

Invasion Dynamics in the Biased Voter Process

Loke Durocher¹, Panagiotis Karras¹, Andreas Pavlogiannis¹ and Josef Tkadlec²

¹Aarhus University, Aabogade 34, Aarhus, Denmark

²Harvard University, 1 Oxford Street, Cambridge, USA

{panos,pavlogiannis}@cs.au.dk, tkadlec@math.harvard.edu

Abstract

The *voter process* is a classic stochastic process that models the invasion of a mutant trait A (e.g., a new opinion, belief, legend, genetic mutation, magnetic spin) in a population of agents (e.g., people, genes, particles) who share a resident trait B , spread over the nodes of a graph. An agent may adopt the trait of one of its neighbors at any time, while the *invasion bias* $r \in (0, \infty)$ quantifies the stochastic preference towards ($r > 1$) or against ($r < 1$) adopting A over B . Success is measured in terms of the *fixation probability*, i.e., the probability that eventually all agents have adopted the mutant trait A . In this paper we study the problem of *fixation probability maximization* under this model: given a budget k , find a set of k agents to initiate the invasion that maximizes the fixation probability. We show that the problem is NP-hard for both $r > 1$ and $r < 1$, while the latter case is also inapproximable within any multiplicative factor. On the positive side, we show that when $r > 1$, the optimization function is submodular and thus can be greedily approximated within a factor $1 - 1/e$. An experimental evaluation of some proposed heuristics corroborates our results.

1 Introduction

Invasion processes are a standard framework for modeling the spread of novel traits in populations. The population structure is represented as a graph, with nodes occupied by agents, while the edge relation captures some sort of agent connection, such as spatial proximity, friendship, direct communication, or particle interaction. The invasion dynamics define the rules of trait diffusion from an agent to its neighbors. These dynamics raise a natural optimization problem that calls to choose a set of at most k initial trait-holders that maximizes a goal involving the eventual spread of the trait in the population. This type of problem has been studied extensively with a focus on influence spread on social network dynamics [Domingos and Richardson, 2001; Kempe *et al.*, 2003; Mossel and Roch, 2007]; see [Li *et al.*, 2018] for a survey.

The *voter process* is a fundamental and natural probabilis-

tic diffusion model that arises in evolutionary population dynamics, where the diffused trait is a mutation and the initial trait-holders are mutants. Also known as imitation updating in evolutionary game theory [Ohtsuki *et al.*, 2006], it was initially introduced to study particle interactions [Liggett, 1985] and territorial conflict [Clifford and Sudbury, 1973]; since then, it has found numerous applications in the spread of genetic mutations (the death-birth Moran process) [Hindersin and Traulsen, 2015; Tkadlec *et al.*, 2020], social dynamics [Castellano *et al.*, 2009], robot coordination and swarm intelligence [Talamali *et al.*, 2021]. In the long run, the process leads to a *consensus state* [Even-Dar and Shapira, 2007], in which the mutant trait is either fixated (i.e., adopted by all agents), or gone extinct. The calibre of the invasion is quantified in terms of its *fixation probability* [Kaveh *et al.*, 2015; Brendborg *et al.*, 2022]; thus, the goal of optimization in this context is to maximize the probability that the mutant trait spreads in the whole population. Unlike other models of spread, the voter process accounts for *active* resistance to the invasion: an agent carrying the mutant trait may not only forward it to its neighbors, but also lose it due to exposure to the resident trait from its neighbors. This model exemplifies settings such as individuals that may switch opinions, magnetic spins that may undergo transient states, and sub-species that compete in a gene pool.

In this paper, we study the problem of *fixation probability maximization* in the context of the *biased voter process*: find a set of k agents initially having a *mutant trait* A that maximizes the fixation probability of A in a population with a pre-existing *resident trait* B . The *invasion bias* $r \in (0, \infty)$ quantifies the stochastic preference of each agent towards ($r > 1$) or against ($r < 1$) adopting A over B . Despite extensive interest in this type of optimization problem and the appeal and generality of the voter process, to our knowledge, this problem in the context of the voter process has only been studied in the restricted case $r = 1$, which admits a simple polynomial-time algorithm [Even-Dar and Shapira, 2007]. Yet biased invasions naturally occur in most settings: for genetic mutations, r captures the fitness (dis)advantage conferred by the new mutation [Allen *et al.*, 2020]; for magnetic spins, r reveals the presence of a magnetic field [Antal *et al.*, 2006]; for social traits, r accounts for the societal predisposition towards a new idea [Bhat and Redner, 2019].

Our contributions can be summarized as follows:

1. We show that computing the fixation probability on undirected graphs admits a fully polynomial-time approximation scheme (FPRAS).
2. We show that, notwithstanding the neutral case where $r = 1$, fixation maximization is NP-hard both for $r > 1$ and for $r < 1$, while the latter case is also inapproximable within any multiplicative factor.
3. On the positive side, we show that, when $r > 1$, the optimization function is submodular and thus can be greedily approximated within a factor $1 - 1/e$.

