- 1 Title: Genetic Algorithms as a method to study adaptive walks in biological landscapes
- 2 Short title: Genetic Algorithms and evolutionary landscapes
- 3 Authors and affiliations:
- 4 Invernizzi, Edith^a and Ruxton, Graeme D^b
- 5 ^a Sir Harold Mitchell Building, School of Biology, University of St Andrews, KY16 9TH, St Andrews, United Kingdom;
- 6 email: ei21@st-andrews.ac.uk
- 7 ^b Dyers Brae Building, School of Biology, University of St Andrews, KY16 9TH, St Andrews, United Kingdom; email:
- 8 graeme.ruxton@st-andrews.ac.uk
- 9 Corresponding author for all stages of refereeing, publication and post-publication:
- 10 E. Invernizzi (ORCHID number: 0000-0002-3234-1016) is the corresponding author.
- 11 Address for correspondence:
- 12 Sir Harold Mitchell Building, School of Biology, University of St Andrews, KY16 9TH, St Andrews, United Kingdom;
- 13 email: ei21@st-andrews.ac.uk
- 14 Declaration of interests: None.
- 15

16 Abstract

17 The metaphor of fitness landscapes is common in evolutionary biology, as a way to visualise the change in allele or
 18 phenotypic frequencies of a population under selection. Understanding how different factors in the evolutionary process
 19 affect the trajectory of the population across the landscape is of interest to both theoretical and empirical evolutionary
 20 biologists. However, fitness landscape studies often have to rely heavily on mathematical methods that are not easy to
 21 access by biologically trained researchers. Here, we used a method borrowed from engineering - genetic algorithms - to
 22 simulate the evolutionary process and study how different components affect the path taken through a phenotypic fitness
 23 landscape. In a simple study, we compare five selection models that reflect different degrees of dependency of fitness on
 24 trait quality: this includes strengths of selection, trait-quality dependent reproductive hierarchy and the amount of
 25 stochasticity in the reproductive process. We include an analysis of other evolutionary variables such as population size
 26 and mutation rate. We analyse a game theory problem, as a test landscape, that lends itself to analysis through a
 27 deterministic mathematical simulation, which we use for comparison. Our results show that there are differences in the
 28 speed with which different models of selection lead to the fitness optimum.

29 Author summary

30 Evolution and adaptation in biology occurs in *fitness landscapes*, multidimensional spaces representing all possible
 31 genotypic or phenotypic combinations, where population adapt by following the cline of the fitness dimension. The study
 32 of adaptation on complex fitness landscapes has so far been limited by the need for mathematically heavy methods. Here,
 33 we present a simulation modelling framework, genetic algorithms, that can be used for evolutionary simulations of a
 34 population on a fitness landscape of chosen features and with custom evolutionary parameters.

35 Keywords

36 evolutionary modelling; evolutionary theory; adaptation; evolutionary algorithm; Sir Philip Sidney game

37 1. Introduction

38 1.1 Fitness landscapes and evolutionary biology

39 The evolutionary trajectory of a trait within a population is often represented as a walk in a *fitness landscape* (1) of as
 40 many dimensions as the number of (either phenotypic or genetic) components underlying fitness, plus one. The additional
 41 dimension is the fitness value that the combination confers to the carrier and is what we use to estimate the evolutionary
 42 trajectory of a population. For example, if there are two components then their values can be thought of as orthogonal x
 43 and y coordinates and fitness can be envisaged as the height of a landscape above the x - y plan. The shape of the landscape
 44 will influence the direction of a notional hill-climber who looks ever to be moving upward. This hill-climbing analogy
 45 has been useful in representing and studying evolutionary paths (despite the challenges of visualising a mapping of low
 46 and high fitness areas in higher dimensions) and has given rise to several formal models on the correspondence between
 47 genotype/phenotype and fitness (2,3). The study of these models has generated fundamental underpinning to the theory
 48 of adaptation, such as the expected distribution of mutation size (4–6). Recent technological developments in genomics
 49 and transcriptomics have allowed us to breakdown and reconstruct empirical fitness landscapes (7,8) that are answering
 50 our questions regarding their expected characteristics.

