorphic states, i.e. the strategy profiles consisting of just a single unique strategy. Thus, states s and s' are evolutionarily equivalent if they are symmetric relative to the monomorphic states. On the other hand, evolutionary processes with strategy mutations are typically irreducible (and have unique stationary distributions); in these processes, the notions of symmetry and evolutionary equivalence coincide by proposition 2.3.

In order to state the definition of a *homogeneous* evolutionary process, we first need some notation. For $s, s' \in S$, we denote by $s_{(s', i), s}$ the state in S^N whose i th coordinate is s' and whose j th coordinate for $j \neq i$ is s ; that is, all players are using strategy s except for player i , who is using s' .

Definition 2.4 (Homogeneous evolutionary process). An evolutionary process on S^N is *homogeneous* if for each $s, s' \in S$, the states $s_{(s', i), s}$ and $s_{(s', j), s}$ are evolutionarily equivalent for each $i, j = 1, \dots, N$. An evolutionary process is *heterogeneous* if it is not homogeneous.

In other words, an evolutionary process is homogeneous if, at the level of the Markov chain it defines, any two states consisting of a single mutant in an otherwise-monomorphic population appear to be relabellings of one another. As noted in §2.1, *all* quantities with which one could measure evolutionary success are the same for these single-mutant states if the process is homogeneous.

3. Evolutionary games on graphs

We consider evolutionary games in graph-structured populations. Unless indicated otherwise, a 'graph' means a directed, weighted graph on N vertices. A directed graph is one in which the edges have orientations, meaning there may be an edge from i to j but not from j to i . Moreover, the edges carry weights, which we assume are non-negative real numbers. A directed, weighted graph is equivalent to a non-negative $N \times N$ matrix, \mathcal{D} , where there is an edge from i to j if and only if $\mathcal{D}_{ij} \neq 0$. If there is such an edge, then the weight of this edge is simply \mathcal{D}_{ij} . Since there is a one-to-one correspondence between directed, weighted graphs on N vertices and $N \times N$ real matrices, we refer to graphs and matrices using the same notation, describing \mathcal{D}




Figure 1. Three different levels of symmetry for connected graphs. Regular graphs have the property that the degrees of the vertices are all the same. Vertex-transitive graphs look the same from each vertex and are necessarily regular. Symmetric (arc-transitive) graphs look the same from any two (directed) edges. Each of these containments is strict; there exist graphs that are regular but not vertex-transitive (figure 2) and vertex-transitive but not symmetric (figure 7a). (Online version in colour.)

as a graph but using the matrix notation \mathcal{D}_{ij} to indicate the weight of the edge from vertex i to vertex j .

Every graph considered here is assumed to be connected (strongly), which means that for any two vertices, i and j , there is a (directed) path from i to j . This assumption is not that restrictive in evolutionary game theory since one can always partition a graph into its strongly connected components and study the behaviour of an evolutionary process on each of these components. Moreover, for evolutionary processes on graphs that are not strongly connected, it is possible to have both (i) recurrent non-monomorphic states in processes without mutations and (ii) multiple stationary distributions in processes with mutations. Some processes (such as the death–birth process) may not even be defined on graphs that are not strongly connected. Therefore, we focus on strongly connected graphs and make no further mention of the term ‘connected’.

Since our goal is to discuss symmetry in the context of evolutionary processes, we first describe several notions of symmetry for graphs. The three types of graphs we treat here are *regular*, *vertex-transitive* and *symmetric*. Informally speaking, a regular graph is one in which each vertex has the same number of neighbouring vertices (and this number is known as the *degree* of the graph). A vertex-transitive graph is one that looks the same from any two vertices; based on the graph structure alone, a player cannot tell if he or she has been moved from one location to another. A symmetric (or *arc-transitive*) graph is one that looks the same from any two edges. That is, if two players are neighbours and are both moved to another pair of neighbouring vertices, then they cannot tell that they have been moved based on the structure of the graph alone. We recall in detail the formal definitions of these terms in appendix B. The relationships between these three notions of symmetry, as well as some examples, are illustrated in figure 1.

We now focus on evolutionary processes in graph-structured populations.

3.1. The Moran process

Consider the Moran process on a graph, \mathcal{D} . Lieberman *et al.* [3] show that if \mathcal{D} is regular, then the fixation probability of a

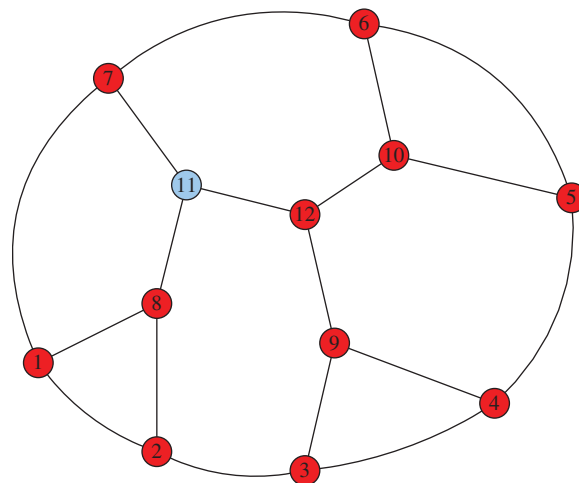


Figure 2. A single mutant (cooperator) at vertex 11 of the Frucht graph. In the Snowdrift Game, the probability that cooperators fixate depends on the initial location of this mutant on the Frucht graph (even if the intensity of selection is weak). (Online version in colour.)

randomly placed mutant is given by equation (1.1), the fixation probability of a single mutant in the classical Moran process. This result (known as the Isothermal Theorem) proves that, in particular, this fixation probability does not depend on the initial location of the mutant. (We refer to this latter statement as the ‘weak’ version of the Isothermal Theorem.) Our definition of homogeneity in the context of evolutionary processes (definition 2.4) is related to this independence of initial location and has nothing to do with fixation probabilities in the classical Moran process. Naturally, the Isothermal Theorem raises the question of whether or not this location independence extends to absorption times (average number of steps until an absorbing state is reached) when \mathcal{D} is regular.

Suppose that \mathcal{D} is the Frucht graph of figure 2. The Frucht graph is an undirected, unweighted, regular (but not vertex-transitive) graph of size 12 and degree 3 [24]. The fixation probabilities and absorption times of a single mutant in a wild-type population are given in figure 3 as a function of the initial location of the mutant. The fixation probabilities do not depend on the initial location of the mutant, as predicted by the Isothermal Theorem, but the absorption times *do* depend on where the mutant arises. In fact, the absorption time is distinct for each different initial location of the mutant. The details of these calculations are in appendix C. Therefore, even the weak form of the Isothermal Theorem fails to hold for absorption times. In particular, the Moran process on a regular graph need not define a homogeneous evolutionary process.

This set-up involving two types of players, frequency-independent interactions, and a population structure defined by a single graph, can be generalized considerably:

3.2. Symmetric games

A powerful version of evolutionary graph theory uses *two* graphs to define relationships between the players: an *interaction graph*, \mathcal{E} , and a *dispersal graph*, \mathcal{D} [25–29]. These graphs both have non-negative weights. As an example of how these two graphs are used to define an evolutionary process, we consider a birth–death process based on two-player, symmetric interactions:




Figure 3. (a) Fixation probability and (b) absorption time versus initial vertex of mutant for the Moran process on the Frucht graph. In both figures, the mutant has fitness $r = 2$ relative to the wild-type. As predicted by the Isothermal Theorem, the fixation probability does not depend on the initial location of the mutant. The absorption time (measured in number of updates) is different for each initial placement of the mutant on the Frucht graph. The precise values of these fixation probabilities and absorption times are in appendix C. (Online version in colour.)

Example 3.1. Consider a symmetric matrix game with n strategies, A_1, \dots, A_n , and pay-off matrix

$$\begin{array}{c}
 A_1 \quad A_2 \quad \cdots \quad A_n \\
 \begin{array}{c}
 A_1 \\
 A_2 \\
 \vdots \\
 A_n
 \end{array}
 \begin{pmatrix}
 a_{11} & a_{12} & \cdots & a_{1n} \\
 a_{21} & a_{22} & \cdots & a_{2n} \\
 \vdots & \vdots & \ddots & \vdots \\
 a_{n1} & a_{n2} & \cdots & a_{nn}
 \end{pmatrix}.
 \end{array} \quad (3.1)$$

If $(s_1, \dots, s_N) \in \{1, \dots, n\}^N$, then the total pay-off to player i is

$$u_i(s_1, \dots, s_N) := \sum_{j=1}^N \mathcal{E}_{ij} a_{s_i s_j}. \quad (3.2)$$

If $\beta \geq 0$ is the intensity of selection, then the fitness of player i is

$$f_\beta(u_i(s_1, \dots, s_N)) := \exp\{\beta u_i(s_1, \dots, s_N)\}. \quad (3.3)$$

In each time step, a player (say, player i) is chosen for reproduction with probability proportional to fitness. With probability $\varepsilon \geq 0$, the offspring of this player adopts a novel strategy uniformly at random from $\{A_1, \dots, A_n\}$; with probability $1 - \varepsilon$, the offspring inherits the strategy of the parent. Next, another member of the population is chosen for death, with the probability of player j dying proportional to \mathcal{D}_{ij} . The offspring then fills the vacancy created by the deceased neighbour, and the process repeats. \mathcal{E} is called the ‘interaction’ graph since it governs pay-offs based on encounters, and \mathcal{D} is called the ‘dispersal’ graph since it is involved in strategy propagation.

3.2.1. Heterogeneous evolutionary games

We now explore the ways in which population structure and strategy mutations can introduce heterogeneity into an evolutionary process. Consider the Snowdrift Game with strategies C (cooperate) and D (defect) and pay-off matrix

$$\begin{array}{c}
 C \quad D \\
 \begin{array}{c}
 C \\
 D
 \end{array}
 \begin{pmatrix}
 5 & 3 \\
 7 & 0
 \end{pmatrix}.
 \end{array} \quad (3.4)$$

Suppose that both \mathcal{E} and \mathcal{D} are the (undirected, unweighted) Frucht graph (figure 2). If the selection intensity is $\beta = 1$, then the fixation probability of a single cooperator in a population of defectors in a death–birth process depends on the initial location of the cooperator (figure 4). Since the Frucht graph is regular (but not vertex-transitive), this example demonstrates that the Isothermal Theorem does not extend to frequency-dependent games. In particular, symmetric games on regular graphs can be heterogeneous, and regularity of the graph does not imply that the ‘fixation probability of a randomly placed mutant’ is well defined. This dependence of the fixation probability on the initial location of the mutant is not specific to the Snowdrift Game or the death–birth update rule; one can show that it also holds for the Donation Game in place of the Snowdrift Game or the birth–death rule in place of the death–birth rule, for instance.

With $\beta = 1$, the selection intensity is fairly strong, which raises the question of whether or not these fixation probabilities still differ if selection is weak. In fact, our observation for this value of β is not an anomaly: suppose that s and s' are states (indicating some non-monomorphic initial configuration of strategies), and that s_i and s_j are monomorphic absorbing states (indicating states in which each player uses the same strategy). Let $\rho_{s,i}$ denote the probability that state i is reached after starting in state s , and let \mathbf{t}_s denote the average number of updates required for the process to reach an absorbing state after starting in state s . Each of $\rho_{s,i}$ and \mathbf{t}_s may be viewed as functions of β , and we have the following result:

Proposition 3.2. Each of the equalities

$$\rho_{s,i} = \rho_{s',j} \quad (3.5a)$$

and

$$\mathbf{t}_s = \mathbf{t}_{s'} \quad (3.5b)$$

holds for either (i) every $\beta \geq 0$ or (ii) at most finitely many $\beta \geq 0$. Thus, if one of these equalities fails to hold for even a single value of β , then it fails to hold for all sufficiently small $\beta > 0$.

