The Evolutionary Price of Anarchy: Locally Bounded Agents in a Dynamic Virus Game

Laura Schmid ![ORCID]
IST Austria, Klosterneuburg, Austria
lschmid@ist.ac.at

Krishnendu Chatterjee
IST Austria, Klosterneuburg, Austria
krishnendu.chatterjee@ist.ac.at

Stefan Schmid ![ORCID]
Faculty of Computer Science, University of Vienna, Austria
stefan_schmid@univie.ac.at

─────── **Abstract** ───────

The Price of Anarchy (PoA) is a well-established game-theoretic concept to shed light on coordination issues arising in open distributed systems. Leaving agents to selfishly optimize comes with the risk of ending up in sub-optimal states (in terms of performance and/or costs), compared to a centralized system design. However, the PoA relies on strong assumptions about agents' rationality (e.g., resources and information) and interactions, whereas in many distributed systems agents interact locally with bounded resources. They do so repeatedly over time (in contrast to "one-shot games"), and their strategies may evolve.

Using a more realistic evolutionary game model, this paper introduces a realized *evolutionary Price of Anarchy (ePoA)*. The ePoA allows an exploration of equilibrium selection in dynamic distributed systems with multiple equilibria, based on *local* interactions of simple *memoryless* agents.

Considering a fundamental game related to virus propagation on networks, we present analytical bounds on the ePoA in basic network topologies and for different strategy update dynamics. In particular, deriving stationary distributions of the stochastic evolutionary process, we find that the Nash equilibria are not always the most abundant states, and that different processes can feature significant off-equilibrium behavior, leading to a significantly higher ePoA compared to the PoA studied traditionally in the literature.

1 Introduction

The performance and efficiency of large distributed systems, such as open peer-to-peer networks which any user can join, often critically depend on cooperation and can suffer if users behave selfishly, e.g.: consume but not contribute resources [16], choose routes [21] and neighbors [7] to optimize their *personal* benefits, etc. Non-cooperative behavior can also affect security. For example, if security mechanisms such as the installation of anti-virus software are employed just for self-protection, a virus may spread further than in cooperative environments [8], while at the same time increasing global security investment costs [2].

The *Price of Anarchy (PoA)* [14] is a game-theoretic concept which allows to assess to which extent a distributed system is affected negatively by non-cooperative behavior. Essentially, the PoA compares the optimal social welfare resp. cost to the welfare resp. cost in the worst Nash equilibrium: an equilibrium in which no selfish agent, henceforth called *player*, has an incentive to change its behavior. If a distributed system has a large PoA, this means that the performance resp. cost can be far from optimal: the system may require a redesign or at least strong incentive mechanisms.

However, Nash equilibria are often not a good model for real-world distributed systems, for several reasons. In particular:

1. **Dynamic character:** Distributed systems typically are not based on "one-shot games" but rely on dynamic interactions over time: e.g., peers (i.e., *players*) in peer-to-peer systems such as BitTorrent interact repeatedly, for example using tit-for-tat strategies, leading to *repeated games* [29].

2. **Local information:** Players in distributed systems often only have local information about the network, based on interactions (e.g., with neighboring players). Obtaining global information is typically infeasible, especially in large-scale distributed systems.

3. **Bounded resources:** Players typically also have only limited resources, e.g., in terms of memory or in terms of the complexity of the kind of algorithms they can execute.

This paper is motivated by the desire to extend the concept of Price of Anarchy to account for these important characteristics of distributed systems. While the research community is currently discussing alternatives to Nash equilibria such as Bayes-Nash equilibria [23] for games with incomplete information, it is believed that such extensions are complex due to having to take into account players' belief systems, amongst other issues, and thus introduce major research challenges [20]. In contrast, we aim to port the PoA to *evolutionary games* in purely *local information* scenarios: games embedded in dynamical processes, where simple, *memoryless* players interact repeatedly and locally, and can update their strategies over time. This allows understanding a system's behavior over time, and helps answer questions about equilibrium selection, convergence, and possible off-equilibrium behavior.

However, analyzing equilibrium selection in stochastic processes described by evolutionary games is already challenging in 2-player games in a population with m pure strategies [10, 9]. Games *on graphs*, while also highly useful in verification and synthesis of (possibly distributed [19]) reactive systems [5, 17], are often particularly difficult, also when considering evolutionary games [1], due to additional dependencies on the possible interactions.

Our contributions.   This paper extends the notion of price of anarchy to evolutionary games, introducing the *evolutionary Price of Anarchy (ePoA)*. In particular, we are interested in the equilibrium behavior of simple memoryless players, which repeatedly and *locally* interact on a graph. For this setting, the ePoA is a more natural measure of efficiency than the static PoA: it allows considering agents that don't have perfect information when making decisions, and don't necessarily always reach equilibria in the game they are playing. We note that the ePoA is essentially a framework; it can be used to study games under different evolutionary dynamics, with different parameters and different ways in which players evolve their strategies.

To shed light on how the evolutionary perspective can affect the conclusion on the possible impact of selfish behavior in distributed systems, we consider a concrete case study: the well-known virus propagation game introduced by Aspnes et al. [2]. We present an analysis of the evolutionary dynamics of this game for the three fundamental dynamic models

(related to genetic evolution and imitation dynamics) and different basic network topologies. Interestingly, while the analysis of such evolutionary games is challenging in general, we are able to provide an exact characterization of the long-run frequencies of configurations for these scenarios.

We make several interesting observations. We find that the evolutionary dynamics of this game give rise to a rich variety of behaviors. In particular, the ePoA can be significantly worse than the classic PoA, for reasonable (i.e. not too high) mutation rates. We also find that Nash equilibria are not always the most frequent, i.e., *abundant*, states, and different processes can feature significant off-equilibrium behavior.

Our analytical results are complemented with simulations, also of more complicated topologies. For the full version of the paper, including technical details, as well as a more detailed discussion of related work, see [22]. Full results and our implementation are available upon request.

2   Preliminaries

Before launching into a full description and analysis of our model, we first revisit the virus inoculation game which will serve us as a case study in this paper. We also give a short introduction to evolutionary dynamics, evolutionary games and evolutionary graph theory.