51 So far, no study of evolutionary processes and fitness landscape trajectories has looked at whether varying the relationship
 52 between the relative quality of a member of the population and their relative contribution to the next generation affects
 53 the evolutionary walk. For example, in a simplified scenario where we measure fitness based on one trait with multiple
 54 underlying genotypic or phenotypic components, the same trait variant is likely to achieve different reproductive fitness
 55 depending on the reproductive hierarchy within the population. This is likely to influence adaptation in many ecological
 56 systems.

57 The relationship between trait quality and reproductive fitness can be broken down into two key dynamics: 1. the
 58 relationship between trait quality and *expected* reproductive fitness (including, specifically in the factors analysed here,
 59 the strength of selection acting on the trait and the reproductive hierarchy of the population) and 2. the strength and nature
 60 of stochasticity in the reproductive process. Here, we call the interplay of these two factors *selection model* and we study
 61 its influence on population-level evolutionary dynamics by means of a simulation method derived from engineering:
 62 genetic algorithms.

63 1.2 Genetic algorithms

64 The exploration *versus* exploitation problem has been explored extensively in a particular subfield of evolutionary
 65 computations: genetic algorithms (GAs (9)). GAs were invented in 1960 by John Holland and colleagues as a part of the
 66 developing field of evolutionary computations (10). They exploit a natural selection-like process to find optimal solutions

to complex problems by evolving a starting population of solutions through selecting, mutating and recombining higher quality variants. This type of heuristic optimisation method is valuable when the solution space analysed is too big to be fully explored and it exploits the correlation in the landscape's ruggedness. Effectively, it looks for optimal solutions to problems where solutions are given by the non-linear interaction of multiple traits – similarly to a population's walk in the fitness landscape.

In evolutionary biology, GAs have been applied to problems whose solution by other means might be challenging: searching for efficient parameters for machine learning systems (such as neural networks, *e.g.* Montana and Davis, 1989), identifying or evolving solutions to game-theory problems (*e.g.* (12,13) and ecological niche models (such as GAs for Rule Set Production, or GARP (14)). They have also been deployed in the simulation of evolutionary or evolution-like dynamics (*e.g.* evolution of cognitive processes (15)), although the limitations in understanding and representing analytically what happens during the optimisation process have somewhat limited this latter application.

GAs in evolutionary biology may be usefully applied to investigate how different assumptions about the selection model and, virtually, any other component of the evolutionary process (mutation, sexual or asexual reproduction, structure of the reproductive interaction network within a group, just to name a few) affect the evolutionary trajectory of a trait. Effectively, we can use this tool to extrapolate common rules in the way these components affect the trajectory and to study whether the trajectory holds consistently (or varies consistently) across different fitness landscapes. This knowledge can then inform our prediction of the evolutionary trajectory in natural populations and also improve the theoretical study of evolutionary scenarios. GAs can specifically inform the study of systems where the evolutionary trajectory depends on non-linear (including epistatic) interactions between traits, and thus cannot be captured by deterministic methods (such as adaptive dynamics (16,17)).

Here, we start by analysing the effect of the selection model on a small phenotypic fitness landscape that is the product of the interaction of biallelic loci. We assume throughout that reproductive potential across the population is limited to a fixed number of reproductive slots at each generation. We use the five most common models of selection and reproductive slot allocation used in the GA literature as proxies for different biological contexts. These models are: k-tournament, truncation, proportional selection, linear and exponential ranking. We summarise in this introduction how each of these models influences the exploration-exploitation trade-off and we illustrate it with an equivalent biological scenario. We do not claim that our findings obtained by this method can be generalised to all evolutionary adaptive landscapes: as with every simulation modelling method, we identify dynamics that are the result of modelling assumptions – here, the characteristics of the fitness landscape under study. This method allows to the same analysis on different landscapes withing a fixed modelling framework that reduces unwanted sources of variation.