For a proof of proposition 3.2, see appendix A. This result allows one to conclude that if there are differences in fixation probabilities or times for large values of β (where these

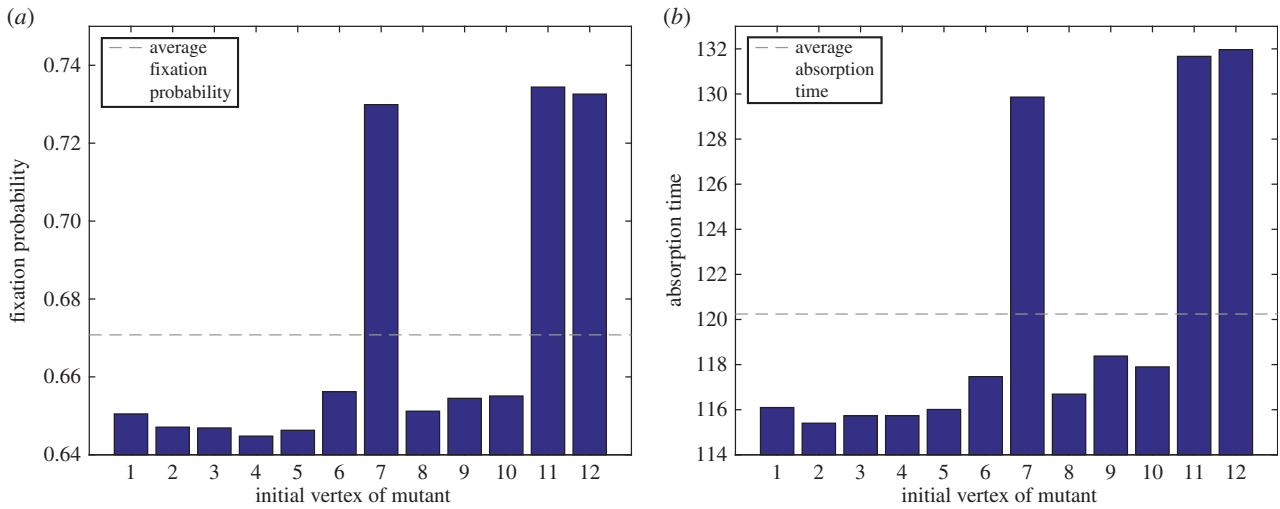


Figure 4. (a) Fixation probability and (b) absorption time versus initial vertex of mutant (cooperator) for a death–birth process on the Frucht graph. In both figures, the game is a Snowdrift Game whose pay-offs are described by pay-off matrix (3.4), and the selection intensity is $\beta = 1$. Unlike the Moran process, this process is frequency-dependent, and it is evident that *both* fixation probabilities and absorption times (measured in number of updates) depend on the initial placement of the mutant. See appendix C for details. Notably, the three vertices (7, 11 and 12) for which both fixation probability and absorption time are highest are the only vertices in the Frucht graph not appearing in a three-cycle. (Online version in colour.)

differences are more apparent), then there are corresponding differences in the limit of weak selection.

Even if a symmetric game is played in a well-mixed population, heterogeneous strategy mutations may result in heterogeneity of the evolutionary process. Consider, for example, the pairwise comparison process [17,18] based on the symmetric Snowdrift Game, (3.4), in a well-mixed population with $N = 3$ players. We model this well-mixed population using a complete, undirected, unweighted graph of size 3 for each of \mathcal{E} and \mathcal{D} (figure 5). For $i \in \{1, 2, 3\}$, let $\varepsilon_i \in [0, 1]$ be the strategy-mutation (‘exploration’) rate for player i . These strategy mutations are incorporated into the process as follows: at each time step, a focal player (player i) is chosen uniformly at random to update his or her strategy. A neighbour (one of the two remaining players) is then chosen randomly as a model player. If β is the selection intensity, π_f is the pay-off of the focal player, and π_m is the pay-off of the model player, then the focal player imitates the strategy of the model player with probability

$$\frac{1 - \varepsilon_i}{1 + e^{-\beta(\pi_m - \pi_f)}}. \quad (3.6)$$

and chooses to retain his or her strategy with probability

$$\frac{1 - \varepsilon_i}{1 + e^{-\beta(\pi_i - \pi_m)}}. \quad (3.7)$$

With probability ε_i , the focal player adopts a new strategy uniformly at random from the set $\{C, D\}$, irrespective of the current strategy. Provided at least one of ε_1 , ε_2 and ε_3 is positive, the Markov chain on $\{C, D\}^3$ defined by this process is irreducible and has a unique stationary distribution, μ . Let $\varepsilon_1 = 0.01$ and $\varepsilon_2 = \varepsilon_3 = 0$. Since the mutation rate depends on the location, i , the strategy mutations are *heterogeneous*. If the selection intensity is $\beta = 1$, then a direct calculation (to four significant figures) gives

$$\mu(C, D, D) = 0.005812 \neq 0.0004897 = \mu(D, C, D), \quad (3.8)$$

where $\mu(C, D, D)$ (resp. $\mu(D, C, D)$) is the mass μ places on the state (C, D, D) (resp. (D, C, D)). Therefore, by proposition 2.3

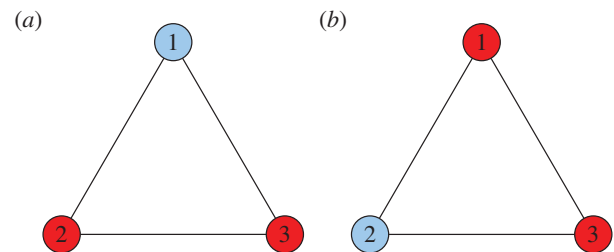


Figure 5. Two states consisting of a single cooperator (mutant) among defectors (the wild-type) in a well-mixed population of size $N = 3$. Despite the spatial symmetry of this population, an evolutionary game on this graph can be heterogeneous as a result of heterogeneous strategy mutations or asymmetric pay-offs. (Online version in colour.)

and definition 2.4, this evolutionary process is *not* homogeneous, despite the fact that the population is well mixed and the game is symmetric. This result is not particularly surprising, but it clearly illustrates the effects of heterogeneous strategy-mutation rates on symmetries of the overall process.

3.2.2. Homogeneous evolutionary games

The behaviour of an evolutionary process sometimes depends heavily on the choice of update rule. As a result, a particular problem in evolutionary game theory is often stated (such as the evolution of cooperation) and subsequently explored separately for a number of different update rules. For example, consider the Donation Game (an instance of the Prisoner’s Dilemma) in which cooperators pay a cost, c , in order to provide the opponent with a benefit, b . Defectors pay no costs and provide no benefits. On a large regular graph of degree k , Ohtsuki *et al.* [10] show that selection favours cooperation in the death–birth process if $b/c > k$, but selection *never* favours cooperation in the birth–death process. Therefore, the approach of exploring a problem in evolutionary game theory separately for several update rules has its merits. On the other hand, one might expect that high degrees of symmetry in the population structure, pay-offs and strategy mutations induce symmetries in an evolutionary game for a variety of update rules.

Before stating our main theorem for symmetric matrix games, we must first understand the basic components that make up an evolutionary game. Evolutionary games generally have two timescales: interactions and updates. In each (discrete) time step, every player in the population has a strategy, and this strategy profile determines the *state* of the population. Neighbours then interact (quickly) and receive pay-offs based on these strategies and the game(s) being played. The total pay-off to a player determines his or her fitness. In the update step of the process, the strategies of the players are updated stochastically based on the fitness profile of the population, the population structure and the strategy mutations. Popular examples of evolutionary update rules are birth–death, death–birth, imitation, pairwise comparison and Wright–Fisher. Since interactions happen much more quickly than updates, there is a separation of timescales.

The most difficult part of an evolutionary game to describe in generality is the update step. If S is the strategy set of the game and N is the population size, then a state of the population is simply an element of S^N , i.e. a specification of a strategy for each member of the population. Implicit in the state space of the population being S^N is an enumeration of the players. That is, if $s \in S^N$ is an N -tuple of strategies, then this profile indicates that player i uses strategy s_i . For our purposes, we need only one property to be satisfied by the update rule of the process, which we state here as an axiom of an evolutionary game:

Axiom. The update rule of an evolutionary game is independent of the enumeration of the players.

Remark 3.3. As an example of what this axiom means, consider a death–birth process in which a player is selected uniformly at random for death and is replaced by the offspring of a neighbour. A neighbour is chosen for reproduction with probability proportional to fitness, and the offspring of this neighbour inherits the strategy of the parent with probability $1 - \varepsilon$ and takes on a novel strategy with probability ε for some $\varepsilon > 0$. If all else is held constant (fitness, mutations, etc.), the fact that a player is referred to as the player at location i is irrelevant: let \mathfrak{S}_N be the symmetric group on N letters. If $\pi \in \mathfrak{S}_N$ is a permutation that relabels the locations of the players by sending i to $\pi^{-1}(i)$, then the strategy of the player at location $\pi^{-1}(i)$ after the relabelling is the same as the strategy of the player at location i before the relabelling. In particular, if $s \in S^N$ is the state of the population before the relabelling, then $\pi(s) \in S^N$ is the state of the population after the relabelling, where $\pi(s)_i = s_{\pi(i)}$. The probability that player $\pi^{-1}(i)$ is selected for death and replaced by the offspring of player $\pi^{-1}(j)$ after the relabelling is the same as the probability that player i is selected for death and replaced by the offspring of player j before the relabelling. Thus, for this death–birth process, the probability of transitioning between states s and s' before the relabelling is the same as the probability of transitioning between states $\pi(s)$ and $\pi(s')$ after the relabelling. In this sense, a relabelling of the players induces an automorphism of the Markov chain defined by the process (in the sense of definition 2.1), and the axiom states that this phenomenon should hold for *any* evolutionary update rule.

In order to state our main result for symmetric games, we note that an evolutionary graph, Γ , in this setting consists of *two* graphs: \mathcal{E} and \mathcal{D} . We say that Γ is regular if both \mathcal{E} and \mathcal{D} are regular. For vertex-transitivity of Γ (resp. symmetry of Γ), we require slightly more than each \mathcal{E} and \mathcal{D} being

vertex-transitive (resp. symmetric); we require them to be *simultaneously* vertex-transitive (resp. symmetric). First of all, we need to define what an *automorphism* of Γ is. For $\pi \in \mathfrak{S}_N$, let $\pi\mathcal{E}$ be the graph defined by $(\pi\mathcal{E})_{ij} := \mathcal{E}_{\pi(i)\pi(j)}$ for each i and j . Using this action, we define an automorphism of an evolutionary graph as follows:

Definition 3.4 (Automorphism of an evolutionary graph). An automorphism of $\Gamma = (\mathcal{E}, \mathcal{D})$ is an action, $\pi \in \mathfrak{S}_N$, such that $\pi\mathcal{E} = \mathcal{E}$ and $\pi\mathcal{D} = \mathcal{D}$. We denote by $\text{Aut}(\Gamma)$ the set of automorphisms of Γ .

We now have the definitions of vertex-transitive and symmetric evolutionary graphs:

Definition 3.5. $\Gamma = (\mathcal{E}, \mathcal{D})$ is vertex-transitive if for each i and j , there exists $\pi \in \text{Aut}(\Gamma)$ such that $\pi(i) = j$.