2.1   The virus inoculation game

In the classic one-shot virus inoculation game [2], nodes must choose between installing anti-virus software (inoculating themselves) at a cost, or risk infection by a virus that spreads from a random location and can reach a node if there is a path of not inoculated nodes in between. The network is modeled by an undirected graph $G = (V, E)$ with N nodes. Nodes correspond to players in the game. Every node is equipped with a strategy a_i that denotes its propensity to inoculate itself. If a_i is 0 or 1, the strategy is called *pure*. Every node takes an action according to its strategy; the overall configuration of nodes is reflected in the strategy profile $\vec{a} \in [0, 1]^n$. Inoculation costs V. After everyone has made their choice, the adversary picks a random node as the starting point of infection. The virus propagates through the network, infecting all non-inoculated nodes that have a direct path to an infected node. That is, propagation happens on an "attack graph" $G_{\vec{a}}$, where inoculated nodes have been removed, and only insecure nodes remain. Being infected then comes with a cost $I > V$. Inoculation prevents infection as well as further virus transmission by the inoculated node.

The cost of a mixed strategy for a node i in this model is therefore

$$\text{cost}_i(\vec{a}) = a_i V + (1 - a_i) I \cdot p_i(\vec{a}), \tag{1}$$

where $p_i(\vec{a})$ is the probability of node i being infected given the strategy profile $\vec{a}$ and the condition that i did not inoculate. The goal of each player is to minimize its own cost, while it does not take the resulting social cost to the system in total into account. This social cost is simply

$$\text{cost}(\vec{a}) = \sum_{j=0}^{N-1} \text{cost}_j(\vec{a}). \tag{2}$$

Aspnes et al. then showed the following characterization of pure equilibria (for the proof and the extension to mixed equilibria, cf. [2]):

▶ **Corollary 1** (Characterization of pure equilibria). *Fix V and I, and let the threshold be $t = VN/I$. A strategy profile $\vec{a}$ is a pure Nash equilibrium if and only if:*
(a) *Every component in the attack graph $G_{\vec{a}}$ has at most size t.*
(b) *Inserting any secure node j and its edges into $G_{\vec{a}}$ yields a component of size at least t.*

2.2   Evolutionary dynamics and games

Game theory considers individuals that consciously aim to reach the best outcome for them in a strategic decision situation. Its classic framework, using concepts like Nash equilibria usually makes some key assumptions about the players' rationality, their beliefs and cognitive abilities. In contrast, *evolutionary game theory*, as a generic approach to evolutionary dynamics [24, 25, 13], considers a population of players with bounded rationality instead. Each player adopts a strategy to interact with other population members in a game. The players' payoffs from these interactions – which depend on the actions of the co-players and therefore on the abundance of different strategies – are considered to be their *evolutionary fitness*. Success in the game is translated into reproductive success: good strategies reproduce faster, whereas disadvantageous strategies go extinct. In a nutshell, evolutionary game theory describes dynamics that are dependent on the frequency of different strategies in a population. The evolutionary dynamics then depends on the setup and structure of the population, the underlying game, and the way strategies spread. *Evolutionary graph theory* considers different population structures and how the dynamics are changed by changing the underlying graph [15].

3   The Evolutionary Price of Anarchy

It may not be realistic to assume that nodes in a game such as the virus game in [2] will have perfect information in the first place. In large networks, it is highly unlikely that nodes separated by many neighbors would know G and each others' decision and would then optimally react to this immense amount of information. Rather, it is more natural to think that nodes only have *local* information at most. They can see their neighbors' most recent choices and react to them only when updating their strategy, while being unaware of the choices of nodes with a higher degree of separation. Thus, our model does not use the assumption of nodes having full information or memory; at most, they need to compare their payoff with the payoffs of their neighbors. To model this, we first introduce an evolutionary virus inoculation game and three different kinds of stochastic evolutionary dynamics. We then define a general notion of the evolutionary price of anarchy.

3.1   The evolutionary virus inoculation game

We consider an evolutionary process on a static graph $G = (V, E)$ with V the set of vertices (players) and E the set of edges. The $N = |V|$ vertices are occupied by players in the game, and we say that two nodes/players are *neighbors* if there is an edge connecting them. One iteration of the evolutionary process consists of three stages:
1. *Decision making.* All players make a decision whether to inoculate themselves against possible virus infections in the network. In case they choose to do so, they pay a cost V, and pay nothing otherwise. Players' propensity to inoculate is encoded in their strategy a_i.
2. *Virus propagation.* After everyone has concurrently made a decision, expected costs of the nodes when the system is attacked by a virus are calculated. To do so, we use a process with n steps: in each step, the virus starts at a different node of the graph and

spreads throughout the network. Inoculated players are unaffected by the virus and cannot transmit it to their neighbors, while unprotected players pay a cost once they become infected, and spread the virus to their other insecure neighbors. Uninfected players who are not inoculated do not pay anything. We will use the term "realized cost vector" to describe the vector $ct = [I, V, L = 0]$, where infected nodes pay I, inoculated nodes pay V, and insecure but uninfected nodes pay nothing. Once the virus has swept over the system, infecting unprotected players and their unprotected neighbors, costs are recorded, infection status is reset, and the next virus starts at the next node, until every node has served as a starting point. Alternatively, the process can be randomized by letting sufficiently many viruses start at random nodes. Once this has happened, cumulative costs are averaged, giving expected negative payoffs for the current strategy configuration, and the next stage commences.

3. *Evolution of strategies.* After receiving their expected negative payoff, players can assess the damage done and subsequently change their strategy based on comparisons with their neighbors' payoffs, before they again decide which action to take, and the next game begins. It is important to realize that this updating process is based on purely local information: nodes only need to know their neighbors' payoffs in order to make their decisions. This means that also the outcomes that can be realized may differ from the Nash equilibria that are found in the perfect information model of Aspnes et al. We consider dynamics that can take both *selection* and *mutation* into account: strategies typically evolve according to their (frequency-dependent) fitness, but they can also randomly change with a small probability μ, which we will refer to as the *mutation rate*. This prevents absorbing states in the process and lets us compute a *unique* stationary distribution that gives the long-term frequencies of system configurations. In the limiting case of $\mu \to 0$, the process always terminates and reaches a state where either all nodes are inoculated, or none are.

We differentiate here between two kinds of well known memoryless evolutionary dynamics, but are not restricted by them; in general we can configure our framework for general monotone imitation dynamics as described in [9]:

(a) Genetic evolution: On one hand, we consider genetic evolution as described by the *Moran process.* In this context, we analyze two different variants: a *death-birth (DB)* and *birth-death (BD)* scenario, respectively (cf [26] and [22]). In the DB scenario, a node is picked to die in each time step of the process. The vacancy is then filled by a copy of one of its neighbors, with the probability of one node being the replacement in some (possibly non-linear) way proportional to its payoff, such that nodes with higher payoffs (or rather, fewer losses) have a higher chance of being chosen as the replacement. In the BD scenario, meanwhile, first a node is picked for reproduction in each round with probability proportional to its payoff. This node subsequently chooses one of its neighbors uniformly at random and replaces it with a copy of itself. After every update, payoffs are recomputed. To visualize an example of such a process, we illustrate the DB scenario in Fig. 1.