1.3 Selection model and the fitness landscape

Recent analyses of empirical landscapes show varying (often considerable) levels of ruggedness, at least for small genomes or gene complexes (18) (with *ruggedness* indicating the amount of epistatic interactions in the genetic architecture (19)). There is, essentially, in biological landscapes a degree of correlation that fits our theoretical understanding of evolutionary landscapes: on average, two neighbouring points in the landscape have similar fitness, while, the more mutations separate them (*i.e.* the further away they are in the landscape), the less likely their fitnesses are to be similar, until we reach a distance at which their fitnesses are uncorrelated (19). If the landscape is large enough, multiple fitness peaks might exist far apart. While theory predicts that it is possible (20), if not likely, that a path connecting any two viable genotypes always exist, this path might cross areas of lower fitness, so that a population might find a local peak (a *local optimum*), with higher fitness than the region it started in, but fail to discover the highest peak in the landscape (the *global optimum*). When mutations are not or cannot be large enough to reach distant areas of the landscape, the evolutionary process must rely on strong enough selection to reliably climb fitness peaks, while at the same time allowing exploration of less-fit phenotypes.

Varying selection models might have two distinct consequences on the movement of a population in the fitness landscape and we illustrate them here. Let us take a classic case study evolutionary scenario used in the evolutionary literature: a discrete fitness landscape with neighbouring points assumed to be at one-mutation distance from each other and with mutations assumed to be rare enough to arise one at a time. If the evolutionary process relies on a small group of high-quality individuals to reproduce, then the population will follow a highest-fitness neighbour trajectory and is likely to land on the closest fitness peak in the landscape. If this point is a local optimum, that optimum is where the population will stay until the next change in selective pressure that modifies the landscape. Maintaining sufficient variation in the population, by allowing lower quality individuals to reproduce, on the contrary facilitates the emergence of new variants away from the current optimum. A second consequence of the relationship between trait quality and reproductive fitness is the change in the speed with which the population moves along the trajectory: the higher the proportion of reproductive fitness assigned to the highest-quality variants, the quicker the movement. The higher the proportion that lower-quality variants benefit from, the more time the population needs to move to a point with higher average quality.

1.4 Explorative and exploitative selection models

We define *exploitative* systems as those biological scenarios where the highest-quality individual (or a small élite of high quality individuals) obtains a share of reproductive fitness that is disproportionately large compared to what would be expected should reproduction be directly proportional to the relative quality of the individual with respect to the average. An exploitative system is, for example, a non-eusocial group with a single reproductive male or female which suppresses subordinate reproduction and sires all new offspring. Conversely, an *explorative* system corresponds to a scenario where

most individuals have at least as much reproductive fitness as what would be expected from their relative quality in the population. In the absence of sexual selection (or in the presence of moderate sexual selection), biological scenarios in which individuals compete for each reproductive opportunity (for example, species in which reproductive pairs form and dissolve yearly) can be considered explorative and the amount of exploration is increased by stochastic events interfering with the fitness-proportional distribution of reproductive chances (*e.g.* in a species with a 20 year reproductive span, the individual with highest fitness expectations might stochastically die after only 1 year, but the individual with the lowest expectations could, stochastically, reproduce for 11 years).

In the study described below, we investigate how different systems affect the evolutionary walk of a population in a phenotypic landscape from game theory, the SPS game.

1.5 Landscape of the Sir Philip Sidney game

The fitness landscape we study here, the Sir Philip Sydney (SPS) game theory problem (21), has been chosen for ease of access and understanding at the time this study began and it cannot be described purely in terms of well-established formal models of fitness landscapes. Nevertheless, those shortcomings do not prevent it from being a useful representation of a biological landscape. It is a phenotypic gambit scenario of six independent traits contributing to the total reproductive fitness of the individual. Mutations are additive at the phenotypic level, but non-additive at the fitness level, meaning that the traits interact non-linearly to determine fitness. Our landscape differs from the assumptions commonly made in the literature in two ways. Firstly, it is discrete. Mathematically, this is the equivalent of any continuous landscape whose optima and minima are all represented on our discrete landscape. Secondly, it is a multi-phenotype frequency-dependent landscape. Although this is not a commonly studied landscape scenario, it is found in several real-world evolutionary case studies, such as apostatic selection or microbial competition experiments where multiple strains grow on a medium with multiple carbon sources.