2 Preliminaries

Structured populations. Following the literature standard, we represent a population structure by a (generally, directed and weighted) graph $G = (V, E, w)$ of $n = |V|$ nodes. Each node stands for a location in some space, while a population of n agents is spread on G with one agent per node. Given a node u , we denote by $N(u) = \{v \in V : (v, u) \in E\}$ the neighbors of u , and by $\deg(u) = |N(u)|$ its degree. The weight function $w: E \rightarrow \mathbb{R}_{\geq 0}$ maps every edge to a non-negative number. As a convention, we write $w(v, u) = 0$ if $(v, u) \notin E$. We consider graphs for which the sub-graph induced by the support of w is strongly connected — which is necessary for the voter process to be well-defined. We say that G is *undirected* if E is symmetric and irreflexive, and $w(v, u) = 1$ for all $(v, u) \in E$. For simplicity, when G is undirected we denote it by $G = (V, E)$.

The voter invasion process. The classic voter process models the invasion of a mutant trait A into a homogeneous population [Clifford and Sudbury, 1973; Liggett, 1985]. The process distinguishes between two types of agents, *residents*, who carry a resident trait B , and *mutants*, who carry the mutant trait A . A *seed set* $S \subseteq V$ defines the agents that are initially mutants. The *invasion bias* $r > 0$ is a real-valued parameter that defines the relative competitive advantage of mutants in this invasion. When $r > 1$ (resp., $r < 1$), the mutants have a competitive advantage (resp., competitive disadvantage), and the process is biased towards (resp., against) invasion, while $r = 1$ yields the neutral case. A *configuration* $X \subseteq V$ defines the set of mutants at a given time. The *fitness* of the agent occupying a node $u \in V$ is defined as

$$f_X(u) = \begin{cases} r, & \text{if } u \in X \\ 1, & \text{otherwise.} \end{cases}$$

The discrete-time *voter invasion process* [Antal et al., 2006; Even-Dar and Shapira, 2007], a.k.a. *Moran death-birth process* [Tkadlec et al., 2020; Allen et al., 2020], with seed set S is a stochastic process $(\mathcal{X}_i)_{i \geq 0}$, where $\mathcal{X}_i \subseteq V$ is a random configuration. Initially, $\mathcal{X}_0 = S$, and, given the current configuration X , we obtain \mathcal{X}_{i+1} in two stochastic steps.

1. (*Death event*): a focal agent on some node u is chosen with uniform probability $1/n$.
2. (*Birth event*): the agent on u adopts the trait of one of its

neighbors, on v , with probability:

$$\frac{f_X(v) \cdot w(v, u)}{\sum_{x \in N(u)} f_X(x) \cdot w(x, u)}.$$

Note that the focal agent may assume its pre-existing trait, and generally, the mutant set may grow or shrink in each step. We often identify agents with the nodes they occupy and use expressions such as “node v reproduces onto u ” to indicate that the agent occupying node u adopts the trait of the agent at node v . See Fig. 1 for an illustration.




Figure 1: Initially, each node is occupied by a resident (red), except for a seed set S of nodes occupied by mutants (blue). In each step of the voter process with bias r , one random node adopts the type of one of its neighbors, with mutants having a relative (dis)advantage r .

Fixation probability. As G is strongly connected, one type will spread to all nodes in finitely many steps with probability 1. The mutant invasion succeeds if the mutant trait spreads, an event called *fixation*. Given a graph G , a bias r , and an initial seed set S , the *fixation probability*, denoted as $\text{fp}_r^G(S)$, is the probability that the mutants starting from S eventually fixate on G . See Fig. 2 for an illustration.




Figure 2: A graph $G = (V, E)$ of 7 nodes and fixation probability $\text{fp}_r^G(S)$ for 5 different subsets S of 2 nodes vs. bias $r \in [0.01, 100]$. We have $\text{fp}_{r \rightarrow 0}^G(S) \rightarrow 0$, unless S forms a vertex cover (Lemma 6) and $\text{fp}_{r \rightarrow \infty}^G(S) \rightarrow 1$ iff S contains two adjacent nodes.

Whereas the computational complexity of determining $\text{fp}_r^G(S)$ is unknown, the function is approximately computable by a simulation of the invasion process. In the next section we will show that when G is undirected, such a simulation yields an efficient approximation scheme.

Fixation maximization. We formulate and study the optimization problem of selecting a seed set that leads to a most successful invasion, i.e., maximizes the fixation probability, subject to cardinality constraints. Formally, given a graph G , a bias r and a budget k , the task is to compute the set

$$S^* = \arg \max_{S \subseteq V, |S|=k} \text{fp}_r^G(S) \quad (1)$$

The corresponding decision question is: Given a budget k and a threshold $\alpha \in [0, 1]$ determine whether there exists a seed set S with $|S| = k$ such that $\text{fp}_r^G(S) \geq \alpha$. As Fig. 2 illustrates, different seed sets yield considerably different fixation probabilities, while the optimal choice varies depending on r , hence the intricacy of the problem.