(b) Cultural evolution: On the other hand, we also consider "cultural" evolution through imitation dynamics in the form of a standard pairwise comparison process ([27] and [22]). In this well established model of evolution, a focal player picks a neighboring "role model" node uniformly at random in every time step, observes its payoff, and switches to its strategy with probability

$$\varrho = \frac{1}{1 + e^{-\beta(\pi' - \pi)}} \tag{3}$$

where π' is the payoff of the neighbor and π the node's own payoff. This function is parameterized by the *selection strength* $\beta \geq 0$, which is a measure for the noise in the update probability, and with it, how much the payoff difference influences the dynamics. Thus, for $\beta = 0$, updating is random with probability $\varrho = 1/2$, whereas for $\beta > 0$, strategies with lower costs are preferred.

These processes are simulated until *timeout*: once there have been k update steps (where k is large to ensure convergence), we calculate the average welfare of the population (which is the average sum of payoffs), as well as the average count of how often the system visited the different states, and return.
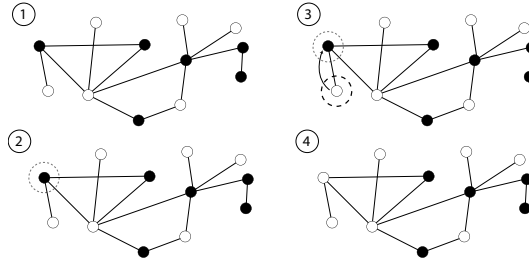
■ Figure 1 We illustrate the evolutionary dynamics given by the Moran Death-Birth process. Step 1: The nodes in the network shown use either Strategy 1 (white) or Strategy 2 (black). Step 2: One random node – in our example, a black one – is picked for death (visualized by the grey dashed line). Step 3: One of the nodes that neighbor the new vacancy is picked for reproduction, depending on its payoff. Here, this is a white node. Step 4: The reproducing node has passed on its strategy, such that there is one more white node on the graph.

3.2   The evolutionary price of anarchy

In the analysis of an evolutionary game, a fundamental question of interest is to which distribution different dynamics converge to, for various topologies and parameters. Such a stationary distribution of the Markov chain underlying the evolutionary dynamics contains the long-run frequencies of states. This corresponds to the long-run probabilities of finding the system in the different states. Since we obtain an ergodic process on the space of all possible population compositions (hereon called configurations) by our elementary updating rules, this limiting distribution of inoculation states exists and is unique. It is also called the *selection-mutation equilibrium* $\mathbf{x}$ of a given evolutionary process, and forms the basis of further analysis - we emphasize here that our results are therefore based on an actual limit distribution and not an arbitrary cutoff time. We note that it is the nonzero mutation rate $\mu > 0$ that provides the necessary ergodicity for a unique stationary distribution to exist – otherwise, the Markov chain would have absorbing states where all nodes are inoculated or none are.

We can subsequently find the average social cost $\hat{S}$ for any of the dynamics we use, either by averaging over the total cost in each round (when the process is simulated), or multiplying $\mathbf{x}$ with the vector $\mathbf{R}$ containing the cost of all possible configurations, such that

$$\hat{S} = \mathbf{x} \cdot \mathbf{R} = \sum_i x_i \, R_i, \tag{4}$$

where $R_i = \sum_{j=1}^{N} \hat{\pi}_j^i$, and $\hat{\pi}_j^i$ is the average payoff of player j in configuration i. We measure the efficiency of a process by comparing the average social cost $\hat{S}$ with the optimum Ω. At this point, we introduce the concept of the *evolutionary price of anarchy ePoA* as the ratio

of the average social cost (or payoffs) of a process against the social optimum. Similarly to the static PoA, we hence define, for a particular evolutionary dynamics and assuming negative payoffs:

$$ePoA = \hat{S}/\Omega. \tag{5}$$

For positive payoffs, we define $ePoA = \Omega/\hat{S}$ and note that in both cases, $ePoA \geq 1$, as the static PoA. In general, this quantity gives an indication which processes are most conducive to spending large fractions of time in configurations that minimize overall cost (or maximize positive payoffs). We also note that in principle, the evolutionary PoA can be both smaller or larger than the static PoA.

Note that the concept of an evolutionary price of anarchy is neither bound to a particular game nor a particular evolutionary dynamics: it can serve as a general framework for analyzing a rich variety of dynamic games with arbitrary update dynamics.

4 Results and Analysis

In the following, we will consider pure strategies and the analysis of a setting with positive mutation rate $\mu > 0$. We will first show how to exactly calculate the selection-mutation equilibrium **x** of the evolutionary process for two fundamental graphs, and then use this to show how the ePoA can differ from a static analysis.

4.1 Analytical results

For simple graphs and pure strategies, $a_i \in \{0, 1\}$, we can calculate the stationary distribution of the underlying Markov chain under the different dynamics. We consider two instructive cases here (similar cases have been studied also in [2]), situated at the opposite ends of the spectrum: the *clique* (which results in perfectly global information) and the *star* (which is fully local). In these cases, we find exact results.

4.1.1 Clique

In a clique, the $(N + 1)$ states of the Markov chain are $i = 0, \ldots, N$, denoting the number of currently inoculated nodes. Here, the threshold for infection of an arbitrary insecure node is 1, as an infected node automatically spreads the virus to all other not inoculated nodes. We use the entries of the cost vector $ct = [I, V, L = 0]$ as the realized negative payoffs of infected (I), inoculated (V) and unaffected ($L = 0$) nodes. For the expected payoffs $\hat{\pi}_X^i$ of nodes using strategy X, in the state i, with X either C ($a_i = 1$) or D ($a_i = 0$), we then find the following simple expressions:

$$\hat{\pi}_C^i = V, \text{ and } \hat{\pi}_D^i = \frac{i}{N}L + \frac{N-i}{N}I = \frac{N-i}{N}I \tag{6}$$

The next step is calculating the transition probabilities of the resulting Markov chain, $p_{i,i+1}$ and $p_{i,i-1}$, describing changes from states i to $i+1$ or $i-1$ (see [22] for details; for a more thorough introduction to the Moran process in evolutionary games, see e.g. [26]). To make the chain ergodic, we include mutation terms with $\mu > 0$ that prevent states 0 and n from being absorbing. Mutation means that with probability $\frac{1}{2}\mu$, the state switches. The expressions for these probabilities including mutation are as follows:

$$P_{i,i+1} = \frac{N-i}{N}\mu\frac{1}{2} + (1-\mu)p_{i,i+1}, \text{ and } P_{i,i-1} = \frac{i}{N}\mu\frac{1}{2} + (1-\mu)p_{i,i-1} \tag{7}$$

The terms $p_{i,j}$, i.e. the probabilities of state transitions without mutation, depend on which dynamics are used. That is, we will find different expressions for a Moran process as compared to a pairwise comparison. One caveat is that the virus inoculation game leads to expected payoffs $\hat{\pi}_X^i \leq 0$. To be able to use these terms in the equations for the Moran process probabilities, we use the standard assumption of an *exponential fitness function* (see [28]): expected payoffs are mapped to a fitness with the function $F(x) = e^x$, such that the fitness becomes $f_X^i = e^{s\hat{\pi}_X^i}$. We subsequently set the parameter $s = 1$, as is common in the literature. This quantity is now always positive, is larger for smaller costs (or equivalently, larger payoffs), and can be used in the standard Moran probabilities (cf. [18] and [22]).