Definition 3.6. $\Gamma = (\mathcal{E}, \mathcal{D})$ is symmetric if $\mathcal{E} = \mathcal{D}$ and \mathcal{E} is a symmetric graph.

Finally, using the notion of an automorphism of Γ , we have our main result:

Theorem 3.7. Consider an evolutionary matrix game on a graph, $\Gamma = (\mathcal{E}, \mathcal{D})$, with symmetric pay-offs and homogeneous strategy mutations. If $\pi \in \text{Aut}(\Gamma)$, then the states with a single mutant at vertex i and $\pi(i)$, respectively, in an otherwise-monomorphic population, are evolutionarily equivalent. That is, in the notation of definition 2.4, the states $s_{(s',i),s}$ and $s_{(s',\pi(i)),s}$ are evolutionarily equivalent for each $s, s' \in S$.

The proof of theorem 3.7 may be found in appendix B. The proof relies on the observation that the hypotheses of the theorem imply that any two states consisting of a single A -player in a population of B -players can be obtained from one another by relabelling the players. Thus, in light of the argument in Remark (3.3), relabelling the players induces an automorphism on the Markov chain defined by the evolutionary game. Since any relabelling of the players leaves the monomorphic states fixed, there is an evolutionary equivalence between any two such states in the sense of definition 2.2. Note that this theorem makes no restrictions on the selection strength or the update rule.

Corollary 3.8. An evolutionary game on a vertex-transitive graph with symmetric pay-offs and homogeneous strategy mutations is itself homogeneous.

Remark 3.9. By theorem 3.7, two mutants appearing on a graph might define evolutionarily equivalent states even if the graph is not vertex-transitive. For example, the Tietze graph [30], like the Frucht graph, has 12 vertices and is regular but not vertex-transitive. However, unlike the Frucht graph, the Tietze graph has some non-trivial automorphisms. By theorem 3.7, any two vertices in the Tietze graph that are related by an automorphism have the property that the two corresponding single-mutant states are indistinguishable. An example of two evolutionarily equivalent states on this graph is given in figure 6. In appendix C, for the Snowdrift Game with death–birth updating, we give the fixation probabilities and absorption times for all configurations of a single cooperator among defectors, which further illustrates the effects of graph symmetries on an evolutionary process.

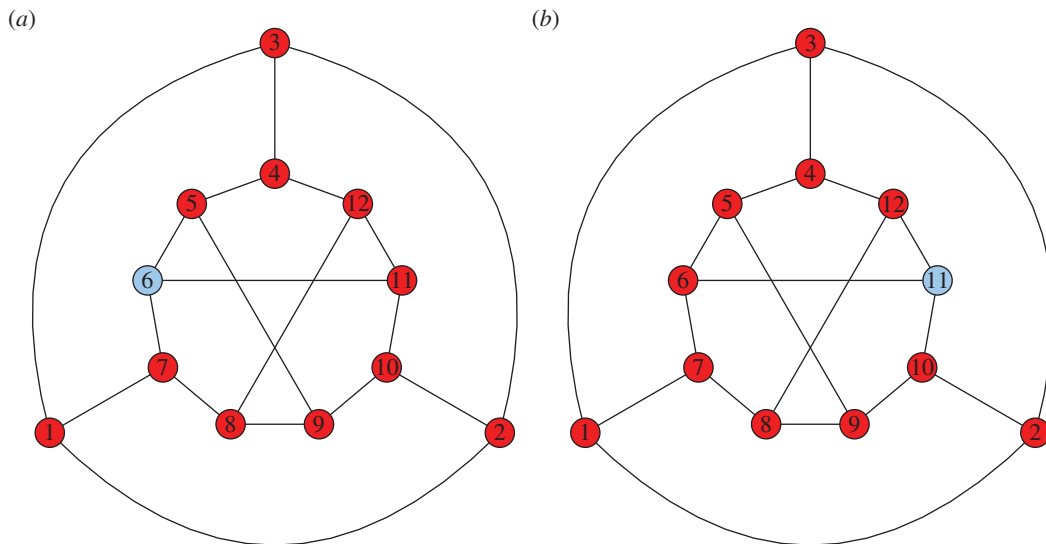


Figure 6. The Tietze graph with two different initial configurations. Like the Frucht graph, the Tietze graph is regular of degree $k = 3$ (but not vertex-transitive) with $N = 12$ vertices. Unlike the Frucht graph, the Tietze graph possesses non-trivial automorphisms. In (a), a cooperator is at vertex 6 and all other players are defectors. In (b), a cooperator is at vertex 11 and, again, the other players are defectors. Despite the fact that the Tietze graph is not vertex-transitive, the single-mutant states defined by (a) and (b) are evolutionarily equivalent. Graphically, this result is clear since one obtains (a) from (b) by flipping the graph (i.e. applying an automorphism), and such a difference between the two states does not affect fixation probabilities, times, etc. However, it is not true that any two single-mutant states are evolutionarily equivalent. For example, in the Snowdrift Game with $\beta = 0.1$ and death–birth updating, the single-mutant state with a cooperator at vertex 1 (resp. vertex 6) has a fixation probability of 0.3777 (resp. 0.4186). Therefore, the two single-mutant states with cooperators at vertices 1 and 6, respectively, are not evolutionarily equivalent, so this process is not homogeneous. (Online version in colour.)

Remark 3.10. For a given population size, N , and network degree, k , there may be many vertex-transitive graphs of size N with degree k . For each such graph, the fixation probability of a randomly occurring mutant is independent of where on the graph it occurs by theorem 3.7. However, this fixation probability depends on more than just N and k ; it also depends on the configuration of the network. For example, figure 7 gives two vertex-transitive graphs of size $N = 6$ and degree $k = 3$. As an illustration, consider the Snowdrift Game (with pay-off matrix (3.4)) on these graphs with birth–death updating. If the selection intensity is $\beta = 0.1$, then the fixation probability of a single cooperator in a population of defectors is 0.1632 in (A) and 0.1722 in (B) (both rounded to four significant figures). These two fixation probabilities differ for all but finitely many $\beta \geq 0$ by proposition (3.2).

Until this point, our focus has been on states consisting of just a single mutant in an otherwise-monomorphic population. One could also inquire as to when any two states consisting of two (or three, four, etc.) mutants are evolutionarily equivalent. It turns out that the answer to this question is simple: in general, the population must be well mixed in order for any two m -mutant states to be evolutionarily equivalent if $m > 1$. The proof that this equivalence holds in well-mixed populations follows from the argument given to establish theorem 3.7 (see appendix B). On the other hand, if the population is not well mixed, then one can find a pair of states with the first state consisting of two mutants on neighbouring vertices and the second state consisting of two mutants on non-neighbouring vertices. In general, the mutant type will have different fixation probabilities in these two states. For example, in the Snowdrift Game on the graph of figure 7b, consider the two states, s and s' , where s consists of cooperators on vertices 1 and 2 only and s' consists of cooperators on vertices 1 and 3 only. If $\beta = 0.1$, then the fixation probability of cooperators under death–birth updating when starting at s (resp. s') is

0.3126 (resp. 0.2607). Therefore, despite the arc-transitivity of this graph, it is not true that any two states consisting of exactly two mutants are evolutionarily equivalent. Only in well-mixed populations are we guaranteed that any two such states are equivalent.

3.3. Asymmetric games

One particular form of pay-off asymmetry appearing in evolutionary game theory is *ecological asymmetry* [21]. Ecological asymmetry can arise as a result of an uneven distribution of resources. For example, in the Donation Game, a cooperator at location i might provide a benefit to his or her opponent based on some resource derived from the environment. Both this resource and the cost of donating it could depend on i , which means that different players have different pay-off matrices. These pay-off matrices depend on both the location of the focal player and the locations of the opponents. Thus, pay-offs for a player at location i against an opponent at location j in an n -strategy ‘bimatrix’ game [21,31,32] are given by the asymmetric pay-off matrix

$$\mathbf{M}^{ij} = \begin{matrix} & A_1 & A_2 & \dots & A_n \\ \begin{matrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{matrix} & \begin{pmatrix} a_{11}^{ij}, a_{11}^{jj} & a_{12}^{ij}, a_{21}^{jj} & \dots & a_{1n}^{ij}, a_{n1}^{jj} \\ a_{21}^{ij}, a_{12}^{jj} & a_{22}^{ij}, a_{22}^{jj} & \dots & a_{2n}^{ij}, a_{n2}^{jj} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}^{ij}, a_{1n}^{jj} & a_{n2}^{ij}, a_{2n}^{jj} & \dots & a_{nn}^{ij}, a_{nn}^{jj} \end{pmatrix} \end{matrix}. \quad (3.9)$$

Similar to equation (3.2), the total pay-off to player i for strategy profile $(s_1, \dots, s_N) \in \{1, \dots, n\}^N$ is

$$u_i(s_1, \dots, s_N) := \sum_{j=1}^N \mathcal{E}_{ij} a_{s_i s_j}^{ij}. \quad (3.10)$$

We saw in §3.2 an example of a heterogeneous evolutionary game in a well-mixed population with symmetric pay-offs. Rather than looking at a symmetric game with heterogeneous

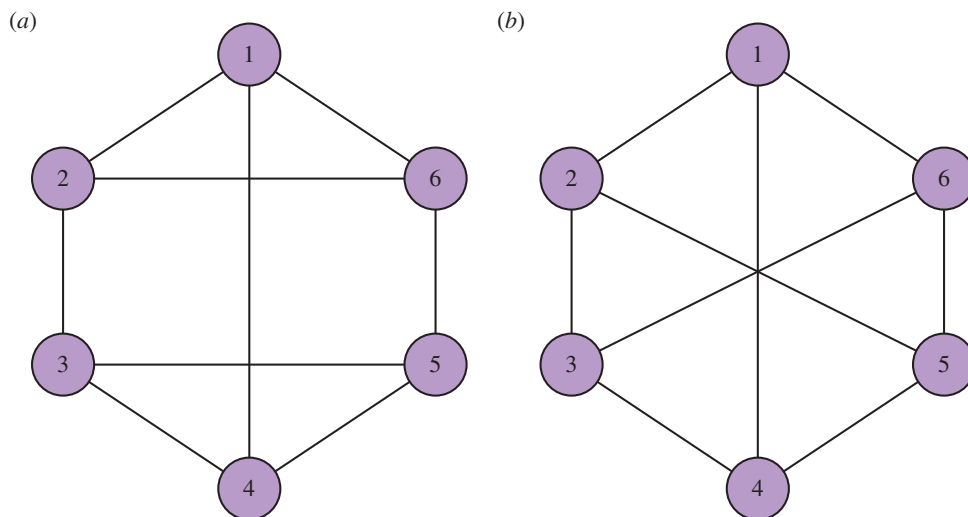


Figure 7. Undirected, unweighted, vertex-transitive graphs of degree $k = 3$ with $N = 6$ vertices; (b) is a symmetric (arc-transitive) graph and (a) is not. (Online version in colour.)

strategy mutations, we now look at an asymmetric game with homogeneous strategy mutations. Consider the ecologically asymmetric Donation Game on the graph of figure 5 (both \mathcal{E} and \mathcal{D}) with a death–birth update rule. In this asymmetric Donation Game, a cooperator at location i donates b_i at a cost of c_i ; defectors donate nothing and incur no costs. If $\beta = 0.1$, $b_1 = b_2 = b_3 = 4$, $c_1 = 1$ and $c_2 = c_3 = 3$, then the fixation probability of a single cooperator at location 1 (figure 5a) is 0.2232, while the fixation probability of a single cooperator at location 2 (figure 5b) is 0.1842 (both rounded to four significant figures). Therefore, even in a well-mixed population with homogeneous strategy mutations (none, in this case), asymmetric pay-offs can prevent an evolutionary game from being homogeneous.