3 Computing the Fixation Probability

For $r = 1$, the fixation probability is additive [Broom *et al.*, 2010]. Thus, given a graph $G = (V, E, w)$ and a budget k , the set S^* that maximizes $\text{fp}_{r=1}^G(S)$ is obtained by selecting the k largest values from the vector $\langle \text{fp}_{r=1}^G(\{v\}) \mid v \in V \rangle$, determined by solving a system of n linear equations [Allen *et al.*, 2015]; thus the problem can be solved in polynomial time. When G is undirected, that system admits a closed-form solution [Broom *et al.*, 2010]: $\text{fp}_{r=1}^G(\{v\}) = \frac{\deg(v)}{2|E|}$.

For $r \neq 1$, the complexity of determining $\text{fp}_r^G(S)$ is open. Here we show that, for undirected graphs, it can be efficiently approximated by a fully-polynomial randomized approximation scheme (FPRAS). The key component of our result is a polynomial upper bound on the expected number of steps $T_r^G(S)$ until one of the types fixates; an analogous approach has been used for the so-called Moran Birth-death process [Díaz *et al.*, 2014].

Theorem 1. *Let $G = (V, E)$ be an undirected graph with $|V| = n$ nodes, $|E| = m$ edges, and maximum degree Δ . Let $S \subseteq V$ be the initial seed set and $r \neq 1$. Then $T_r^G(S) = \mathcal{O}(\Delta nm) = \mathcal{O}(n^4)$.*

Proof. Without loss of generality, suppose $r > 1$ (otherwise we swap the role of mutants and residents). Given a current configuration X , we define a potential $\phi(X): 2^V \rightarrow \mathbb{R}$ as $\phi(X) = \sum_{v \in X} \deg(v)$. We show that, as we run the process, $\phi(X)$ increases in a controlled manner. To that end, we distinguish two types of configurations: we say that a configuration is *bad* if every edge $uv \in E$ connects nodes of different types, and *good* otherwise. Note that bad configurations exist iff G is bipartite, in which case there are exactly two of them. In [Durocher *et al.*, 2022], we show that:

1. When at a bad configuration, the next configuration is good and the potential does not change in expectation.
2. When at a good configuration, in one step the potential increases by at least $s = (r - 1)/(r\Delta n)$ in expectation.

All in all, this implies that, in expectation, the potential increases by at least s over any two consecutive steps (until fixation occurs). Since at all times the potential is non-negative and upper-bounded by $U = \sum_{v \in V} \deg(v) = 2m$, by a standard drift theorem for submartingales [He and Yao, 2001;

Lengler, 2020] the expected number of steps is at most:

$$T_r^G(S) \leq 2 \frac{U}{s} = \frac{4r}{r-1} \Delta nm. \quad \square$$

Theorem 1 yields the following corollary based on the Monte Carlo method.

Corollary 1. *When $G = (V, E)$ is undirected, $S \subseteq V$, the function $\text{fp}_r^G(S)$ admits a FPRAS for any r .*

4 Invasion with Favorable Bias

In this section we study the voter process for $r > 1$. We show that the problem is **NP**-hard (Section 4.1) but the fixation probability is monotone and submodular and can thus be efficiently approximated (Section 4.2).

4.1 Hardness of Optimization

Here we consider the weakly biased process, i.e., the case $r = 1 + \varepsilon$ for small $\varepsilon > 0$. Since $\text{fp}_r^G(S)$ is a continuous and smooth function of r , its Taylor expansion around $r = 1$ is:

$$\text{fp}_{r=1+\varepsilon}^G(S) = \text{fp}_{r=1}^G(S) + \varepsilon \cdot c(G, S) + \mathcal{O}(\varepsilon^2), \quad (2)$$

where $c(G, S)$ is a constant independent of ε . Recall that $\text{fp}_{r=1}^G(S) = \sum_{v \in S} \frac{\deg(v)}{2|E|}$ [Broom *et al.*, 2010]. For d -regular graphs, this yields $\text{fp}_{r=1}^G(S) = \frac{|S|}{2|E|}$, hence the first term in Eq. (2) is constant across all equal-sized seed sets S . We thus have the following lemma.

Lemma 1. *For any undirected regular graph G , there exists a small constant $\varepsilon > 0$ such that:*

$$\arg \max_{S \subseteq V, |S|=k} \text{fp}_{r=1+\varepsilon}^G(S) \subseteq \arg \max_{S \subseteq V, |S|=k} c(G, S).$$

In light of Lemma 1, we establish the hardness of fixation maximization by computing $c(G, S)$ analytically on regular graphs, and showing that maximizing it is **NP**-hard. McAvoy and Allen [2021] recently showed that, given G and S , the constant $c(G, S)$ can be defined through a system of n^2 linear equations. Here we determine $c(G, S)$ explicitly for regular graphs. For each node $v \in V$, let $x_v = 1$ iff $v \in S$. First, we restate [McAvoy and Allen, 2021, Proposition 2] in the special case of d -regular graphs and adapt it to our notation.