Meanwhile, for the pairwise comparison – imitation dynamics, we can still use the payoffs themselves without transforming them (see [22]).

From these transition probabilities, we can calculate the stationary distribution of the process: it is the normalized left eigenvector of the transition matrix, which is the tridiagonal matrix $\mathbf{P}$, see [22].

The mutation-selection equilibrium is then the solution to

$$\mathbf{xP} = \mathbf{x}.$$

It is the limit distribution (which exists as per the ergodicity of the underlying Markov chain) that gives the long-run frequencies of all possible states. These frequencies can also be obtained by simulating the process for long enough and counting how often each state occurs.

To be able to compare the evolutionary price of anarchy with the static price of anarchy, we first need to describe the Nash equilibria of the system. For the complete graph, by using Corollary 1.1, the static analysis predicts one equivalence class of Nash equilibria, $\mathcal{N}$, where exactly $N - t = N - VN/I$ nodes are inoculated. We denote the efficiency of these equilibria as PoA, the static price of anarchy.

In order to calculate $ePoA$, we first calculate the average social cost $\hat{S}$. We do so either by averaging over the total cost in each round in our simulations, or taking $\hat{S} = \mathbf{xR}$ (cf. Eq. 4). For the complete graph, the vector $\mathbf{R}$ containing the total system cost in all possible configurations, with $i = 0, ..., N$ inoculated nodes, has the components

$$R_i = i\hat{\pi}_C^i + (N - i)\hat{\pi}_D^i = iV + (N - i)\frac{N - i}{N}I. \tag{8}$$

We also know the cost of the optimum; it is attained in the state with $i^* = \frac{N(2L-V)}{2L}$, which is the number of inoculated nodes where $R_i^* = \max R_i$ holds. The optimal cost is then $\Omega = i^*V + \frac{(N - i^*)^2}{N}I$. With this, we can use Eq. 5 to measure the efficiency of the different dynamics by finding their corresponding evolutionary price of anarchy as $\hat{S}/\Omega$. We present our analysis below.

Using our evolutionary process, our findings can now be summarized in the following lemma:

▶ **Lemma 2.** *For a fixed cost vector $ct = [V, I, 0]$, large $N \gtrsim 30$, any reasonable mutation rates $0 < \mu < 0.5$, and intermediate to large selection strength $\beta > 1$ (for the pairwise comparison process), we always recover the predicted equivalence class of Nash equilibria, $\mathcal{N}$, as the most abundant state in the selection-mutation equilibrium of both types of processes. That is, the process spends the largest fraction of time in the Nash equilibria, where exactly $t = VN/I$ nodes are inoculated.*

We note that there is also substantial weight on neighboring states of the Nash equilibria (with $t \pm i$, where $i \in \{1, 2, 3, ...\}$) with worse total welfare, due to the stochasticity of the process. However, the average social cost $\hat{S}$ is not substantially different from the cost of the Nash equilibria, since the weight on neighboring states is symmetrically distributed. The numerical results these insights are based on are provided in Figs. 3 and 4, as well as in [22] and upon request in full.

What this means for the evolutionary price of anarchy is expressed in the following corollary.

▶ **Corollary 3** (Evolutionary Price of Anarchy for Cliques.). *The evolutionary price of anarchy $ePoA_{Clique}$ in a clique with N nodes approaches the efficiency of the Nash equilibrium (the price of anarchy PoA_{Clique}), as $\mu \to 0$ and N grows large, such that $\lim_{N \to \infty} |ePoA_{Clique} - PoA_{Clique}| = 0$.*

A straightforward argument shows why Lemma 2 and Corollary 3 hold.

For the base case of the unstructured population on a complete graph, the perfect information setting corresponds to the local information setting, as each node only has distance 1 from every other node. Furthermore, the Markov chain underlying the stochastic evolutionary process is ergodic by $\mu > 0$, such that there exists a stationary distribution and the process converges to it. This stationary distribution then places most weight on the Nash equilibrium described in [2], as it is the configuration where nodes have no incentive to switch their strategy, and are aware of this fact, just as in the static case. The stochastic noise becomes smaller as the number of nodes grows larger (which inhibits the system spending too much time in the two extremal states of all nodes inoculated and all nodes insecure), and as the mutation rate becomes smaller (which sharpens the peak of the distribution at the equilibrium). This lets us recover the equilibrium results of [2].

4.1.2 Star graph

For star graphs $K_{1,N-1}$, we can also numerically compute the expected payoffs and the Markov chain properties of the process. The $2N$ states in this case are of the form (t, l), $t \in \{0, 1\}$ and $l \in \{0, \ldots, N-1\}$. The parameter t denotes the inoculation state of the center node, whereas l gives the number of inoculated leaf nodes.

By using Corollary 1.1 again, we find two equivalence classes of Nash equilibria: class $\mathcal{N}_1$ has $N - t = N - \lfloor V N/I \rfloor$ inoculated leaf nodes (which in our notation is the state $(0, N-t)$), whereas $\mathcal{N}_2$ contains the optimal equilibrium, which features the center as the only inoculated node (state $(1, 0)$). We will show that for this highly structured population, the outcome can be quite different from the predictions of the static analysis of the one-shot game. Local evolutionary processes prevent the system from spending too much time in either of the equilibria classes. We detail this in the following paragraphs.

To see this, we first compute the expected payoffs of leaf nodes ($\hat{\pi}_X^{t,l}$, with $X \in \{C, D\}$) and the center node ($\hat{\pi}_{Center}^{t,l}$) in the configurations (t, l):

$$\hat{\pi}_C^{t,l} = \hat{\pi}_{Center}^{1,l} = V, \ \hat{\pi}_D^{0,l} = \hat{\pi}_{Center}^{0,l} = \frac{N-l}{N}I, \text{ and } \hat{\pi}_D^{1,l} = \frac{1}{N}I. \tag{9}$$

For the Moran and pairwise comparison – imitation dynamics we derive probabilities $p_{n,o}^{k,m}$ that describe the transitions $(k, n) \to (m, o)$ without mutation, again with fitness $f_X^{t,l} = e^{\hat{\pi}_X^{t,l}}$. The exact expressions can be found in [22].