1.6 Details of analysed selection models

We call selection model the set of assumptions in an evolutionary simulation that define the relationship between trait quality and reproductive success of the individual, as a function of the two dynamics listed in paragraph 1.1 (relationship to expected reproductive fitness and amount of stochasticity in the system). Different biological scenarios often imply different selection models. Below, we provide a brief summary of the models analysed here.

k-tournament: This selection model assumes that competition for each reproductive slot occurs in small, randomly formed groups, where the individual with the highest-quality trait wins with a specified probability. The stochasticity introduced by randomly selecting competitors and the uncertainty in the probability that the individual with the highest-quality

157 phenotype wins means that k -tournament allows for some stochasticity in the selection of the traits transmitted to the next
158 generation. This method is equivalent, in a biological scenario, to a group of individuals competing for a mate after being
159 drawn to the same location by a mating signal: competitors are brought together by contingency and the strongest does
160 not necessarily, albeit is more likely to, prevail.

161 *Truncation*: Truncation sets a fixed threshold to the proportion of the population that can reproduce and makes this
162 threshold dependent on trait quality. For example, only the top 10% of individuals in terms of trait quality may have the
163 opportunity to reproduce. Among this reproductive elite, chances of reproducing are then uniform (22). Making a parallel
164 with a biological scenario, we can compare this method of selection to dynamics in an animal group where a reproductive
165 hierarchy is established, by which only top-ranking individuals mate and reproduce.

166 *Proportional selection*: Proportional selection allocates reproductive slots in a manner that is statistically proportional to
167 trait quality, but it introduces some stochasticity in the allocation. In a biological equivalent, we can think of this method
168 of selecting reproductives as a system where individual reproductive success is strongly dependent on the absolute quality
169 of one trait – but failure of a high-quality individual to reproduce is still possible (simply through bad luck).

170 *Ranking*: We can visualise the ranking selection model as a population organised as a list of individuals ranked from
171 highest to lowest trait quality. Individuals are then assigned reproductive slots according to the place held in the ranking
172 (plus some additional stochasticity), rather than to the absolute trait quality value. The advantage of this method is that
173 the within-population variance in the expected number of reproductive slots remains fixed, evening out reproductive
174 chances within each generation and effectively enhancing exploration. In fact, because the range of the number of
175 allocated slots is independent of population trait-quality mean and variance at that time, the highest quality phenotype in
176 the population will receive the same expected reproductive proportion regardless of how far away it stands from average
177 population trait quality. Conversely, lower quality individuals will enjoy reproductive opportunities as long as their
178 ranking position is not too far below the rest of the population. A biological equivalent of this method is a scenario with
179 a hierarchically organised population where reproductive opportunities are proportional to the rank and where the number
180 of reproductive positions available remains relatively fixed across generations.

181 We also use this study to test the effect of other evolutionary parameters on the evolutionary trajectory. The complete list
182 of evolutionary parameters tested (including selection method parameters) is given in **Table 1**. We explain here the use
183 of one particular method, derived from GA literature, and its meaning as a biological parameter: *replacement*.

184 In a GA, replacement is the process by which new individuals replace members of the parent generation in a fixed-size
185 population. It relies on two features with important evolutionary implications in affecting the exploration-exploitation
186 balance: replacement size (or number of reproductive slots M relative to population size, $M:N$) and replacement criterion.

187 It is either generational, where the offspring completely replaces parents in the new generation ($M = N$), or overlapping,
188 where only a fixed proportion of the population is replaced ($M < N$). In biological terms, this is the equivalent of studying
189 the effect of the reproductive overlap between generations. Overlapping replacement allows trait variants already present
190 in the previous population to compete with newly-generated variants. Increasing M considerably reduces variance in the
191 growth curve of the best-variant proportion of the population, with generational replacement having the least variance
192 (23). It thus reliably produces quality-proportional strategy fitness, but relative to the quality of already-existing strategies
193 – the local fitness cline is exploited. The second critical feature is how strategies are selected for replacement. In the First-
194 In-First-Out (FIFO) GA method, individuals are replaced in chronological “birth” order, while in Worst-In-First-Out
195 (WIFO) replacement occurs either by quality-dependent proportional selection or by deterministic death of the worst
196 individual. This is the equivalent of studying the effect of the relationship (or absence thereof) between trait quality and
197 probability of death on the evolutionary trajectory. FIFO favours drift and thus exploration, granting each strategy the
198 same reproductive time frame, while WIFO is exploitation-based.