Lemma 2. *Fix an undirected d -regular graph $G = (V, E)$ on n nodes and an initial configuration $S \subseteq V$. Then*

$$c(G, S) = \frac{1}{2d^2n^2} \sum_{u, v \in V} |N(u) \cap N(v)| \cdot \eta_{uv}, \quad (3)$$

where η_{uv} comprise the unique solution to the linear system

$$\eta_{uv} = \begin{cases} \frac{n}{2} |x_u - x_v| + \frac{1}{2d} \sum_{w \in N(u)} \eta_{uw} \\ \quad + \frac{1}{2d} \sum_{w \in N(v)} \eta_{vw} & \text{if } u \neq v, \\ 0 & \text{if } u = v. \end{cases} \quad (4)$$

Intuitively, each η_{uv} expresses the expected number of steps in which nodes u and v are occupied by a mutant and resident, respectively. We say that an edge $e \in E$ is *active* if its endpoints are occupied by heterotypic agents. The following lemma establishes that, on regular graphs, $c(G, S)$ is decreasing in the number of active edges induced by S .

Lemma 3. Let $G = (V, E)$ be a d -regular graph on n nodes and $S \subseteq V$ an initial configuration of size $|S| = k$. Let $a(S)$ be the number of active edges induced by S . Then

$$c(G, S) = \frac{k(n-k)}{2n} - \frac{a(S)}{2dn}.$$

Proof. Let $|uv|$ be the shortest-path distance between u and v in G . We proceed in two steps. First, we claim that summing all the n^2 equations in the system of Eq. (4), we obtain

$$\sum_{u,v \in V} \eta_{uv} = \frac{n}{2} \cdot 2k(n-k) + \sum_{u,v \in V} \eta_{uv} - \frac{1}{2d} \sum_{|uv|=1} 2\eta_{uv}. \quad (5)$$

Indeed:

1. We have $\sum_{u,v \in V} |x_u - x_v| = 2k(n-k)$, since each ordered pair of heterotypic agents is counted once.
2. For a fixed pair (u, v) with $|uv| \geq 2$, term $\frac{1}{2d}\eta_{uv}$ appears on the right-hand side of $2d$ equations, namely d equations with η_{uv} on the left-hand side and $u \in N(w)$ and d equations with η_{uv} on the left-hand side and $v \in N(w)$.
3. For a fixed pair (u, v) with $|uv| = 1$, term $\frac{1}{2d}\eta_{uv}$ appears on the right-hand side of $2d - 2$ equations, namely the above $2d$ equations, except for 2 cases where $w \in \{u, v\}$.

Rearranging Eq. (5), we obtain

$$\sum_{|uv|=1} \eta_{uv} = dnk(n-k). \quad (6)$$

Second, we claim that summing the $2|E|$ equations in the system of Eq. (4) where $|uv| = 1$, we obtain

$$\sum_{|uv|=1} \eta_{uv} = \frac{n}{2} \cdot 2a(S) + \frac{1}{2d} \sum_{u,v \in V} 2|N(u) \cap N(v)| \cdot \eta_{uv},$$

ergo, using Eq. (6)

$$\sum_{u,v \in V} |N(u) \cap N(v)| \cdot \eta_{uv} = d^2nk(n-k) - dn \cdot a(S). \quad (7)$$

Indeed,

1. We have $\sum_{|uv|=1} |x_u - x_v| = 2a(S)$, since each active edge is counted exactly twice, once in each direction.
2. For a fixed pair (u, v) with $|uv| = 1$, let $N(u, v) = N(u) \cap N(v)$ be the set of common neighbors of u and v . The term $\frac{1}{2d}\eta_{uv}$ then appears on the right-hand side of precisely $2|N(u, v)|$ equations, namely $|N(u, v)|$ equations with η_{uv} on the left-hand side and $|N(u, v)|$ equations with η_{uv} on the left-hand side, where $w \in N(u, v)$.

From Lemma 2 and Eq. (7), we obtain the desired

$$\begin{aligned} c(G, S) &= \frac{1}{2d^2n^2} \sum_{u,v \in V} |N(u) \cap N(v)| \cdot \eta_{uv} \\ &= \frac{d^2nk(n-k) - a(S)dn}{2d^2n^2} = \frac{k(n-k)}{2n} - \frac{a(S)}{2dn}. \end{aligned} \quad \square$$

Thus, given a budget k , we maximize the fixation probability at the limit $r \rightarrow 1$, by minimizing the number of active edges. We are now ready to establish the main result in this section.

Theorem 2. Fixation maximization in the biased voter process with $r > 1$ is **NP-hard**, even on regular undirected graphs.