We can now get the transition matrix, with its entries $P_{l,l}^{0,1}$ [22], and subsequently the selection-mutation equilibrium $\mathbf{x}$ of the process with its corresponding average system cost.

Having calculated/simulated the stationary distribution, we observe the following (see [22]):

- No matter the network size or the process, the Nash equilibria in $\mathcal{N}_1, \mathcal{N}_2$ are not abundant states in the stationary distribution, with up to a factor 10^3 in weight difference to the more frequent states. We instead find a high abundance of costly non-equilibrium states $\mathcal{X} = \{(0, N - t - i)\}$ for some integers $t > i > 0$. There is also substantial weight on the beneficial configurations with low cost $(1, N - t - i)$ for the same values of i.
- The equilibrium $\mathcal{N}_1 = (0, N-t)$ is typically of far lower frequency than the non-equilibrium states . But it still plays more of a role overall than the optimum $\mathcal{O} = \mathcal{N}_2 = (1, 0)$, which is a rare state at stationarity and almost never visited in the long run.

We will now argue why the process exhibits this off-equilibrium behavior. First of all, starting from the above observations, it is straightforward to show why the optimum, that is, the state $\mathcal{O} = (1, 0)$, cannot be an abundant state in the mutation-selection equilibrium of a star graph.

▶ **Lemma 4** (The optimal Nash equilibrium is rare in the mutation-selection equilibrium.). *Consider a star graph, with fixed but arbitrary number of nodes N. For arbitrary mutation rates $\mu > 0$, arbitrary $|V| < |I|$, and any of the three evolutionary processes we consider, the optimal Nash equilibrium $\mathcal{O} = (1, 0)$ cannot be an abundant state in the mutation-selection equilibrium.*

Proof. In fact, the equilibrium is not even reached if not for mutation. To see this, consider the states $a = (0, 0)$ and $b = (1, 1)$, and suppose $\mu = 0$. While both these states only differ in one node's strategy from the Nash equilibrium, they cannot serve as starting points for a transition. State a is absorbing in this scenario, as there is no inoculated node to pass on its strategy. Meanwhile, in the state b, the one inoculated leaf node cannot change its strategy without the center node being not inoculated – the terms $p_{l,l-1}^{1,1}$ are always zero. This however leads to the opposite of the Nash equilibrium we are trying to reach. Thus, only a nonzero mutation rate can provide access to this state. At the same time, it is clear that the transitions from $\mathcal{O}$ do not need mutation to reach the neighboring states a and b, which leads to a higher probability to leave the state than to reach it. This makes $\mathcal{O}$ unsustainable in the long run.   ◀

The following lemma states the corresponding result for the other Nash equilibria.

▶ **Lemma 5** (The Nash equilibria from the equivalence class $\mathcal{N}_1$ cannot form an abundant state in the mutation-selection equilibrium.). *Consider a star graph, with fixed but arbitrary number of nodes N. For arbitrary mutation rates μ, arbitrary $|V| < |I|$, and any of the three evolutionary processes we consider, the Nash equilibria of the form $(0, N - t)$ cannot be an abundant state in the mutation-selection equilibrium. Instead, pairs of states of the form $(0, N - t - i)$ and $(1, N - t - i)$, $i \in \mathcal{I}$ with the set $\mathcal{I} \subseteq \mathbb{N}_0$ depending on N, V, I, μ and the evolutionary dynamics used, act as sinks of the process.*

We illustrate our argument by Fig. 2, where we show the case of the pairwise comparison process on a star graph with $N = 12$ and $V/I = 1/3$.

Consider the process as a two-dimensional random walk, defined by our transition probabilities $p_{n,o}^{k,m}$ (see [22]) in the limit of the mutation rate $\mu \to 0$. Let $t = VN/I$; the Nash equilibrium $\mathcal{N}_1$ is then the state $(0, N - t)$, as discussed above. In the example, we have $t = 4$, such that $\mathcal{N}_1 = (0, 8)$.

For easier readability, we use the notation $\{u_{(i,i+1)}, q_{(i,i-1)}, r_i, s_i\}$ for $\{p_{N-t-i,N-t-i-1}^{0,0}, p_{N-t-i,N-t-i+1}^{1,1}, p_{N-t-i,N-t-i}^{0,1}, p_{N-t-i,N-t-i}^{1,0}\}$ (cf. [22]). That is, u gives the probability of moving one further step away from the Nash equilibrium $\mathcal{N}_1$ by one leaf node
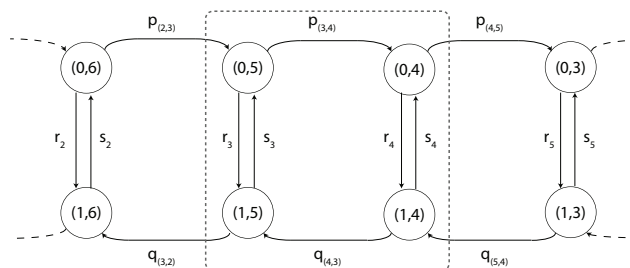
■ **Figure 2** Crucial section of the Markov chain underlying evolutionary processes with local information on the star graph $K_{1,11}$. The ratios r/u, s/q, and u/q of transition probabilities and their dependency on the distance i from the Nash equilibrium $(0,8)$ has the two-dim. random walk oscillate between pairs of states $\{(0,5),(1,5)\}$ and $\{(0,4),(1,4)\}$, and thus be trapped in increasingly constricted cycles.

switching to insecure; q the probability of moving one step closer to $\mathcal{N}_1$ by one leaf node switching to secure; r the probability of the center switching to secure, and s the probability of the center switching to insecure. The parameter i can be thought to be the distance to $\mathcal{N}_1$. Note that this two-dimensional random walk has a defined direction; there is no possibility to increase i in the lower level of the chain (where the center is inoculated), and no possibility to decrease i in the upper level (where the center is insecure).

We show in [22] that due to these transition probabilities, the random walk gets trapped in increasingly constricting cycles as it moves away from $\mathcal{N}_1$. In the example of Fig. 2, these states form the set $\{(0,5),(0,4),(1,5),(1,4)\}$, with the most weight on $(0,4)$ and $(0,5)$.

   We thus have shown that using local information only, the system spends a high fraction of time in states that are not Nash equilibria, and will not reach the optimum. What does this mean for the evolutionary price of anarchy?