199 In this paper, we describe the results obtained from evolving a population towards a solution to the SPS game, under each
200 selection model and across different parameter values for selection, mutation, population size and replacement. Multiple
201 solutions to the SPS game have been found by different methods (see **Box 2**) and are considered the expected endpoints
202 of an evolving population. We compare the reliability with which each model reaches the evolutionary stable solutions
203 previously identified for this game: the evolutionary stable strategy (ESS) analytically found by Johnstone and Grafen
204 (24) and the evolutionary stable set (ES) obtained by Hamblin and Hurd (25).

Parameter	Description	Values or variants tested in our simulations
N	population size	{50,100,200}
M	number of available reproductive slots per generation	$\left\{1, \frac{N}{20}, \frac{N}{10}, \frac{N}{5}, \frac{N}{4}, N\right\}$
μ	mutation rate (<i>per locus</i> or <i>per chromosome</i>)	{0.001,0.003,0.005,0.007,0.01}
a	in the starting population, proportion occupied by the ancestral strategy [0,0,0,0,0]	{0.05, 0.10, 0.15, 0.20}

l	Used in <i>truncating selection</i> . Top proportion of the population (ordered by rank) that reproduces.	{0.1, 0.2, 0.3, 0.5}
k	Number of individuals selected to compete for a reproductive slot at each round, in <i>tournament selection</i> .	{2,3,5,7}
α	Used in <i>tournament selection</i> . Probability of fittest competitor to win.	1
v	Used in <i>proportional selection</i> . A strategy with quality equal to the population average can be expected to realise a number of reproductive slots equal to $\frac{vM}{N}$ in the next generation.	{0.70, 0.80, 0.90, 1}
z	Used in <i>linear ranking selection</i> . The highest-quality solution has a maximum of $\frac{2zM}{N}$ reproductive slots that they can be expected to realise in the next generation.	{0.55, 0.60, 0.65, 0.70}
c	Used in <i>exponential ranking selection</i> . The factor by which the number of expected reproductive slots decreases per every unit of rank.	{0.55, 0.60, 0.80, 0.90}
mutation type	whether μ indicates the value for each genotype or at each locus of a genotype	{per chromosome, per locus}
replacement type	method with which one member of the parent generation is selected for replacement by one offspring	{WIFO, FIFO}

205 **Table 1. Description of parameter tested with listed values.**

206 **Box 1: Individual, genotype and phenotype in a GA and in biology**

207 In a GA, populations consist of multiple possible solutions to the problem under study that are undergoing a selection
208 process for a high-quality solution. This population is the equivalent of a biological population of individuals displaying
209 variation in the trait (solution) under selection. Each solution is encoded by a genotype, consists of a string of values, one
210 (*locus*) for each of the parameters or features that contribute to a solution. For example, in a GA genotype modelling

leopard predatory effectiveness as a physiological trait, the genotype might consist in a series of loci encoding the parameter values for: muscular mass, fibre elasticity and the parameter defining the developmental and physiological trade-off between the two. The solution built under this genotype is a combination of muscular power and elasticity used to respond to predatory encounters, in which prey size and speed are drawn from fixed probability distributions. While phenotype and genotype overlap in one-locus traits, the difference becomes apparent wherever we introduce an interaction between loci. In our predatory effectiveness trait, it is the trade-off parameter that causes a phenotype to be different than the effect of the power and elasticity loci taken independently. This resembles the complexity observed in real biological traits. Similarly to biology, moreover, an additional layer of complexity can be introduced at the strategy-generation step if noise or plasticity (learning) mechanisms influence the trait. Multiple individuals in the same population might carry the same genotype and, indeed, the change in genotype frequency in a population over generations is one currency for measuring its success.

In the game theory problem modelled here, we are looking at the evolution of a behavioural strategy, in a signalling-for-resources scenario. Each strategy is a combination of behavioural responses, each encoded at a separate locus. Here, we call the behavioural strategy a *strategy (solution)* and we follow GA terminology in using *genotype* to indicate the combination of behavioural responses that make up a strategy, despite them being the phenotypes in the phenotypic landscape.

2. Methods

To test the effect of selection models over evolutionary trajectories, we evolved solutions to the Sir Philip Sidney game under different models, using a genetic algorithm. We applied five different selection models, each evaluated across a range of parameter values. **Table 1** summarises the parameters analysed and lists the values tested for each.

2.1. Genetic Algorithm

The workflow of a GA is depicted in **Figure 1**.