Proof. Lemma 1 and Lemma 3 imply that, for any regular undirected graph G , there exists a small $\varepsilon > 0$ for which

$$\arg \max_{S \subseteq V, |S|=k} \text{fp}_{r=1+\varepsilon}^G(S) \subseteq \arg \min_{S \subseteq V, |S|=k} a(S)$$

i.e., the fixation probability is maximized by a seed set that minimizes the number of active edges. For $k = n/2$, this is known as the minimum bisection problem, which is **NP-hard** even on 3-regular graphs [Bui *et al.*, 1987]. \square

4.2 Monotonicity and Submodularity

A real-valued set function f is called monotone if for any two sets A and B with $A \subseteq B$, we have $f(A) \leq f(B)$. Further, f is called submodular if for all sets A and B we have

$$f(A) + f(B) \geq f(A \cup B) + f(A \cap B). \quad (8)$$

Monotone submodular functions, even if hard to maximize, admit efficient approximations. Here we complement the hardness of Theorem 2 by showing that $\text{fp}_{r \geq 1}^G(S)$ is monotone and submodular. Although submodularity is well-known for some *progressive* invasion processes [Kempe *et al.*, 2003], recall that our setting is *non-progressive*, i.e., the mutant set may grow or shrink in any step. This non-progressiveness renders submodularity a non-trivial property.

Node duplication. Towards establishing monotonicity and submodularity, it is convenient to view the process in a slightly different but equivalent way, via *node duplication*. Consider the voter process \mathcal{M} on a graph G with bias $r > 1$. The node duplication view of \mathcal{M} is another stochastic process $\mathcal{M}' = (\mathcal{X}'_i)_i$ on G , defined as follows. Let $X \subseteq V$ be the current configuration.

1. A node u is chosen for death uniformly at random (i.e., this step is identical to \mathcal{M}).
2. We modify $G = (V, E, w)$ to $G' = (V', E', w')$ as follows. For every node $v \in N(u) \cap X$, we create a duplicate v' of v , and insert an edge (v', u) in G' with weight $w'(v', u) = w(v, u)$. We associate v' with fitness $f'_X(v') = r - 1$, while every other node $v \in V$ gets fitness $f'_X(v) = 1$. Finally, we execute a stochastic birth step as in the standard voter process, i.e., we choose a node v to propagate to u with probability

$$\frac{f'_X(v) \cdot w'(v, u)}{\sum_x f'_X(x) \cdot w'(x, u)}$$

The new configuration \mathcal{X}'_{i+1} contains u if either some mutant node v or its duplicate v' was chosen for reproduction. It is straightforward to see that this modified birth step preserves the probability distribution of random configurations, and thus \mathcal{M} and \mathcal{M}' are equivalent.

Monotonicity. Using the equivalence between the voter process and its variant with node duplication, we now establish the monotonicity of the fixation probability.

Lemma 4. *For any graph G and bias r , for any two seed sets S_1 and S_2 with $S_1 \subseteq S_2$, we have $\text{fp}_r^G(S_1) \leq \text{fp}_r^G(S_2)$.*

Proof. First assume $r > 1$. Consider two processes $\mathcal{M}_1 = (\mathcal{X}_i^1)_i$ and $\mathcal{M}_2 = (\mathcal{X}_i^2)_i$ starting from $\mathcal{X}_0^1 = S_1$ and $\mathcal{X}_0^2 = S_2$, respectively. We establish a coupling between \mathcal{M}_1 and \mathcal{M}_2 that satisfies $\mathcal{X}_i^1 \subseteq \mathcal{X}_i^2$ for all $i \geq 0$, from which the lemma follows.

In each step, we first choose the same node u for death in \mathcal{M}_2 and \mathcal{M}_1 with uniform probability. Then, we execute a birth step in \mathcal{M}_2 . If the reproducing node v is present in \mathcal{M}_1 , we perform the same update in \mathcal{M}_1 . Otherwise, if v is a mutant duplicate absent from \mathcal{M}_1 , we perform an independent birth step in \mathcal{M}_1 . This coupling maintains the invariant $\mathcal{X}_i^1 \subseteq \mathcal{X}_i^2$, as whenever \mathcal{M}_1 takes an independent step, a mutant has reproduced in \mathcal{M}_2 . Moreover, the coupling is indeed transparent to \mathcal{M}_1 and \mathcal{M}_2 , as every birth event occurs with probability proportional to its agent’s fitness.

The monotonicity for all r follows by symmetry, as when $r < 1$ we can exchange the roles of mutants and residents. Using the above argument, it follows that the fixation probability of the residents grows monotonically in $V \setminus S$ and, in reverse, decreases monotonically as $V \setminus S$ shrinks, hence the fixation probability of the mutants grows monotonically in S . \square

Submodularity. We now show submodularity, i.e., we have

$$\text{fp}_r^G(A) + \text{fp}_r^G(B) \geq \text{fp}_r^G(A \cup B) + \text{fp}_r^G(A \cap B). \quad (9)$$

To this end, it is convenient to consider the following more refined view of the invasion process. Given an initial seed set $S = A \cup B$, we keep track not only of the mutant set, but also the subset of mutants that are copies of those agents initially in A , and those initially in B . Consider this refined view, and execute the invasion beyond the point where mutants fixate. With probability 1, every execution that results in the fixation of S eventually leads to the fixation of A or B (possibly both, assuming that $A \cap B \neq \emptyset$). We can thus compute the fixation probability of S by summing over each execution that ends in the fixation of at least one of A, B . This “extended” process is instrumental in proving submodularity.