Asymmetric matrix games in large populations reduce to symmetric games if selection is weak [21]. In the limit of weak selection, McAvoy & Hauert [33] establish a selection condition for asymmetric matrix games in finite graph-structured populations that extends the condition (for symmetric games) of Tarnita *et al.* [34]:

Theorem 3.11 (McAvoy & Hauert [33]). *There exists a set of structure coefficients, $\{\sigma_1^{ij}, \sigma_2^{ij}, \sigma_3^{ij}\}_{i,j}$, independent of pay-offs, such that weak selection favours strategy $r \in \{1, \dots, n\}$ if and only if*

$$\sum_{i,j=1}^N \mathcal{E}_{ij}(\sigma_1^{ij}(a_{rr}^{ij} - \bar{a}_{**}^{ij}) + \sigma_2^{ij}(\bar{a}_{r*}^{ij} - \bar{a}_{*r}^{ij}) + \sigma_3^{ij}(\bar{a}_{r*}^{ij} - \bar{a}^{ij})) > 0, \quad (3.11)$$

where $\bar{a}_{**}^{ij} = (1/n) \sum_{s=1}^n a_{ss}^{ij}$, $\bar{a}_{r*}^{ij} = (1/n) \sum_{s=1}^n a_{rs}^{ij}$, $\bar{a}_{*r}^{ij} = (1/n) \sum_{s=1}^n a_{sr}^{ij}$ and $\bar{a}^{ij} = (1/n^2) \sum_{s,t=1}^n a_{st}^{ij}$.

Strictly speaking, theorem 3.11 is established for \mathcal{E} and \mathcal{D} undirected, unweighted and satisfying $\mathcal{E} = \mathcal{D}$. However, the proof of theorem 3.11 extends immediately to the case with \mathcal{E} and \mathcal{D} directed, weighted and possibly distinct, so we make no restrictive assumptions on \mathcal{E} and \mathcal{D} in the statement of this theorem here. In the simpler case $n = 2$, condition 3.11 takes the form

$$\sum_{i,j=1}^N \mathcal{E}_{ij}(\tau_1^{ij}(a_{11}^{ij} - a_{22}^{ij}) + \tau_2^{ij}(a_{12}^{ij} - a_{21}^{ij})) > 0, \quad (3.12)$$

for some collection $\{\tau_1^{ij}, \tau_2^{ij}\}_{i,j}$. For the death–birth process with \mathcal{E} and \mathcal{D} the graph of figure 7a, we calculate *exact* values for all of these structure coefficients (see appendix C). In particular, we find that $\tau_1^{12} = 707905/931552$ and $\tau_1^{14} = 16291/194074$, so vertex-transitivity does not guarantee that the structure coefficients are independent of i and j . For the same process on the graph in figure 7b, we find that $\tau_1^{ij} = \tau_2^{ij} = 2189/27728$ for each i and j , so these coefficients do not depend on i and j . (In general, even for well-mixed populations, τ_1 and τ_2 need not be the same; for the same process studied here but on the graph of figure 5, $\tau_1^{ij} = 33/1616$ and $\tau_2^{ij} = 99/1616$ for each i and j .) This lack of dependence on i and j is due to the fact that the graph of figure 7b is *symmetric*, and it turns out to be a special case of a more general result:

Theorem 3.12. *Suppose that an asymmetric matrix game with homogeneous strategy mutations is played on an evolutionary graph, $\Gamma = (\mathcal{E}, \mathcal{D})$. For each $\pi \in \text{Aut}(\Gamma)$, $k \in \{1, 2, 3\}$ and $i, j \in \{1, \dots, N\}$,*

$$\sigma_k^{ij} = \sigma_k^{\pi(i)\pi(j)}. \quad (3.13)$$

The proof of theorem 3.1 may be found in appendix B. The following corollary is an immediate consequence of theorem 3.12:

Corollary 3.13. *If $\mathcal{E} = \mathcal{D}$ and \mathcal{E} is a symmetric graph (i.e. Γ is a symmetric evolutionary graph), then the structure coefficients are independent of i and j .*

Since symmetric graphs are also regular, we have:

Corollary 3.14. *If $\mathcal{E} = \mathcal{D}$ and \mathcal{E} is a symmetric graph (i.e. Γ is a symmetric evolutionary graph), then strategy r is favoured in the limit of weak selection if and only if*

$$\sigma_1(a_{rr} - \bar{a}_{**}) + \sigma_2(\bar{a}_{r*} - \bar{a}_{*r}) + \sigma_3(\bar{a}_{r*} - \bar{a}) > 0, \quad (3.14)$$

where $\bar{\mathbf{M}} = (a_{st})_{1 \leq s, t \leq n}$ is the spatial average of the matrices \mathbf{M}^{ij} , i.e.

$$\bar{\mathbf{M}} := \frac{1}{kN} \sum_{i,j=1}^N \mathcal{E}_{ij} \mathbf{M}^{ij}, \quad (3.15)$$

where k is the degree of the graph, Γ .

Remark 3.15. Equation (3.14) is just the selection condition of Tarnita *et al.* [34] for symmetric matrix games.

It follows from corollary 3.14 that asymmetric matrix games on arc-transitive (symmetric) graphs can be reduced to symmetric games in the limit of weak selection.

4. Discussion

Evolutionary games in finite populations may be split into two classes: those with absorbing states ('absorbing processes') and those without absorbing states. In absorbing processes, the notion of fixation probability has played a crucial role in quantifying evolutionary outcomes, but fixation probabilities are far from the only measure of evolutionary success. Much of the literature on evolutionary games with absorbing states has neglected other metrics such as the time to absorption or the time to fixation conditioned on fixation occurring ('conditional fixation time'). This bias towards fixation probabilities has resulted in certain evolutionary processes appearing more symmetric than they actually are. We have illustrated this phenomenon using the frequency-independent Moran process on graphs: the Isothermal Theorem guarantees that, on regular graphs, a single mutant cannot distinguish between initial locations in the graph if the only metric under consideration is the probability of fixation. However, certain initial placements of the mutant may result in faster absorption times than others if the graph is regular but not vertex-transitive, and the Frucht graph exemplifies this claim. The same phenomenon also holds for conditional fixation times.

The Frucht graph, which is a regular structure with no non-trivial symmetries (there are no two vertices from which the graph 'looks' the same), also allowed us to show that the Isothermal Theorem of Lieberman *et al.* [3] does not extend to frequency-dependent evolutionary games. That is, on regular graphs, the probability of fixation of a single mutant may depend on the initial location of the mutant if fitness is frequency-dependent. This claim was illustrated via a death–birth process on the Frucht graph, in which the underlying evolutionary game was a Snowdrift Game. For $\beta = 1$ (strong selection), the fixation probability of a cooperator at vertex 11 was nearly 14% larger than the fixation probability of a cooperator at vertex 4. Moreover, we showed that if the fixation probabilities of two initial configurations differ for a single value of β , then they are the same for at most finitely many values of β . In particular, these fixation probabilities differ for almost every selection strength, so our observation for $\beta = 1$ was not an anomaly. Similar phenomena are observed for frequency-dependent birth–death processes on the Frucht graph, for example, and even for frequency-dependent games with the 'equal gains from switching' property, such as the Donation Game.

Theorem 3.7 is an analogue of the Isothermal Theorem that applies to a broader class of games and update rules. The Isothermal Theorem is remarkable since regularity of the population structure implies that the fixation probabilities are not only independent of the initial location of the mutant, they are the *same* as those of the classical Moran process. Our treatment of homogeneous evolutionary processes is focused on when different single-mutant states are equivalent, not when they are equivalent to the corresponding states in the classical

Moran process. Even if the fixation probability of a single mutant does not depend on the mutant's location, other factors (such as birth and death rates) may affect whether or not this fixation probability is the same as the one in a well-mixed population [35,36]. Remark 3.10, which compares the fixation probabilities for the Snowdrift Game on two different vertex-transitive graphs of the same size and degree, shows that the fixation probability of a single mutant—even if independent of the mutant's location—can depend on the configuration of the graph. In light of these results, the symmetry phenomena for the Moran process guaranteed by the Isothermal Theorem do not generalize and should be thought of as properties of the frequency-independent Moran process and not of evolutionary processes in general.

Theorem 3.7, and indeed most of our discussion of homogeneity, focused on symmetries of states consisting of just a *single* mutant. In many cases, mutation rates are sufficiently small that a mutant type, when it appears, will either fixate or go extinct before another mutation occurs [22,23]. Thus, with small mutation rates, one need not consider symmetries of states consisting of more than one mutant. However, if mutation rates are larger, then these multi-mutant states become relevant. Our definition of evolutionary equivalence (definition 2.2) applies to these states, but, as expected, the symmetry conditions on the population structure guaranteeing any two multi-mutant states are equivalent are much stronger. In fact, as we argued in §3.2.2, the population must, in general, be well-mixed even for any pair of states consisting of *two* mutants to be evolutionarily equivalent. Consequently, our focus on single-mutant states allowed us to simultaneously treat biologically relevant configurations (assuming mutation rates are small) and obtain non-trivial conditions guaranteeing homogeneity of an evolutionary process.

The counterexamples presented here could be defined on sufficiently small population structures, and thus all calculations (fixation probabilities, structure coefficients, etc.) are exact. However, these quantities need not always be explicitly calculated in order to prove useful: in our study of asymmetric games, we concluded that an asymmetric game on an arc-transitive (symmetric) graph can be reduced to a symmetric game in the limit of weak selection. (The graph of figure 7a demonstrates that vertex-transitivity alone does not guarantee that an asymmetric game can be reduced to a symmetric game in this way.) This result was obtained by examining the qualitative nature of the structure coefficients in the selection condition 3.11, but it did not require explicit calculations of these coefficients. Therefore, despite the difficulty in actually calculating these coefficients, they can still be used to glean qualitative insight into the dynamics of evolutionary games.

On large random regular graphs, the dynamics of an asymmetric matrix game are equivalent to those of a certain symmetric game obtained as a 'spatial average' of the individual asymmetric games [21]. Corollary 3.14 is highly reminiscent of this type of reduction to a symmetric game. For large populations, this result is obtained by observing that large random regular graphs approximate a Bethe lattice [4] and then using the pair approximation method [37] to describe the dynamics. The pair approximation method is *exact* for a Bethe lattice [10], so, from this perspective, corollary 3.14 is not that surprising since Bethe lattices are arc-transitive. Of course, a Bethe lattice has infinitely many vertices, and corollary 3.14 is a finite-population analogue of this result.

The term ‘homogeneous’ is used in the literature to refer to several different kinds of population structures. This term has been used to describe well-mixed populations [38,39]. For graph-structured populations, ‘homogeneous graph’ sometimes refers to vertex-transitive graphs [26,40]. In algebraic graph theory, however, the term ‘homogeneous graph’ implies a much higher degree of symmetry than does vertex-transitivity [41]. ‘Homogeneous’ has also been used to describe graphs in which each vertex has the same number of neighbours, i.e. regular graphs [42–44]. In between regular and vertex-transitive graphs, ‘homogeneous graph’ has also referred to large, random regular graphs [45]. As we noted, large, random regular graphs approximate Bethe lattices (which are infinite, arc-transitive graphs), but these approximations need not themselves be even vertex-transitive.