Fig 1. Workflow of a GA. Simplified representation of the processes in a GA (in white), that pool, select and modify the initial *population* (in red) to generate a new population. The *parental pool* are the individuals in the current population that meet the trait-quality requirements to reproduce. They produce copies of themselves (*parental copies*) that undergo mutation and crossover to generate the final *offspring* individual, as in asexual selection. Offspring take the place of individuals of their parent's generation in the population. Each cycle through this sequence of steps constitutes one generation.

239 2.1.1. Chromosomes

240 Strategies are encoded by 6-locus strings, or *chromosomes*, with each locus encoding either a 0 or a 1.

241 2.1.2. Population and replacement size

242 The population consists of N individuals, of which M are replaced by offspring strategies at each generation.

243 2.1.3. Trait quality evaluation

244 This is where the quality of the strategy generated by each individual n with genotype g_i in the population is tested. We
 245 define trait quality as the average outcome of the test; this value is converted into fitness at later stages. In this simulation,
 246 quality is evaluated through competition against randomly selected strategies from the same population over five rounds,
 247 as in (25). At each round, an “opponent” strategy is chosen among the $N - 1$ remaining and takes the complimentary role
 248 (*i.e.* donor or signaller) to the strategy evaluated. Role, health state and degree of relatedness are assigned with probability
 249 equal to population frequencies in the equilibrium parameter range of (24). Performance is calculated from the interaction
 250 of the two strategies in each round and trait quality as the average of all rounds.

251 2.1.4. Selection

252 We assume throughout that reproduction is asexual, requiring only a single parent. Selection defines which subset of the
 253 parental population will be passed on to the next generation and how many breeding slots (proportion of the offspring
 254 generation) each parent solution receives.

255 *k-tournament*: k randomly selected individuals are compared in each of N rounds. In each round, the individual with
 256 highest quality is selected with probability α . If the highest-quality individual is rejected, then the second fittest is chosen
 257 with probability α , *etc.* Selected individuals re-enter the population pool after reproduction and can thus be selected
 258 multiple times.

259 *Truncation*: Individuals are sorted by decreasing quality and parents selected from the l topmost proportion of the list are
 260 selected for reproduction. Note that individuals with equal trait quality might be separated by the threshold and some not
 261 enter the selection pool. Parent individuals are selected with uniform probability from the selection pool to enter the parent
 262 pool.

263 *Proportional selection*: Individuals are assigned an expected value of reproductive slots, which is, the statistical mean of
 264 the number they should receive in an infinite number of trials. The assignment function used to derive expected values is

265

$$E[m_{n,t}] := \begin{cases} \left(1 + \frac{h_n(t) - \bar{h}(t)}{2\sigma(t)}\right) \frac{M}{N}, & \sigma(t) \neq 0 \\ \frac{M}{N}, & \sigma(t) = 0 \end{cases}$$

from (26). The equation represents the expected number of reproductive slots $m_{n,t}$ that individual n will be allocated at time t , given its estimated quality at that time step, $h_{n,t}$, and the average quality in the population at that time step, $\bar{h}(t)$, and adjusted according to the standard deviation of quality within the population at that time step, $\sigma(t)$. Variance-scaling is applied to limit the effect of drift from the highest quality value, making reproductive chances more widely distributed when variance is high and less when it is low. We modify (26) to control how reproductive slot assignment depends on the quality value. Via proportional selection, we can set the expected value for an individual with average quality through a parameter v . This is done by multiplying:

$$E[m_{n,t}] := \left(1 + \frac{h_n(t) - \bar{h}(t)}{2\sigma(t)}\right) \frac{vM}{N}.$$

When $v = 1$, an average quality individual receives as many reproductive slots as if assigned by a uniform quality-independent distribution. If we set $v < 1$, only higher-than-average individuals will expect at least one reproductive slot. The lower v , the greater the reproductive advantage of high-quality individuals over those of average quality. Thus, the lower v , the higher the exploitation.