Lemma 5. *For any graph G and bias $r \geq 1$, the fixation probability $\text{fp}_r^G(S)$ is submodular.*

Proof. Consider any two sets $A, B \subseteq V$. Let $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$ and \mathcal{M}_4 be the four (extended) invasion processes, under node duplication, that correspond to initial seeds $A, B, A \cup B$ and $A \cap B$, respectively. For $j \in [4]$, let \mathcal{X}_i^j be the i -th random configuration of \mathcal{M}_j . Moreover, let \mathcal{Y}_i (resp., \mathcal{Z}_i) be the set of nodes of \mathcal{M}_3 in step i that are occupied by mutant agents who are descendants of some agent initially in A (resp., B). We establish a four-way coupling between all \mathcal{M}_j that preserves the following invariants:

$$(i) \mathcal{Y}_i \subseteq \mathcal{X}_i^1 \quad (ii) \mathcal{Z}_i \subseteq \mathcal{X}_i^2 \quad \text{and} \quad (iii) \mathcal{X}_i^4 \subseteq \mathcal{X}_i^1 \cap \mathcal{X}_i^2.$$

Given our extended invasion process, any execution that leads to the fixation of $A \cup B$ in \mathcal{M}_3 will eventually witness the fixation of at least one of A or B . Invariants (i) and (ii), then, guarantee that any fixating execution in \mathcal{M}_3 is also fixating in at least one of \mathcal{M}_1 and \mathcal{M}_2 . Invariant (iii) guarantees that any fixating run execution in \mathcal{M}_4 is fixating in both \mathcal{M}_1 and \mathcal{M}_2 . Hence, the three invariants imply Eq. (9), i.e., submodularity.

We now establish the coupling. In each step i , we choose the same node u in all \mathcal{M}_j to die uniformly at random. Then, as long as $\mathcal{Y}_i, \mathcal{Z}_i \neq \emptyset$ we perform the following steps.

1. Choose a node v to reproduce in \mathcal{M}_3 using node duplication (hence v is either an original or a duplicate node). For each process \mathcal{M}_j in which v is a node (either original or duplicate), propagate the agent from v to u . Note that this step updates \mathcal{M}_3 and at least one of \mathcal{M}_1 and \mathcal{M}_2 , while if it updates \mathcal{M}_4 , it also updates both \mathcal{M}_1 and \mathcal{M}_2 .
2. If v is not a node in one of \mathcal{M}_1 or \mathcal{M}_2 , then perform an independent birth step in that process, choosing some node x (either original or duplicate). If x is present in \mathcal{M}_4 , apply the birth event in \mathcal{M}_4 as well, otherwise apply an independent birth step in \mathcal{M}_4 .

If $\mathcal{Y}_i = \emptyset$ or $\mathcal{Z}_i = \emptyset$, we perform step 2 only in the process not yet terminated. As in Lemma 4, the coupling is indeed transparent to each \mathcal{M}_j and maintains invariants (i)–(iii). \square

Conclusively, monotonicity and submodularity lead to the following approximation guarantee [Nemhauser *et al.*, 1978].

Theorem 3. *Given a graph G and integer k , let S^* be the seed set that maximizes $\text{fp}_r^G(S)$, and S' the seed set constructed by a greedy algorithm opting for maximal returns in each step. Then $\text{fp}_r^G(S') \geq (1 - 1/e) \text{fp}_r^G(S^*)$.*

5 Invasion with Infavorable Bias

We now turn our attention to disadvantageous invasions, i.e., $r < 1$, and show that the problem is intractable. To establish our result we focus on the limit $r \rightarrow 0$, i.e., the process is “infinitely” biased against invasion, and write $\text{fp}_0^G(S) = \lim_{r \rightarrow 0} \text{fp}_r^G(S)$. Our key insight is the following lemma.

Lemma 6. *Let $G = (V, E)$ be an undirected graph on n nodes and $S \subseteq V$ a seed set. If S is a vertex cover on G then $\text{fp}_0^G(S) > 1/n^n$, otherwise $\text{fp}_0^G(S) = 0$.*

Proof. First, suppose S is a vertex cover. Denote the nodes in $V \setminus S$ by u_1, \dots, u_i . With probability $1/n^i \geq 1/n^n$, in the next i steps the agents at nodes u_1, \dots, u_i are selected for death (in this order). Since S is a vertex cover, each of them adopts the mutant trait, and mutants fixate.

Second, suppose S is not a vertex cover, thus we have an edge $(u, v) \in E$ with $u, v \notin S$. The idea is that, in the limit $r \rightarrow 0$, the agents on u and v are overwhelmingly more likely to spread to the rest of the graph than to ever become mutants.