In many of the various uses of the term ‘homogeneous’, a common aim is to study the fixation probability of a randomly placed mutant. Our definition of *homogeneous evolutionary game* formally captures what it means for two single-mutant states to be equivalent, and our explorations of the Frucht graph (in conjunction with theorem 3.7) show that vertex-transitivity, and *not* regularity, is what the term ‘homogeneous’ in graph-structured populations should indicate. We also demonstrated the effects of pay-offs and strategy mutations on the behaviour of these single-mutant states and concluded that the term homogeneous should apply to the entire process rather than to just the population structure. The homogeneity (theorem 3.7) and symmetry (theorem 3.12) results given here do not depend on the update rule, in contrast to results such as the symmetry of conditional fixation times in the Moran process of Taylor *et al.* [46] or the Isothermal Theorem of Lieberman *et al.* [3]. We now know that games on regular graphs are *not* homogeneous, and we know precisely under which conditions the ‘fixation probability of a randomly placed mutant’ is well defined. These results provide a firmer foundation for evolutionary game theory in finite populations and a basis for defining the evolutionary success of the strategies of a game.

Authors’ contributions. A.M. and C.H. designed and wrote the paper. A.M. performed the research.

Competing interests. We declare we have no competing interests.

Funding. A.M. and C.H. acknowledge financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC), grant no. RGPIN-2015-05795, and C.H. from the Foundational Questions in Evolutionary Biology Fund (FQEB), grant no. RFP-12-10.

Acknowledgements. The authors thank Wes Maciejewski for helpful conversations.

Appendix A. Fixation and absorption

Using a method inspired by a technique of Press & Dyson [47], we derive explicit expressions (in terms of the transition matrix) for fixation probabilities and absorption times. Subsequently, we prove a simple lemma that says that Markov chain symmetries preserve the set of a chain’s stationary distributions.

A.1. Fixation probabilities

Suppose that $\{X_n\}_{n \geq 0}$ is a discrete-time Markov chain on a finite state space, \mathcal{S} , that has exactly K (≥ 1) absorbing states, s_1, \dots, s_K . Moreover, suppose that the non-absorbing

states are transient [22]. The transition matrix for this chain, \mathbf{T} , may be written as

$$\mathbf{T} = \begin{pmatrix} \mathbf{I}_K & \mathbf{0} \\ \mathbf{S}_1 & \mathbf{S}_2 \end{pmatrix}, \quad (\text{A } 1)$$

where \mathbf{I}_K is the $K \times K$ identity matrix and $\mathbf{0}$ is the matrix of zeros (in this case, its dimension is $K \times (|\mathcal{S}| - K)$, where $|\mathcal{S}|$ is the number of states in \mathcal{S}). This chain will eventually end up in one of the K absorbing states, and we denote by $\rho_{s,i}$ the probability that state s_i is reached when the chain starts off in states $s \in \mathcal{S}$. Let \mathbf{P} be the $|\mathcal{S}| \times K$ matrix of fixation probabilities, i.e. $\mathbf{P}_{s,i} = \rho_{s,i}$ for each s and i . This matrix satisfies

$$\mathbf{TP} = \mathbf{P}, \quad (\text{A } 2)$$

which is just the matrix form of the recurrence relation satisfied by fixation probabilities (obtained from a ‘first-step’ analysis of the Markov chain). Consider the matrix, $\mathcal{M} = \mathcal{M}(\mathbf{T})$, defined by

$$\mathcal{M} := - \begin{pmatrix} \mathbf{I}_K & \mathbf{0} \\ \mathbf{S}_1 & \mathbf{S}_2 - \mathbf{I}_{|\mathcal{S}-K} \end{pmatrix}. \quad (\text{A } 3)$$

Since $\rho_{j,i} = \delta_{ij}$ for $i, j \in \{1, \dots, K\}$, we see that

$$-\mathcal{M}\mathbf{P} = \mathbf{TP} - \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{|\mathcal{S}-K} \end{pmatrix} \mathbf{P} = \begin{pmatrix} \mathbf{I}_K \\ \mathbf{0} \end{pmatrix}. \quad (\text{A } 4)$$

Moreover, the matrix \mathcal{M} must have full rank since the non-absorbing states are transient; that is,

$$\det \mathcal{M} = (-1)^{|\mathcal{S}|} \det (\mathbf{S}_2 - \mathbf{I}_{|\mathcal{S}-K}) \neq 0. \quad (\text{A } 5)$$

Therefore, by Cramer’s rule,

$$\rho_{s,i} = - \frac{\det \mathcal{M}(s, \mathbf{e}_i)}{\det \mathcal{M}}, \quad (\text{A } 6)$$

where the notation $\mathcal{M}(s, \mathbf{e}_i)$ means the matrix obtained by replacing the column corresponding to state s with the i th standard basis vector, \mathbf{e}_i . Thus, equation (A 6) gives explicit formulae for the fixation probabilities.

A.2. Absorption times

Let \mathbf{t} be the $|\mathcal{S}|$ -vector indexed by \mathcal{S} whose entry t_s is the expected time until the process fixates in one of the absorbing states when started in state $s \in \mathcal{S}$. This vector satisfies $t_i = 0$ for $i = 1, \dots, K$ as well as the recurrence relation $\mathbf{T}(\mathbf{t} + \mathbf{1}) = \mathbf{t} + \sum_{i=1}^K \mathbf{e}_i$. Therefore,

$$\mathcal{M}\mathbf{t} = \sum_{j=K+1}^{|\mathcal{S}|} \mathbf{e}_j, \quad (\text{A } 7)$$

so, by Cramer’s rule,

$$t_s = \sum_{j=K+1}^{|\mathcal{S}|} \frac{\det \mathcal{M}(s, \mathbf{e}_j)}{\det \mathcal{M}}. \quad (\text{A } 8)$$

We now turn to Markov chains defined by evolutionary games. Before proving proposition 3.2, we make two assumptions:

- (i) The pay-off-to-fitness mapping is of the form $f_\beta(\pi) = \exp\{\beta\pi\}$, $f_\beta(\pi) = \beta\pi$, $f_\beta(\pi) = 1 + \beta\pi$ or $f_\beta(\pi) = 1 - \beta + \beta\pi$, where β denotes the intensity of selection and π denotes pay-off. (Of course, fitness can be defined in one of the latter three ways only if the pay-offs are such that $f_\beta(\pi) \geq 0$.)

- (ii) The update probabilities are rational functions of the fitness profile of the population.

Remark A.1. Assumptions (i) and (ii) are not at all restrictive in evolutionary game theory. Any process in which selection occurs with probability proportional to fitness will satisfy this rationality condition, and indeed all of the standard evolutionary processes (birth–death, death–birth, imitation, pairwise comparison, Wright–Fisher, etc.) have this property. The four pay-off-to-fitness mappings are also standard.

With assumptions (i) and (ii) in mind, we have:

Proposition (3.2). *Each of the equalities*

$$\rho_{s,i} = \rho_{s',j} \quad (\text{A } 9a)$$

and

$$\mathbf{t}_s = \mathbf{t}_{s'} \quad (\text{A } 9b)$$

holds for either (i) every $\beta \geq 0$ or (ii) at most finitely many $\beta \geq 0$. Thus, if one of these equalities fails to hold for even a single value of β , then it fails to hold for all sufficiently small $\beta > 0$.

Proof. Suppose that $s, s' \in S$ and that s_i and s_j are absorbing states. By equation (A 6),

$$\begin{aligned} \rho_{s,i} = \rho_{s',j} &\Leftrightarrow \frac{\det \mathcal{M}(s, \mathbf{e}_i)}{\det \mathcal{M}} = \frac{\det \mathcal{M}(s', \mathbf{e}_j)}{\det \mathcal{M}} \\ &\Leftrightarrow \det \mathcal{M}(s, \mathbf{e}_i) = \det \mathcal{M}(s', \mathbf{e}_j). \end{aligned} \quad (\text{A } 10)$$

Similarly, by equation (A 8),

$$\begin{aligned} \mathbf{t}_s = \mathbf{t}_{s'} &\Leftrightarrow \sum_{j=K+1}^{|S|} \frac{\det \mathcal{M}(s, \mathbf{e}_j)}{\det \mathcal{M}} = \sum_{j=K+1}^{|S|} \frac{\det \mathcal{M}(s', \mathbf{e}_j)}{\det \mathcal{M}} \\ &\Leftrightarrow \sum_{j=K+1}^{|S|} \det \mathcal{M}(s, \mathbf{e}_j) = \sum_{j=K+1}^{|S|} \det \mathcal{M}(s', \mathbf{e}_j). \end{aligned} \quad (\text{A } 11)$$

Assuming (i) and (ii), equations (A 10) and (A 11) are equivalent to polynomial equations in either β or $\exp\{\beta\}$. Either way, since non-zero polynomial equations have at most finitely many solutions, we see that the equalities $\rho_{s,i} = \rho_{s',j}$ and $\mathbf{t}_s = \mathbf{t}_{s'}$ each hold for either (i) every β or (ii) finitely many values of β . Thus, if $\rho_{s,i} \neq \rho_{s',j}$ (resp. $\mathbf{t}_s \neq \mathbf{t}_{s'}$) for even a single selection intensity, then these fixation probabilities (resp. absorption times) differ for almost every selection intensity. In particular, they differ for all sufficiently small β . ■

Appendix B. Symmetry and evolutionary equivalence

B.1. Symmetries of graphs

Here we recall some standard notions of symmetry for graphs. Although we treat directed, weighted graphs in general, throughout the main text we give several examples of *undirected* and *unweighted* graphs, which are defined as follows:

Definition B.1 (Undirected graph). A graph, \mathcal{D} , is undirected if $\mathcal{D}_{ij} = \mathcal{D}_{ji}$ for each i and j .

Definition B.2 (Unweighted graph). A graph, \mathcal{D} , is unweighted if $\mathcal{D}_{ij} \in \{0, 1\}$ for each i and j .

Since our goal is to discuss symmetry in the context of evolutionary processes, we first describe several notions of symmetry for graphs. In a graph, \mathcal{D} , the *indegree* and *outdegree* of vertex i are $\sum_{j=1}^N \mathcal{D}_{ji}$ and $\sum_{j=1}^N \mathcal{D}_{ij}$, respectively. With these definitions in mind, we recall the definition of a *regular* graph:

Definition B.3 (Regular graph). \mathcal{D} is regular if and only if there exists $k \in \mathbb{R}$ such that

$$\sum_{j=1}^N \mathcal{D}_{ji} = \sum_{j=1}^N \mathcal{D}_{ij} = k, \quad (\text{B } 1)$$

for each i . If \mathcal{D} is regular, then k is called the *degree* of \mathcal{D} .

Let \mathfrak{S}_N denote the symmetric group on N letters; that is, \mathfrak{S}_N is the set of all bijections $\pi: \{1, \dots, N\} \rightarrow \{1, \dots, N\}$. Each $\pi \in \mathfrak{S}_N$ extends to a relabelling action on the set of directed, weighted graphs defined by $(\pi\mathcal{D})_{ij} = \mathcal{D}_{\pi(i)\pi(j)}$. In other words, any relabelling of the set of vertices results in a corresponding relabelling of the graph. The *automorphism group* of \mathcal{D} , written $\text{Aut}(\mathcal{D})$, is the set of all $\pi \in \mathfrak{S}_N$ such that $\pi\mathcal{D} = \mathcal{D}$. We now recall a condition slightly stronger than regularity known as vertex-transitivity:

Definition B.4 (Vertex-transitive graph). \mathcal{D} is vertex-transitive if for each i and j , there exists $\pi \in \text{Aut}(\mathcal{D})$ such that $\pi(i) = j$.