   Seeing that the abundant states carry a high social cost compared to the optimum and also the worse equilibria in $\mathcal{N}_1$, it is already intuitive that $ePoA$ will be larger than PoA in the star graph, as long as the mutation rate μ is sufficiently small ($\mu \lessapprox 0.005$). Indeed, using the stationary distribution $\mathbf{x}$ to calculate the social cost $\hat{S}$ as in Eq. 4, and then the evolutionary price $ePoA$ as in Eq. 5, we find that the evolutionary process has to settle for a relatively high $\hat{S}$, and with it, a high $ePoA$ (see also Figs. 3- 4). We summarize this in the following corollary:

▶ **Corollary 6** (Evolutionary Price of Anarchy for Star Graphs.). *For small mutation rates* $\mu \lessapprox 0.005$, *arbitrary* N *and arbitrary* $|V| < |I|$, *the evolutionary price of anarchy* $ePoA_{Star}$ *in a star graph* $K_{1,N-1}$ *with* N *nodes is at least equal to or higher than the static price of anarchy* PoA_{Star}. *That is,* $ePoA_{Star} - PoA_{Star} \geq 0$.

We note that the exact evolutionary price in relation to the static price of anarchy is determined by parameter choices. This means that with a mutation rate $\mu \gtrapprox 0.005$, for some choices of cost vector and network size, it is possible to achieve a slightly lower *average* cost than can be achieved by the worst possible solution to the one-shot game due to higher mutation facilitating contributions from states that lead to a high total payoff (see Fig. 3).

   However, our results let us conjecture that for a local information model with reasonably small mutation rates, we cannot hope to do much better on average than the worst Nash equilibrium in highly structured networks (like a star), much less reach the optimum, such that paying (at least) the price of anarchy is not only a theoretical possibility, but also a realized fact.

 Simulation of more complex topologies

With the algorithms introduced above, we are able to simulate the process also for more complicated graphs. While numerical analysis is usually impossible in these cases - it is intuitive that obtaining numerically precise results on the stochastic process will be harder with increasing graph size and less inherent symmetry in the graph we can always compute the average social welfare by simulating the different dynamics long enough. We can even feasibly find the stationary distribution for graphs that have some inherent symmetry.

5.1   2-clique network

This graph consists of two cliques, connected by a single path. The states of the Markov chain are now (a, b, c, d), with $a, c \in \{0, 1\}$ denoting the inoculation state of the two nodes where the cliques are joined (subsequently called the hubs), and $b, d \in \{0, ...N - 1\}$ denoting how many of the $N - 1$ remaining nodes on each side are inoculated. For this graph, seeing that the underlying Markov chain of the process only has N^2 states, it is indeed possible to also find both the selection-mutation equilibrium and the average social cost $\hat{S}$ by simulation. To be able to calculate $ePoA$, and compare it to the efficiencies of other equilibria as well as the static price of anarchy PoA, we also compute the social welfare in all possible configurations (vector $\vec{R}$), and find the optimum as the maximum over $\vec{R}$. This section describes our findings after running our simulations with $\mu = 0.001$ for the network sizes $N = \{10, 12, 20, 50\}$ and the realized cost vectors $ct = \{[-2, -1, 0], [-3, -1, 0]\}$.

1. First considering static Nash equilibria as the baseline, we again use Corollary 1.1 with $t = V/I$ and find two equivalence classes $\mathcal{N}_1, \mathcal{N}_2$ of Nash equilibria. Class $\mathcal{N}_1$ contains the states $(0, p, 0, q)$, with $p + q = N - t$, which have both hubs insecure, and p respective q inoculated nodes in the remainders of the two cliques. The other class of equilibria, $\mathcal{N}_2$, is composed of the states $(1, N/2 - t - 1, r, s)$ and $(r, s, 1, N/2 - t - 1)$ with $r + s = N/2 - t$, which are the states where at least one hub is always inoculated. In the case where N is divisible by I, the two equilibria classes are equivalent, giving the same cost. However, when N is not divisible by I, $\mathcal{N}_2$ is more efficient with respect to the overall cost, such that it is the cost of the equilibria in $\mathcal{N}_1$ which is used to calculate the PoA.

2. In evolution, we find an $ePoA > 1$ for all three network sizes and all three evolutionary processes – the optimum is not an abundant state. It is also the case that $ePoA \geq PoA$, making the average cost slightly higher than the static price of anarchy, even though $ePoA$ becomes smaller with increasing network size. In all the tested scenarios, the system does not spend a substantial fraction of time close to social optima and does not make up for partially costly off-equilibrium behavior it exhibits otherwise. We again find that the Moran processes show behavior that is slightly dependent on the network size N, whereas the pairwise comparison process gives more consistent results even for smaller N, and also exhibits the smallest value of $ePoA$.

We have seen in Corollary 3 that for one clique on its own, we recover the Nash equilibria in our evolutionary process. However, the behavior of such a process on the 2-clique graph leads to an outcome that differs both from the optimum and the static predictions for Nash equilibria. Intuitively, this is due to the link between the two hubs acting as a bottleneck for information flow.

5.2   2-star network

Another example of a symmetric graph is two stars, joined at their hubs. Here, we find:

1. We again begin with the static Nash equilibria and the static PoA. In this case, there are three equivalence classes of equilibria: $\mathcal{N}_1$ corresponds to its counterpart in the 2-clique graph – it contains states $(0, p, 0, q)$ (defined analogously to above), with $p + q = N - t$, and is again the class with the most costly equilibria. It therefore is used for the static price of anarchy. Class $\mathcal{N}_2$ consists of the states $(1, 0, s, 0)$, with $s = 0$ if $t \geq N/2$ and $s = 1$ otherwise. Lastly, $\mathcal{N}_3$ exists if $t < N/2$, and features states of the type $(1, 0, 0, N/2 - t + 1)$. Here, we can also explicitly characterize the optimum: it is always the state $O = (1, 0, 1, 0)$.