Consider a node with a mutant and $b \geq 1$ resident agents among its neighbors. Note that if an agent at such a node dies, the node gets occupied by a resident agent with probability $\frac{b}{b+a} \geq \frac{1}{n} =: p_{\text{res}}$ and by a mutant agent with prob-




Figure 3: Experimental comparison of heuristics on four graphs.

ability $\frac{ar}{b+ar} \leq nr =: p_{\text{mut}}$. Since G is connected, there exists an ordering (u_1, \dots, u_k) of the nodes in S such that node u_i has a neighbor among $\{u_1, \dots, u_{i-1}\} \cup (V \setminus S)$, for each $i = 1, \dots, k$. Consider the first k steps. With probability $p_{\text{ext}} \geq \left(\frac{1}{n} \cdot p_{\text{res}}\right)^k = 1/n^{2k}$, agents at nodes u_1, \dots, u_k die (in this order) and all become residents, thus mutant extinction occurs. On the other hand, by Union Bound, within this time-frame, one of u, v becomes mutant with probability $p_{uv} \leq \frac{2}{n} p_{\text{mut}} \leq 2r$. If neither event occurs, in the next k steps the situation repeats. Thus, mutants fixate with probability at most $\frac{p_{uv}}{p_{uv} + p_{\text{ext}}} \leq \frac{2r}{2r + 1/n^{2k}} \rightarrow_{r \rightarrow 0} 0$. \square

Since vertex cover is known to be hard even on regular graphs, Lemma 6 implies the following theorem.

Theorem 4. *Fixation maximization in the biased voter process with $r < 1$ is NP-hard, even on regular undirected graphs.*

Submodularity considerations. Lemma 6 implies that $\max_{S \subseteq V, |S|=k} \text{fp}_0^G(S)$ cannot be approximated within any multiplicative factor in polynomial time, unless $\mathbf{P} = \mathbf{NP}$, since such an algorithm would decide whether S is a vertex cover on G . That is so even though Lemma 5 implies, by symmetrically exchanging mutant and resident roles, that the *extinction probability* $1 - \text{fp}_r^G(S)$ for $r < 1$ is submodular. Fixation maximization with $r < 1$ is thus equivalent to minimizing a submodular function subject to the constraint $|S| = k$. Although the minimization of submodular functions without constraints is polynomially solvable, simple cardinality constraints render it intractable [Svitkina and Fleischer, 2011].

6 Experiments

Here we present some selective case studies on the performance of various heuristics solving the fixation maximization problem for the biased voter process. Our studies do not aim to be exhaustive, but rather offer some insights that may drive future work. We consider six heuristics for the choice of the seed set, namely (i) high degree, (ii) high temperature, defined as $\mathcal{T}(u) = \sum_v w(v, u)$, (iii) high pagerank, (iv) high eigenvalue centrality, (v) high betweenness centrality, (vi) high closeness centrality, and the Greedy algorithm (Theorem 3)

that provides an approximation guarantee for $r \geq 1$. We evaluate all methods on four networks from the Netzschleuder database [Peixoto, 2020], chosen arbitrarily. In each case, we choose a budget k equal to 5% of the nodes of the graph.

Fig. 3 shows the performance for various biases $r \geq 1$. The networks Wiki and Polblogs are well-optimizable by almost all heuristics, possibly except high closeness centrality in the second case that performs visibly worse (though not by much). The Twitter graph is more challenging, while the Elegans graph is the most difficult to optimize, with high variability on heuristic performance. Unsurprisingly, the Greedy algorithm dominates the performance on this benchmark, which is expected given its theoretical guarantees (Theorem 3). Among the heuristics, high pagerank appears to be the most consistently performant, and the first one to match the performance of Greedy on the Elegans network.

7 Conclusion

We have studied the optimization problem of selecting a cardinality-constrained *seed set* of mutants that maximizes the fixation probability in the classic biased voter process. We have shown that the problem exhibits intricate complexity properties in the presence of biases $r \neq 1$. In particular, the problem is NP-hard in both regimes $r > 1$ and $r < 1$, while the latter case is also hard to approximate. On the positive side, we have shown that the optimization function is monotone and submodular, and can thus be approximated efficiently within a factor $(1 - 1/e)$. Interestingly, our results on optimization hardness imply that even just computing the fixation probability over randomly chosen seed sets is also NP-hard. Random seed sets arise in genetic/biological settings [Kaveh *et al.*, 2015; Allen *et al.*, 2020; Tkadlec *et al.*, 2020], due to the random nature of novel mutations. Interesting future work includes investigating better approximation factors for $r > 1$, and an in-depth search for good heuristics.