Informally, a graph is vertex-transitive if and only if it ‘looks the same’ from every vertex. If a graph is vertex-transitive, then it is necessarily regular. The strongest form of symmetry for graphs that we consider here is the following:

Definition B.5 (Symmetric graph). \mathcal{D} is symmetric (or arc-transitive) if for each i, j with $\mathcal{D}_{ij} \neq 0$ and i', j' with $\mathcal{D}_{i'j'} \neq 0$, there exists $\pi \in \text{Aut}(\mathcal{D})$ such that $\pi(i) = i'$ and $\pi(j) = j'$.

A graph is symmetric if it ‘looks the same’ from any two directed edges. Arc-transitivity is typically defined for unweighted graphs, i.e. graphs satisfying $\mathcal{D} \in \{0, 1\}^{N \times N}$. For the more general class of weighted graphs, we require that \mathfrak{S}_N act transitively on the set of edges of \mathcal{D} , where ‘edge’ means a pair (i, j) with $\mathcal{D}_{ij} \neq 0$. Thus, all of the edges in a symmetric, weighted graph have the same weight: otherwise, if (i, j) and (i', j') are edges but $\mathcal{D}_{ij} \neq \mathcal{D}_{i'j'}$, then there would exist no permutation, π , sending i to i' , j to j' , and preserving the weights of the graph. Therefore, since the weights of a symmetric graph take one of two values (0 or else the only non-zero weight), such a graph is essentially unweighted.

B.2. Symmetries of evolutionary processes

In §2.1, we defined two states, s and s' , to be *evolutionarily equivalent* if (i) there exists an automorphism of the Markov chain, $\phi \in \text{Aut}(X)$, such that $\phi(s) = s'$, and (ii) this automorphism satisfies $\phi(\mu) = \mu$ for each stationary distribution, μ , of the chain. Condition (i), which means that s and s' are *symmetric*, alone is not quite strong enough to guarantee that s and s' have the same long-run behaviour. To give an example of a symmetry of states that is not an evolutionary equivalence, we consider the neutral Moran process in a well-mixed population of size $N = 3$:

Example B.6. In a well-mixed population of size $N = 3$, consider the (frequency-independent) Moran process with two types of players: a mutant type and a wild-type. Suppose that the mutant type is neutral with respect to the mutant; that is, the fitness of the mutant relative to the wild-type is 1. Since the population is well mixed, the state of the population is given by the number of mutants it contains, $i \in \{0, 1, 2, 3\} =: S$. Consider the map $\phi: S \rightarrow S$ defined by $\phi(i) = 3 - i$. States 0 and 3 are absorbing, and, for $i \in \{1, 2\}$, the transition probabilities of this process are as follows:

$$\mathbf{T}_{i,i-1} = \mathbf{T}_{i,i+1} = \binom{i}{3} \binom{3-i}{3} \quad (\text{B } 2a)$$

and

$$\mathbf{T}_{i,i} = \left(\frac{i}{3}\right)^2 + \left(\frac{3-i}{3}\right)^2. \quad (\text{B } 2b)$$

It follows at once that ϕ preserves these transition probabilities, so ϕ is an automorphism of the Markov chain. Let ρ_i be the probability that mutants fixate given an initial abundance of i mutants. The states 1 and 2 are symmetric since $\phi(1) = 1$, but it is *not* true that $\rho_1 = \rho_2$ since $\rho_1 = 1/3$ and $\rho_2 = 2/3$. The reason for this difference in fixation probabilities is that states 1 and 2, although symmetric, are not evolutionarily equivalent since ϕ swaps the two absorbing states of the process.

In contrast to example B.6, processes with unique stationary distributions have the property that every symmetry of the Markov chain is an evolutionary equivalence (proposition 2.3). The following lemma establishes proposition (2.3):

Lemma B.7. *If $\phi: S \rightarrow S$ is a symmetry of a Markov chain and μ is a stationary distribution of this chain, then $\phi(\mu)$ is also a stationary distribution. In particular, if μ is unique, then $\phi(\mu) = \mu$.*

Proof. If \mathbf{T} is the transition matrix of this Markov chain, then

$$\begin{aligned} [\phi(\mu)^T \mathbf{T}]_s &= \sum_{s' \in S} \phi(\mu)_{s'} \mathbf{T}_{s',s} \\ &= \sum_{s' \in S} \mu_{\phi(s')} \mathbf{T}_{\phi(s'),\phi(s)} \\ &= \sum_{s' \in S} \mu_{s'} \mathbf{T}_{s',\phi(s)} \\ &= [\mu^T \mathbf{T}]_{\phi(s)} \\ &= \mu_{\phi(s)} \\ &= \phi(\mu)_{s'} \end{aligned} \quad (\text{B } 3)$$

so $\phi(\mu)^T \mathbf{T} = \phi(\mu)^T$, which completes the proof. ■

We turn now to the proofs of our main results (theorems 3.7 and 3.12):

Theorem 3.7. *Consider an evolutionary matrix game on a graph, $\Gamma = (\mathcal{E}, \mathcal{D})$, with symmetric pay-offs and homogeneous strategy mutations. If $\pi \in \text{Aut}(\Gamma)$, then the states with a single mutant at vertex i and $\pi(i)$, respectively, in an otherwise-monomorphic population, are evolutionarily equivalent. That is, in the notation of definition 2.4, the states $s_{(s',i),s}$ and $s_{(s',\pi(i)),s}$ are evolutionarily equivalent for each $s, s' \in S$.*

Proof. The state space of the Markov chain defined by this evolutionary game is S^N , where S is the strategy set and N is the population size. For $\pi \in \text{Aut}(\Gamma)$, let π act on the

state of the population by changing the strategy of player i to that of player $\pi^{-1}(i)$ for each $i = 1, \dots, N$. In other words, π sends $s \in S^N$ to $\pi s \in S^N$, which is defined by $(\pi s)_i = s_{\pi^{-1}(i)}$ for each i . Therefore, for $s, s' \in S$, we have

$$\pi s_{(s',i),s} = s_{(s',\pi(i)),s}, \quad (\text{B } 4)$$

for each $i = 1, \dots, N$. Since π is an automorphism of the evolutionary graph, Γ , we have $\pi \mathcal{E} = \mathcal{E}$ and $\pi \mathcal{D} = \mathcal{D}$. Moreover, π preserves the strategy mutations since they are homogeneous. Since the pay-offs are symmetric, π just rearranges the fitness profile of the population: the pay-off of player i becomes the pay-off of player $\pi^{-1}(i)$ (see equation (3.2)), so the same is true of the fitness values. Therefore, applying the map π to S^N is equivalent to applying the map on S^N obtained by simply relabelling the players. Since any such relabelling of the players results in an automorphism of the Markov chain on S^N that preserves the monomorphic absorbing states, it follows that $s_{(s',i),s}$ and $s_{(s',\pi(i)),s}$ are evolutionarily equivalent. ■

Theorem 3.12. *Suppose that an asymmetric matrix game with homogeneous strategy mutations is played on an evolutionary graph, $\Gamma = (\mathcal{E}, \mathcal{D})$. For each $\pi \in \text{Aut}(\Gamma)$, $k \in \{1, 2, 3\}$ and $i, j \in \{1, \dots, N\}$,*

$$\sigma_k^{ij} = \sigma_k^{\pi(i)\pi(j)}. \quad (\text{B } 5)$$

Proof. Let \mathbf{T} be the transition matrix for the Markov chain defined by this process. Since there are non-zero strategy mutations, this chain has a unique stationary distribution, μ . The matrix \mathbf{T} defines a directed, weighted graph on $|S| = |S|^N$ vertices that has an edge from vertex s to vertex s' if and only if $\mathbf{T}_{s,s'} \neq 0$. If there is an edge from s to s' , then the weight of this edge is simply $\mathbf{T}_{s,s'}$. The (outdegree) Laplacian matrix of this graph, $\mathcal{L} = \mathcal{L}(\mathbf{T})$, is defined by $\mathcal{L} = I_{|S|} - \mathbf{T}$ [48]. In terms of this Laplacian matrix, Press & Dyson [47] show that for any vector, v , the stationary distribution satisfies

$$\mu \cdot v = \frac{\det \mathcal{L}(s, v)}{\det \mathcal{L}(s, \mathbf{1})}, \quad (\text{B } 6)$$

for each state, s , where $\mathcal{L}(s, v)$ denotes the matrix obtained from \mathcal{L} by replacing the column corresponding to state s by v . Thus, if ψ_r is the vector indexed by $S = S^N$ with $\psi_r(s)$ being the frequency of strategy r in state s , then the *average abundance* of strategy r is

$$F_r := \mu \cdot \psi_r = \frac{\det \mathcal{L}(s, \psi_r)}{\det \mathcal{L}(s, \mathbf{1})}. \quad (\text{B } 7)$$

Since \mathbf{T} is a function of the pay-offs, $\mathbf{a} = (a_{st}^{ij})_{s,t,i,j}$, we may write $F_r = F_r(\mathbf{a})$. (\mathbf{a} is just an ordered tuple defined by $\mathbf{a}_{st}^{ij} := a_{st}^{ij}$ for each s, t, i and j .) Moreover, since the entries of \mathbf{T} are assumed to be smooth functions of \mathbf{a} [34], F_r is also a smooth function of \mathbf{a} by equation (B7) and the definition of \mathcal{L} . We will show that for each s and t ,

$$\frac{\partial F_r}{\partial a_{st}^{ij}} \Big|_{\mathbf{a}=0} = \frac{\partial F_r}{\partial a_{st}^{\pi(i)\pi(j)}} \Big|_{\mathbf{a}=0}, \quad (\text{B } 8)$$

for each i and j . The theorem will then follow from the derivations of σ_1^{ij} , σ_2^{ij} and σ_3^{ij} in [33] since it is shown there that each σ^{ij} is a function of the elements in the set

$$\left\{ \frac{\partial F_r}{\partial a_{st}^{ij}} \Big|_{\mathbf{a}=0} \right\}_{s,t=1}^n. \quad (\text{B } 9)$$

Table 1. The absorption times of the 12 initial configurations of a single mutant in a wild-type population for the Moran process on the Frucht graph. The fitness of the mutant relative to the wild-type is $r = 2$.

initial vertex of mutant	absorption time
1	238.1836
2	237.0596
3	234.5982
4	235.8447
5	236.5967
6	234.5792
7	231.6988
8	238.0375
9	233.6122
10	235.1514
11	230.1340
12	228.7114