2. An evolutionary process on the 2-star graph for the network sizes and V/I ratios tested again leads to an $ePoA \geq PoA$. Seeing that the basic component of the graph – the simple star – already exhibits off-equilibrium behavior, this is not too surprising. However, as opposed to the single star, we now observe that the Moran Birth-Death scenario is advantageous for all network sizes, as it leads to the lowest overall $ePoA$.
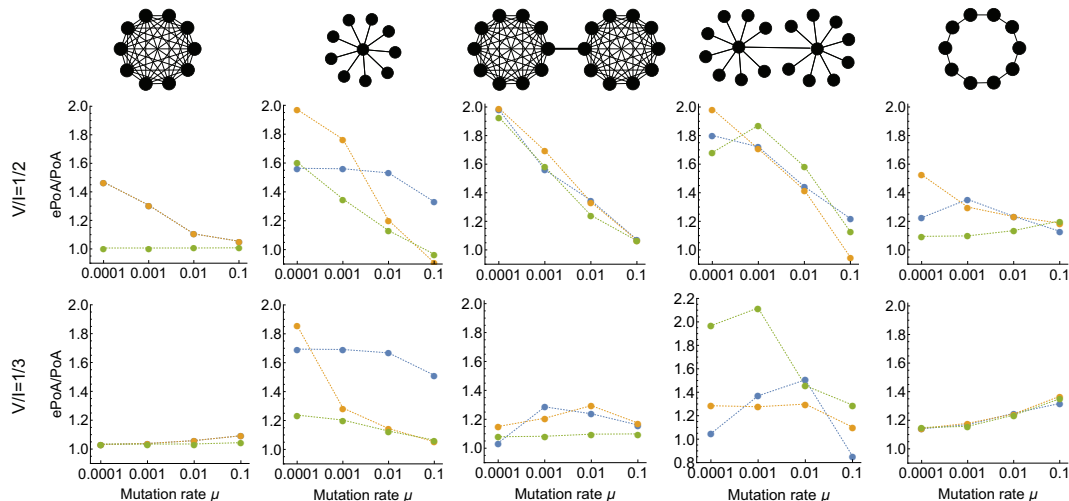

Figure 3 We visualize the ratio of the evolutionary price of anarchy to the price of anarchy, $ePoA/PoA$, for the four discussed topologies, two different values of V/I, and varying *mutation rate*. The network size is kept constant at $N = 20$. We plot the three different evolutionary dynamics together: Moran-DB (blue), Moran-BD (yellow), and pairwise comparison process (green). We see that different processes show different efficiency, depending on topology and mutation rate, and the behavior of the $ePoA$ does not have to be strictly monotonic - there can be "sweet spots" for certain combinations of parameters. Simulations were run for $5 * 10^5$ iterations.

5.3   Cycle graph

Another structure of interest is the cycle. While this topology is quite simple, the process does not show straightforward behavior due to a large state space with few possible transitions.

Due to symmetry considerations, we can encode states corresponding to configurations $\mathbf{a} \in \{0, 1\}^N$ by counting the number of insecure nodes in between every two inoculated nodes, and listing the number of gaps ordered by their size, starting at 1. As an example, we can take a network of size $N = 6$ with $i = 4$ inoculated nodes in the configuration 010111,
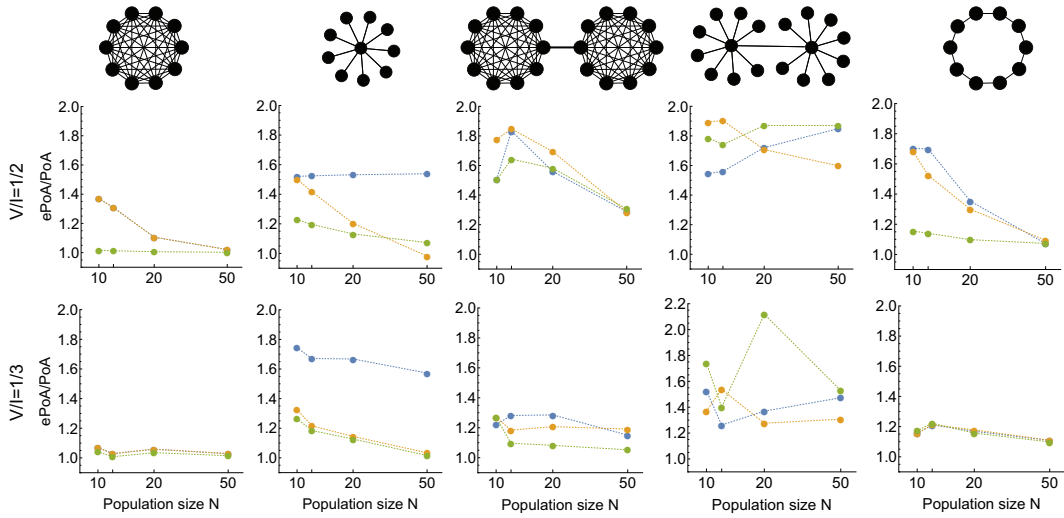
Figure 4 We visualize the ratio of the evolutionary price of anarchy to the price of anarchy, $ePoA/PoA$, for the four discussed topologies, two different values of V/I, and varying *network size*. The mutation rate is kept constant at $\mu = 0.001$. We again plot the three different evolutionary dynamics together: Moran-DB (blue), Moran-BD (yellow), and pairwise comparison process (green). We see that different processes show different efficiency, depending on the network topology and the network size. Again, the behavior of the $ePoA$ does not have to be strictly monotonic - there can be "sweet spots" for certain combinations of parameters. Simulations were run for $5 * 10^5$ iterations.

which corresponds to the state $[4, 2, 0, 0, 0, 0, 0]$, where the first number in the list is simply $i = 4$. This is thereby also equivalent to the configurations 101011, 110101 and 111010. To enumerate and order these states, we can describe the sets of states with i inoculated nodes more generally as $\{i\}$-integer partitions of $N - i$, where an $\{i\}$-partition of x is the set of partitions of x with length $\leq i$. This way, we get a natural ordering on states - from the smallest partition length to the largest-, and find a simple way to enumerate them. In the previous example, the corresponding partition of $N - i = 2$ is $(1, 1)$. The total number of states in a network with N nodes is then $2 + \sum_{i=1}^{N-1} \sum_{k=1}^{i} p_k(N - i)$, where $p_k(x)$ is the partition function, i.e. the number of integer partitions of x with length k.

For a state to be a Nash equilibrium, the largest insecure component cannot be larger than $t = VN/I$. For this condition to hold, the number of inoculated nodes has to be $i^* = \left\lceil \frac{N}{t+1} \right\rceil$, and the Nash equilibrium states is then the $\{i\}$-partition of $N - i^*$, where the largest part is at most t. In our toy example above, using $V/I = 1/2$, those are the states $[2, 1, 0, 1, 0, 0, 0]$ and $[2, 0, 2, 0, 0, 0, 0]$.

We run simulations for $N = \{10, 12, 20, 50\}$ and $\mu = \{0.0001, 0.001, 0.01, 0.1\}$, and again find the $ePoA$ to be higher than the PoA, making the cycle a suboptimal topology for efficient inoculation: clusters of inoculated/insecure nodes can not be broken if not for mutation. Intuitively speaking, a Nash equilibrium with e.g. 2 inoculated nodes that are not next to each other will not easily be reached, as moving an inoculated node from point A to B has to involve mutation: starting from a cluster of two inoculated nodes, two imitation events and two mutation events are necessary to move the second node two steps away from the first one. This becomes more unlikely as $\mu \to 0$. However, a surprising result can be seen in Fig. 3, for $V/I = 1/3$: the three evolutionary processes behave more or less the same in this case, and the efficiency is monotonic in the mutation rate, just as in the complete graph. An interesting direction for future work would thus be a more thorough analysis of the role of regular graphs in the behavior of stochastic evolutionary processes.