References

[Allen *et al.*, 2015] Benjamin Allen, Christine Sample, Yulia Dementieva, Ruben C Medeiros, Christopher Paoletti, and Martin A Nowak. The molecular clock of neutral evolu-

- tion can be accelerated or slowed by asymmetric spatial structure. *PLoS Computat Biol*, 11(2):e1004108, 2015.
- [Allen *et al.*, 2020] Benjamin Allen, Christine Sample, Robert Jencks, James Withers, Patricia Steinhagen, Lori Brizuela, Joshua Kolodny, Darren Parke, Gabor Lippner, and Yulia A. Dementieva. Transient amplifiers of selection and reducers of fixation for death-birth updating on graphs. *PLOS Computat Biol*, 16(1):1–20, 01 2020.
- [Antal *et al.*, 2006] T. Antal, S. Redner, and V. Sood. Evolutionary dynamics on degree-heterogeneous graphs. *Phys. Rev. Lett.*, 96:188104, May 2006.
- [Bhat and Redner, 2019] Deepak Bhat and S. Redner. Nonuniversal opinion dynamics driven by opposing external influences. *Phys. Rev. E*, 100:050301, Nov 2019.
- [Brendborg *et al.*, 2022] Joachim Brendborg, Panagiotis Karras, Andreas Pavlogiannis, Asger Ullersted Rasmussen, and Josef Tkadlec. Fixation maximization in the positional moran process. *arXiv preprint arXiv:2201.02248*, 2022.
- [Broom *et al.*, 2010] M Broom, C Hadjichrysanthou, J Rychtář, and BT Stadler. Two results on evolutionary processes on general non-directed graphs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2121):2795–2798, 2010.
- [Bui *et al.*, 1987] Thang Nguyen Bui, Soma Chaudhuri, Frank Thomson Leighton, and Michael Sipser. Graph bisection algorithms with good average case behavior. *Combinatorica*, 7(2):171–191, 1987.
- [Castellano *et al.*, 2009] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591–646, May 2009.
- [Clifford and Sudbury, 1973] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.
- [Díaz *et al.*, 2014] Josep Díaz, Leslie Ann Goldberg, George B Mertzios, David Richerby, Maria Serna, and Paul G Spirakis. Approximating fixation probabilities in the generalized moran process. *Algorithmica*, 69(1):78–91, 2014.
- [Domingos and Richardson, 2001] Pedro Domingos and Matt Richardson. Mining the network value of customers. In *ACM SIGKDD*, KDD ’01, page 57–66, 2001.
- [Durocher *et al.*, 2022] Loke Durocher, Panagiotis Karras, Andreas Pavlogiannis, and Josef Tkadlec. Invasion dynamics in the biased voter process. *arXiv preprint arXiv:2201.08207*, 2022.
- [Even-Dar and Shapira, 2007] Eyal Even-Dar and Asaf Shapira. A note on maximizing the spread of influence in social networks. In *WINE*, pages 281–286, 2007.
- [He and Yao, 2001] Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial intelligence*, 127(1):57–85, 2001.
- [Hindersin and Traulsen, 2015] Laura Hindersin and Arne Traulsen. Most undirected random graphs are amplifiers of selection for birth-death dynamics, but suppressors of selection for death-birth dynamics. *PLoS Comput Biol*, 11(11):e1004437, 2015.
- [Kaveh *et al.*, 2015] Kamran Kaveh, Natalia L Komarova, and Mohammad Kohandel. The duality of spatial death–birth and birth–death processes and limitations of the isothermal theorem. *Royal Society open science*, 2(4):140465, 2015.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *ACM SIGKDD*, pages 137–146, 2003.
- [Lengler, 2020] Johannes Lengler. Drift analysis. In *Theory of Evolutionary Computation*, pages 89–131. Springer, 2020.
- [Li *et al.*, 2018] Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. Influence maximization on social graphs: A survey. *IEEE TKDE*, 30(10):1852–1872, 2018.
- [Liggett, 1985] T. Liggett. *Interacting Particle Systems*. Classics in mathematics. Springer New York, 1985.
- [McAvoy and Allen, 2021] Alex McAvoy and Benjamin Allen. Fixation probabilities in evolutionary dynamics under weak selection. *Journal of Mathematical Biology*, 82(3):1–41, 2021.
- [Mossel and Roch, 2007] Elchanan Mossel and Sebastien Roch. On the submodularity of influence in social networks. In *STOC*, page 128–134, 2007.
- [Nemhauser *et al.*, 1978] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions–i. *Math. Program.*, 14(1):265–294, December 1978.
- [Ohtsuki *et al.*, 2006] Hisashi Ohtsuki, Christoph Hauert, Erez Lieberman, and Martin A Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502, 2006.
- [Peixoto, 2020] Tiago P. Peixoto. Netzschleuder: the network catalogue, repository and centrifuge. <https://networks.skewed.de>, 2020. Accessed: 2021-11-30.
- [Svitkina and Fleischer, 2011] Zoya Svitkina and Lisa Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. *SIAM Journal on Computing*, 40(6):1715–1737, 2011.
- [Talamali *et al.*, 2021] Mohamed S. Talamali, Arindam Saha, James A. R. Marshall, and Andreagiovanni Reina. When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, 6(56):eabf1416, 2021.
- [Tkadlec *et al.*, 2020] Josef Tkadlec, Andreas Pavlogiannis, Krishnendu Chatterjee, and Martin A. Nowak. Limits on amplifiers of natural selection under death-birth updating. *PLOS Computational Biology*, 16(1):1–13, 01 2020.