For s, t, i and j fixed and $a \in \mathbb{R}$, let $F_r^{s,t,i,j}(a)$ be the function of the vector with a at entry a_{st}^{ij} and 0 in all other entries. Symbolically, if $\delta_{x,y}$ is defined as 1 if $x = y$ and 0 otherwise and

$$\mathbf{a}_0^{(s,t,i,j)}(a) := (a\delta_{s,s'}\delta_{t,t'}\delta_{i,i'}\delta_{j,j'})_{s',t',i',j'} \quad (\text{B } 10)$$

then

$$F_r^{s,t,i,j}(a) := F_r(\mathbf{a}_0^{(s,t,i,j)}(a)). \quad (\text{B } 11)$$

Let $\pi \in \mathfrak{S}_N$ and suppose that $\pi \in \text{Aut}(\Gamma)$; that is, $\pi\mathcal{E} = \mathcal{E}$ and $\pi\mathcal{D} = \mathcal{D}$. π induces a map on the pay-offs, \mathbf{a} , defined by $\pi(a_{st}^{ij})_{s,t,i,j} = (a_{\pi(s)\pi(t)}^{\pi(i)\pi(j)})_{s,t,i,j}$. Let $\text{orb}_{\mathfrak{S}_N}(\mathbf{a})$ denote the orbit of \mathbf{a} under this action, and consider the enlarged state space $S' := S^N \times \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$. Using the Markov chain on S^N coming from the evolutionary process, we obtain a Markov chain on $S^N \times \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$ via the transition matrix, \mathbf{T}' , defined by

$$\mathbf{T}'_{(s,b),(s',b')} := \delta_{b,b'}\mathbf{T}_{s,s'}(b). \quad (\text{B } 12)$$

for $s, s' \in S^N$ and $\mathbf{b}, \mathbf{b}' \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$. (We write $\mathbf{T}(\mathbf{b})$ to indicate the transition matrix as a function of the pay-off values of the game.) π extends to a map on S' defined by $\pi(s, \mathbf{b}) = (\pi s, \pi \mathbf{b})$. Since π preserves \mathcal{E} , \mathcal{D} and the strategy mutations (since they are homogeneous), it follows that the induced map $\pi: S' \rightarrow S'$ is an automorphism of the Markov chain on S' defined by \mathbf{T}' . If μ' is a stationary distribution for the chain \mathbf{T}' , then, for each $s \in S^N$ and $\mathbf{b} \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$,

$$\begin{aligned} \mu_{(s,b)}' &= \sum_{s' \in S^N} \sum_{\mathbf{b}' \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})} \mu'_{(s',b')} \mathbf{T}'_{(s',b'),(s,b)} \\ &= \sum_{s' \in S^N} \sum_{\mathbf{b}' \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})} \mu'_{(s',b')} \delta_{\mathbf{b}',\mathbf{b}} \mathbf{T}_{s',s} \\ &= \sum_{s' \in S^N} \mu'_{(s',\mathbf{b})} \mathbf{T}_{s',s}. \end{aligned} \quad (\text{B } 13)$$

It then follows from the uniqueness of $\mu = \mu(\mathbf{b})$ that there exists $c_{\mu'}(\mathbf{b}) \geq 0$ such that $\mu'_{(s,b)} = c_{\mu'}(\mathbf{b})\mu_s$. If μ' is such a stationary distribution, then, by lemma (B.7), $\pi\mu' = \mu''$ for some other stationary distribution, μ'' , of the chain on S' . This equation implies that $c_{\mu''}(\mathbf{b}) = c_{\mu'}(\pi\mathbf{b})$ for each $\mathbf{b} \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$. Therefore,

$$\pi\mu(\pi\mathbf{b})_s = \mu(\pi\mathbf{b})_{\pi s} = \mu(\mathbf{b})_{s'} \quad (\text{B } 14)$$

Table 2. The fixation probabilities and absorption times of the 12 initial configurations of a single cooperator among defectors for the death–birth process on the Frucht graph. Pay-offs are frequency-dependent and derived from the Snowdrift Game, (3.4). The intensity of selection is $\beta = 1$.

initial vertex of mutant	fixation probability	absorption time
1	0.6505	116.0959
2	0.6471	115.4026
3	0.6469	115.7302
4	0.6448	115.7348
5	0.6463	116.0100
6	0.6562	117.4671
7	0.7299	129.8609
8	0.6512	116.6906
9	0.6545	118.3795
10	0.6551	117.8995
11	0.7344	131.6681
12	0.7326	131.9634

Table 3. The fixation probabilities and absorption times of the 12 initial configurations of a single cooperator among defectors for the death–birth process on the Tietze graph. Pay-offs are frequency-dependent and derived from the Snowdrift Game, (3.4). The intensity of selection is $\beta = 0.1$. These values are illustrated graphically in figure 8.

initial vertex of mutant	fixation probability	absorption time
1	0.3777	70.7869
2	0.3777	70.7869
3	0.3777	70.7869
4	0.4141	76.5048
5	0.4186	77.3094
6	0.4186	77.3094
7	0.4141	76.5048
8	0.4186	77.3094
9	0.4186	77.3094
10	0.4141	76.5048
11	0.4186	77.3094
12	0.4186	77.3094

for each $s \in S^N$ and $\mathbf{b} \in \text{orb}_{\mathfrak{S}_N}(\mathbf{a})$. Consequently, since $\pi\psi_r = \psi_r$,

$$\begin{aligned} F_r^{s,t,i,j}(a) &= F_r(\mathbf{a}_0^{(s,t,i,j)}(a)) \\ &= \mu(\mathbf{a}_0^{(s,t,i,j)}(a)) \cdot \psi_r \\ &= \pi\mu(\pi\mathbf{a}_0^{(s,t,i,j)}(a)) \cdot \psi_r \\ &= \pi\mu(\mathbf{a}_0^{(s,t,\pi(i),\pi(j))}(a)) \cdot \pi\psi_r \\ &= \mu(\mathbf{a}_0^{(s,t,\pi(i),\pi(j))}(a)) \cdot \psi_r \\ &= F_r(\mathbf{a}_0^{(s,t,\pi(i),\pi(j))}(a)) \\ &= F_r^{s,t,\pi(i),\pi(j)}(a). \end{aligned} \quad (\text{B } 15)$$

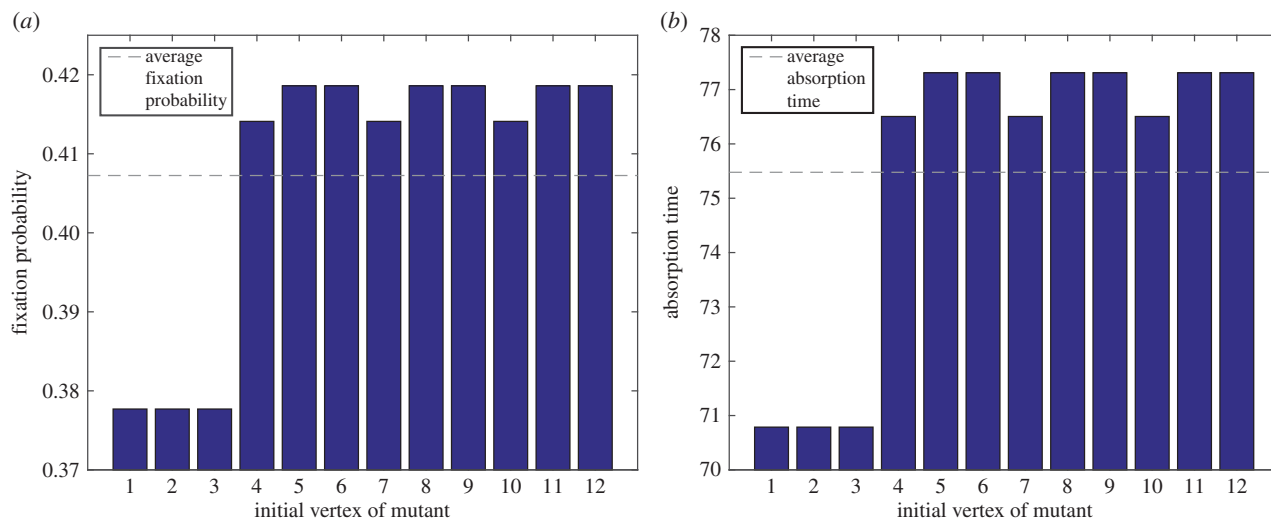


Figure 8. (a) Fixation probability and (b) absorption time versus initial vertex of mutant (cooperator) for a death–birth process on the Tietze graph. In both figures, the game is a Snowdrift Game whose pay-offs are described by pay-off matrix (3.4), and the selection intensity is $\beta = 0.1$. This example illustrates the single-mutant states that are not evolutionarily equivalent in the Tietze graph. Moreover, it happens to be the case that any two of these states with the same fixation probability (or absorption time) are evolutionarily equivalent. (Online version in colour.)

As a result, we have

$$\begin{aligned} \left. \frac{\partial F_r}{\partial a_{st}^{ij}} \right|_{a=0} &= \left. \frac{d}{da} \right|_{a=0} F_r^{s,t,ij} = \left. \frac{d}{da} \right|_{a=0} F_r^{s,t,\pi(i),\pi(j)} \\ &= \left. \frac{\partial F_r}{\partial a_{\pi(i)\pi(j)}} \right|_{a=0}, \end{aligned} \quad (\text{B } 16)$$

so equation (B 8) holds, which completes the proof. ■

Appendix C. Explicit calculations

We now perform explicit calculations using equations (A 6) and (A 8) to show that the Isothermal Theorem extends to neither absorption times nor frequency-dependent games. We also calculate the structure coefficients for the death–birth process on the graph of figure 7a to show that corollary 3.14 does not necessarily hold for graphs that are vertex-transitive but not symmetric.

C.1. The Moran process

Consider the Moran process on the Frucht graph (figure 2), and suppose that the mutant type has fitness $r > 0$ relative to the wild-type. By the Isothermal Theorem of Lieberman *et al.* [3], the fixation probability of a fixed number of mutants is independent of the configuration of those mutants on the graph. For $r = 2$, the absorption times (of configurations of a single mutant in a wild-type population) are listed in table 1. The fixation probability of a single mutant is $(1 - 1/2)/(1 - 1/2^{12}) \approx 0.5001$ for every vertex. Thus, unlike fixation probabilities, absorption times depend on the initial location of the mutant. (Some of the absorption times are similar in this case, but no two are the same.)

C.2. Frequency-dependent games

C.2.1. Symmetric games

Consider the instance of the Snowdrift Game that has for a pay-off matrix (3.4). For this game, table 2 gives the fixation probabilities and the absorption times (rounded to four digits after the decimal point) for the death–birth process on the Frucht graph (figure 2) with $\beta = 1$.

Table 4. The structure coefficients in equation (3.12) for the death–birth process on the vertex-transitive (but not symmetric) graph of figure 7a with homogeneous strategy-mutation rate $\varepsilon = 0.01$.

(i, j)	τ_1^{ij}	τ_2^{ij}
(1, 2)	707 905/9 315 552	32 989/405 024
(1, 4)	16 291/194 074	57 057/776 296
(1, 6)	707 905/9 315 552	32 989/405 024
(2, 1)	707 905/9 315 552	32 989/405 024
(2, 3)	16 291/194 074	57 057/776 296
(2, 6)	707 905/9 315 552	32 989/405 024
(3, 2)	16 291/194 074	57 057/776 296
(3, 4)	707 905/9 315 552	32 989/405 024
(3, 5)	707 905/9 315 552	32 989/405 024
(4, 1)	16 291/194 074	57 057/776 296
(4, 3)	707 905/9 315 552	32 989/405 024
(4, 5)	707 905/9 315 552	32 989/405 024
(5, 3)	707 905/9 315 552	32 989/405 024
(5, 4)	707 905/9 315 552	32 989/405 024
(5, 6)	16 291/194 074	57 057/776 296
(6, 1)	707 905/9 315 552	32 989/405 024
(6, 2)	707 905/9 315 552	32 989/405 024
(6, 5)	16 291/194 074	57 057/776 296

Similarly, for the same game (and update rule) but on the Tietze graph (figure 6) with $\beta = 0.1$, table 3 and figure 8 give the fixation probabilities and absorption times for all possible configurations of a single cooperator among defectors.