6 Related Work

There exists much research on alternative equilibrium concepts which overcome the shortcomings of Nash equilibria. An interesting concept in this context are Bayes-Nash equilibria [23] in games of incomplete information (e.g., [23]), and in particular the Bayes-Nash Price of Anarchy [20]. Another interesting concept, also based on dynamic equilibria, is the stochastic price of anarchy [6]: it addresses the issue that some PoAs are sensitive to small perturbations, by considering more stable concepts. However, these models come with strong assumptions on the agents' information and/or resources. In particular, not much is known today about the price of anarchy in games with partial information *based on local interactions*. The stochastic price of anarchy even requires agents to have information beyond the current game (which is derived implicitly), to decide on their best action: the players need knowledge of the execution *history*. Dynamic games based on learning dynamics [4], such as regret minimization [11] or fictitious play [12], have the same shortcoming: they require players to keep track of historical information. There is work on network creation games with local knowledge [3], however, it only considers the static pice of anarchy and does not embed the game into an evolutionary process. All these results hence do not apply to the evolutionary games studied in this paper. Indeed, to the best of our knowledge, this paper is the first to consider the application of the price of anarchy to evolutionary games where players are memoryless and only have (graph-)local information.

7 Conclusion

This paper introduced the evolutionary price of anarchy to study equilibrium behavior of simple agents interacting in a distributed system, based on local information. We showed that for memoryless agents, the resulting equilibria can be significantly different from their static equivalent, and that Nash equilibria are sometimes assumed only very infrequently.

We believe that our work can provide a novel perspective on the discussion of the impact of limited information in games. In particular, it opens several interesting avenues for future research. Our model still comes with several limitations due to the well-known notorious difficulty of analyzing evolutionary multi-player games on graphs. In particular, it will be interesting to analyze the ePoA for additional topologies and more general interaction models (beyond memoryless), as well as to explore randomized (mixed) strategies. It would also be interesting to prove or disprove our conjecture that processes based on imitation dynamics always result in a ePoA which is higher than the PoA in the virus inoculation game.

───── **References** ─────

1   Benjamin Allen and Martin A Nowak. Games on graphs. *EMS surveys in mathematical sciences*, 1(1):113–151, 2014.
2   James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. Inoculation Strategies for Victims of Viruses and the Sum-of-squares Partition Problem. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 43–52, Philadelphia, PA, USA, 2005.
3   Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Locality-based Network Creation Games. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '14, pages 277–286, New York, NY, USA, 2014. ACM.
4   Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, learning, and games.* Cambridge university press, 2006.
5   Krishnendu Chatterjee and Thomas A Henzinger. A survey of stochastic ω-regular games. *Journal of Computer and System Sciences*, 78(2):394–413, 2012.

6   Christine Chung, Katrina Ligett, Kirk Pruhs, and Aaron Roth. The price of stochastic anarchy. In *International Symposium on Algorithmic Game Theory*, pages 303–314. Springer, 2008.

7   Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H Papadimitriou, and Scott Shenker. On a network creation game. In *Proc. 22nd annual symposium on Principles of distributed computing (PODC)*, pages 347–351. ACM, 2003.

8   Chris Fleizach, Michael Liljenstam, Per Johansson, Geoffrey M Voelker, and Andras Mehes. Can you infect me now?: malware propagation in mobile phone networks. In *Proceedings of the 2007 ACM workshop on Recurring malcode*, pages 61–68. ACM, 2007.

9   Drew Fudenberg and Lorens A. Imhof. Monotone imitation dynamics in large populations. *Journal of Economic Theory*, 140(1):229–245, 2008.

10   Drew Fudenberg, Martin A Nowak, Christine Taylor, and Lorens A Imhof. Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theoretical population biology*, 70(3):352–363, 2006.

11   Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

12   Josef Hofbauer and William H Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002.

13   Josef Hofbauer and Karl Sigmund. *Evolutionary games and population dynamics*. Cambridge university press, 1998.

14   Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science (STOC)*, pages 404–413. Springer, 1999.

15   Erez Lieberman, Christoph Hauert, and Martin A Nowak. Evolutionary dynamics on graphs. *Nature*, 433(7023):312, 2005.

16   Thomas Locher, Patrick Moor, Stefan Schmid, and Roger Wattenhofer. Free Riding in BitTorrent is Cheap. In *Proc. 5th Workshop on Hot Topics in Networks (HotNets)*, pages 85–90, 2006.

17   Swarup Mohalik and Igor Walukiewicz. Distributed Games. In *FST TCS 2003: Foundations of Software Technology and Theoretical Computer Science*, pages 338–351. Springer, 2003.

18   Martin A. Nowak. *Evolutionary Dynamics*, pages 123–143. Harvard University Press, 2006.

19   Amir Pneuli and Roni Rosner. Distributed reactive systems are hard to synthesize. In *Proc. 31st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 746–757, 1990.

20   Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Transactions on Economics and Computation (TEAC)*, 3(1):6, 2015.

21   Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

22   Laura Schmid, Krishnendu Chatterjee, and Stefan Schmid. The Evolutionary Price of Anarchy: Locally Bounded Agents in a Dynamic Virus Game. *arXiv preprint arXiv:1906.00110*, 2019.

23   Satinder Singh, Vishal Soni, and Michael Wellman. Computing approximate Bayes-Nash equilibria in tree-games of incomplete information. In *Proc. 5th ACM conference on Electronic Commerce (EC)*, pages 81–90. ACM, 2004.

24   J Maynard Smith and George R Price. The logic of animal conflict. *Nature*, 246(5427):15, 1973.

25   John Maynard Smith. Evolution and the Theory of Games. In *Did Darwin Get It Right?*, pages 202–215. Springer, 1988.

26   Christine Taylor, Drew Fudenberg, Akira Sasaki, and Martin A Nowak. Evolutionary game dynamics in finite populations. *Bulletin of mathematical biology*, 66(6):1621–1644, 2004.

27   Arne Traulsen, Martin A. Nowak, and Jorge M. Pacheco. Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:011909, July 2006.

28   Arne Traulsen, Noam Shoresh, and Martin A Nowak. Analytical results for individual and group selection of any intensity. *Bulletin of mathematical biology*, 70(5):1410, 2008.

29   Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.