C.2.2. Asymmetric games

For the death–birth process on the graph in figure 7a with homogeneous strategy-mutation rate $\varepsilon = 0.01$, we calculate the complete collection of structure coefficients $\{\tau_1^{ij}, \tau_2^{ij}\}_{i,j}$ (for $r = 1$) as follows: let ψ_1 be the vector indexed by \mathcal{S} with $\psi_1(s)$ being the frequency of strategy 1 in state s , and let 1

be the vector of ones. McAvoy & Hauert [33] show that, for any \mathfrak{s} , equation (3.11) is equivalent to

$$\frac{1}{n} \operatorname{tr} \left(([\mathcal{L}(\mathfrak{s}, \psi_1)]_{\beta=0}]^{-1} - [\mathcal{L}(\mathfrak{s}, \mathbf{1})]_{\beta=0}]^{-1} \right) \frac{d}{d\beta} \Big|_{\beta=0} \mathcal{L}(\mathfrak{s}, \mathbf{0}) > 0, \quad (\text{C } 1)$$

where $\mathcal{L} = \mathcal{L}(\mathbf{T}) = I_{|S|} - \mathbf{T}$ is the (outdegree) Laplacian matrix of the graph defined by \mathbf{T} . If $k \in \{1, 2\}$, i and j are fixed, and we choose the (two-strategy) asymmetric game so that

$$a_{st}^{i'j'} = \begin{cases} 1 & s = 1, t = k, i' = i, j' = j; \\ 0 & \text{otherwise,} \end{cases} \quad (\text{C } 2)$$

then

$$\tau_k^{ij} = \frac{1}{2} \operatorname{tr} \left(([\mathcal{L}(\mathfrak{s}, \psi_1)]_{\beta=0}]^{-1} - [\mathcal{L}(\mathfrak{s}, \mathbf{1})]_{\beta=0}]^{-1} \right) \frac{d}{d\beta} \Big|_{\beta=0} \mathcal{L}(\mathfrak{s}, \mathbf{0}). \quad (\text{C } 3)$$

Using this method, we obtain the structure coefficients for figure 7a listed in table 4. For the same process on the symmetric graph of figure 7b, we find that τ_1^{ij} and τ_2^{ij} are independent of i and j and are both equal to 2189/27 728.

References

- Moran PAP. 1958 Random processes in genetics. *Math. Proc. Camb. Phil. Soc.* **54**, 60. (doi:10.1017/s0305004100033193)
- Nowak MA. 2006 *Evolutionary dynamics: exploring the equations of life*. Cambridge, MA: Belknap Press.
- Lieberman E, Hauert C, Nowak MA. 2005 Evolutionary dynamics on graphs. *Nature* **433**, 312–316. (doi:10.1038/nature03204)
- Bollobás B. 2001 *Random graphs*. Cambridge, UK: Cambridge University Press.
- Maruyama T. 1974 A simple proof that certain quantities are independent of the geographical structure of population. *Theor. Popul. Biol.* **5**, 148–154. (doi:10.1016/0040-5809(74)90037-9)
- Taylor PD, Jonker LB. 1978 Evolutionary stable strategies and game dynamics. *Math. Biosci.* **40**, 145–156. (doi:10.1016/0025-5564(78)90077-9)
- Hofbauer J, Sigmund K. 1998 *Evolutionary games and population dynamics*. Cambridge, UK: Cambridge University Press.
- Nowak MA, Sasaki A, Taylor C, Fudenberg D. 2004 Emergence of cooperation and evolutionary stability in finite populations. *Nature* **428**, 646–650. (doi:10.1038/nature02414)
- Taylor C, Fudenberg D, Sasaki A, Nowak MA. 2004 Evolutionary game dynamics in finite populations. *Bull. Math. Biol.* **66**, 1621–1644. (doi:10.1016/j.bulm.2004.03.004)
- Ohtsuki H, Hauert C, Lieberman E, Nowak MA. 2006 A simple rule for the evolution of cooperation on graphs and social networks. *Nature* **441**, 502–505. (doi:10.1038/nature04605)
- Szabó G, Fáth G. 2007 Evolutionary games on graphs. *Phys. Rep.* **446**, 97–216. (doi:10.1016/j.physrep.2007.04.004)
- Tarnita CE, Antal T, Ohtsuki H, Nowak MA. 2009 Evolutionary dynamics in set structured populations. *Proc. Natl Acad. Sci. USA* **106**, 8601–8604. (doi:10.1073/pnas.0903019106)
- Taylor PD, Irwin AJ, Day T. 2001 Inclusive fitness in finite deme-structured and stepping-stone populations. *Selection* **1**, 153–164. (doi:10.1556/select.1.2000.1-3.15)
- Hauert C, Imhof L. 2012 Evolutionary games in deme structured, finite populations. *J. Theor. Biol.* **299**, 106–112. (doi:10.1016/j.jtbi.2011.06.010)
- Zukewich J, Kurella V, Doebeli M, Hauert C. 2013 Consolidating birth–death and death–birth processes in structured populations. *PLoS ONE* **8**, e54639. (doi:10.1371/journal.pone.0054639)
- Ohtsuki H, Nowak MA. 2006 The replicator equation on graphs. *J. Theor. Biol.* **243**, 86–97. (doi:10.1016/j.jtbi.2006.06.004)
- Szabó G, Töke C. 1998 Evolutionary Prisoner's dilemma game on a square lattice. *Phys. Rev. E* **58**, 69–73. (doi:10.1103/physreve.58.69)
- Traulsen A, Pacheco JM, Nowak MA. 2007 Pairwise comparison and selection temperature in evolutionary game dynamics. *J. Theor. Biol.* **246**, 522–529. (doi:10.1016/j.jtbi.2007.01.002)
- Ewens WJ. 2004 *Mathematical population genetics*. New York, NY: Springer.
- Imhof LA, Nowak MA. 2006 Evolutionary game dynamics in a Wright-Fisher process. *J. Math. Biol.* **52**, 667–681. (doi:10.1007/s00285-005-0369-8)
- McAvoy A, Hauert C. 2015 Asymmetric evolutionary games. *PLoS Comp. Biol.* **11**, e1004349. (doi:10.1371/journal.pcbi.1004349)
- Fudenberg D, Imhof LA. 2006 Imitation processes with small mutations. *J. Econ. Theory* **131**, 251–262. (doi:10.1016/j.jet.2005.04.006)
- Wu B, Gokhale CS, Wang L, Traulsen A. 2011 How small are small mutation rates? *J. Math. Biol.* **64**, 803–827. (doi:10.1007/s00285-011-0430-8)
- Frucht R. 1939 Herstellung von graphen mit vorgegebener abstrakter gruppe. *Compos. Math.* **6**, 239–250.
- Ohtsuki H, Nowak MA, Pacheco JM. 2007 Breaking the symmetry between interaction and replacement in evolutionary dynamics on graphs. *Phys. Rev. Lett.* **98**, 108106. (doi:10.1103/physrevlett.98.108106)
- Taylor PD, Day T, Wild G. 2007 Evolution of cooperation in a finite homogeneous graph. *Nature* **447**, 469–472. (doi:10.1038/nature05784)
- Ohtsuki H, Pacheco JM, Nowak MA. 2007 Evolutionary graph theory: breaking the symmetry between interaction and replacement. *J. Theor. Biol.* **246**, 681–694. (doi:10.1016/j.jtbi.2007.01.024)
- Pacheco JM, Pinheiro FL, Santos FC. 2009 Population structure induces a symmetry breaking favoring the emergence of cooperation. *PLoS Comp. Biol.* **5**, e1000596. (doi:10.1371/journal.pcbi.1000596)
- Débarre F, Hauert C, Doebeli M. 2014 Social evolution in structured populations. *Nat. Commun.* **5**, 3409. (doi:10.1038/ncomms4409)
- Bondy JA, Murty USR. 2008 *Graph theory*. London, UK: Springer.
- Hofbauer J. 1996 Evolutionary dynamics for bimatrix games: a Hamiltonian system? *J. Math. Biol.* **34**, 675–688. (doi:10.1007/bf02409754)
- Ohtsuki H. 2010 Stochastic evolutionary dynamics of bimatrix games. *J. Theor. Biol.* **264**, 136–142. (doi:10.1016/j.jtbi.2010.01.016)
- McAvoy A, Hauert C. 2015 Structure coefficients and strategy selection in multiplayer games. *J. Math. Biol.* (doi:10.1007/s00285-015-0882-3)
- Tarnita CE, Wage N, Nowak MA. 2011 Multiple strategies in structured populations. *Proc. Natl Acad. Sci. USA* **108**, 2334–2337. (doi:10.1073/pnas.1016008108)
- Komarova NL. 2006 Spatial stochastic models for cancer initiation and progression. *Bull. Math. Biol.* **68**, 1573–1599. (doi:10.1007/s11538-005-9046-8)
- Kaveh K, Komarova NL, Kohandel M. 2015 The duality of spatial death–birth and birth–death processes and limitations of the isothermal theorem. *R. Soc. open sci.* **2**, 140465. (doi:10.1098/rsos.140465)
- Matsuda H, Ogita N, Sasaki A, Sato K. 1992 Statistical mechanics of population: the lattice Lotka–Volterra model. *Progr. Theor. Phys.* **88**, 1035–1049. (doi:10.1143/ptp/88.6.1035)
- Assaf M, Mobilia M. 2012 Metastability and anomalous fixation in evolutionary games on scale-free networks. *Phys. Rev. Lett.* **109**, 188701. (doi:10.1103/physrevlett.109.188701)
- Szolnoki A, Perc M, Mobilia M. 2014 Facilitators on networks reveal optimal interplay between information exchange and reciprocity. *Phys. Rev. E* **89**, 042802. (doi:10.1103/physreve.89.042802)
- Tarnita CE, Taylor PD. 2014 Measures of relative fitness of social behaviors in finite structured population models. *Amer. Nat.* **184**, 477–488. (doi:10.1086/677924)
- Beineke LW, Wilson RJ, Cameron PJ (eds). 2004 *Topics in algebraic graph theory*. Cambridge, UK: Cambridge University Press.
- Roca CP, Cuesta JA, Sánchez A. 2009 Evolutionary game theory: temporal and spatial effects beyond

- replicator dynamics. *Phys. Life Rev.* **6**, 208–249. (doi:10.1016/j.plrev.2009.08.001)
43. Hindersin L, Traulsen A. 2014 Counterintuitive properties of the fixation time in network-structured populations. *J. R. Soc. Interface* **11**, 20140606. (doi:10.1098/rsif.2014.0606)
44. Cheng D, He F, Qi H, Xu T. 2015 Modeling, analysis and control of networked evolutionary games. *IEEE Trans. Automatic Control* **60**, 2402–2415. (doi:10.1109/tac.2015.2404471)
45. Traulsen A, Santos FC, Pacheco JM. 2009 Evolutionary games in self-organizing populations. In *Adaptive networks* (eds T Gross, H Sayama), pp. 253–267. Berlin, Germany: Springer Science+Business Media.
46. Taylor C, Iwasa Y, Nowak MA. 2006 A symmetry of fixation times in evolutionary dynamics. *J. Theor. Biol.* **243**, 245–251. (doi:10.1016/j.jtbi.2006.06.016)
47. Press WH, Dyson FJ. 2012 Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl Acad. Sci. USA* **109**, 10 409–10 413. (doi:10.1073/pnas.1206569109)
48. Chung FRK. 1996 *Spectral graph theory*. Providence, RI: American Mathematical Society.