The effective number of slots is assigned using a sampling algorithm: a random integer r between 0 and the sum of expected values is drawn for a number of rounds M equal to the number of reproductive slots and the expected value of individuals summed as the population is looped through in a fixed order, until the expected value of an individual makes the total equal to or higher than r and that individual wins one reproductive slot. Slot allocation following this system is thus statistically proportional to assigned expected values, with drift. *Linear ranking*: Individuals are ranked in increasing order of quality and the rank used to assign the expected value through an assignment function. The linear implementation of the assignment function (27) is:

$$E[m_{n,t}] := \min + (\max - \min) \times \frac{\text{rank}(n,t) - 1}{N - 1},$$

where $\min$ is the expected value allocated to the worst individual in the population at time t and $\max$ that allocated to the best individual. In this type of implementation, it is possible to manipulate the expected fitness of the best individual to adjust the exploitation-exploration balance, in the same way as we changed the expected fitness of the average individual in proportional selection. $\max$ can be set to a different value from its maximum, $2M/N$, through a parameter z similar to v in proportional selection, so that

$$Max = \frac{2zM}{N}$$

(see the derivation in SI.1.).

The effect of z is analogous to that of v in proportional selection, but with opposite effects as it acts on the highest-quality individual: the higher z , the higher the exploitation and the lower z , the higher the expected fitness of lower quality individuals.

Reproductive slots are then assigned with a sampling algorithm identical to that in proportional selection.

Exponential ranking: The exponential implementation of the assignment function takes the form

$$E[m_{n,t}] := c^{N-\text{rank}(n,t)} \quad \text{and} \quad 0 \leq c \leq 1,$$

where c determines the steepness of the exponential increase in expected values with rank. By changing c , we can control the exploration-exploitation balance: a high c assigns exponentially higher advantage to the highest-quality individual, thus creating high exploitation.

The effect of z , v or c is that of fine-tuning the proportion of individuals effectively reproducing (what Baker (27) calls “percentage involvement”).

2.1.5. Mutation

Parents that are assigned breeding slots through selection generate copies of themselves. Such copies undergo mutation to generate new variants and explore the solution space around parent solutions. We do not explicitly assume a “strong selection, weak mutation” (SSWM) regime: we instead model mutation through a mutation rate parameter and a mutation method parameter that span through a range of mutation regimes. In *per chromosome* mutation, one locus on the chromosome is selected with uniform probability and with probability μ_c its value switches to the opposite binary value. In *per locus* mutation, the chromosome is scanned and every locus switches with probability μ_l . As a result, our study investigates SSWM conditions, under low mutation rate, per chromosome method and small reproductive slot size, as well as biologically unlikely high-mutation scenarios, when reproductive slot size is very large or under per locus mutation. We chose not to include crossover in the analysis, because, in small fitness landscapes, mutation should be sufficient to explore the full space. In accordance with this expectation, (25) found that crossover does not affect the endpoint reached. After mutation (and crossover), the offspring set has been generated and is input into the population, entering the next generation.

318 2.1.6. Replacement

319 At the end of each cycle, M individuals among the population's N are replaced by the offspring. Individuals in the parent
320 generation are chosen for replacement either on the basis of quality (Worst-In-First-Out or WIFO), where the M lowest
321 quality individuals are replaced, or in a time-sequence manner, with oldest individuals being replaced (First-In-First-Out
322 or FIFO). We investigate both WIFO and FIFO.

323 2.1.7. Running time and data collection

324 The algorithm is run for 500 generations, as in (25). Frequency of each possible strategy in the population and population
325 fitness values are recorded every 50 generations.

326 2.1.8. Implementation

327 Each combination of selection model and parameter values was run in 10 replicates, each with random number generators
328 seeded with a different integer value between 0-9. The simulations were implemented in Python version 2.7.12 using
329 Cython language version 0.27.3.

330 Box 2: Sir Philip Sidney game

331 The SPS game is a game theory problem in evolutionary biology used (first by (21)) to study the evolution of costly and
332 honest signalling within an inclusive fitness framework. The problem of signalling for help when health is poor is
333 evolutionarily interesting because, once honest signalling has evolved, individuals who are in good health benefit by
334 signalling as if they were poorly and receiving the resource at the expense of the donor. This means that honest signalling
335 is evolutionarily unstable. However, honest signalling might be maintained through a combination of inclusive fitness -
336 conferring an indirect fitness benefit to the helping donor and to the honest signaller - and of signals that are costly to
337 produce - enforcing lower dishonesty. The SPS game recreates a simple case study scenario with a population of
338 individuals of relatedness r who can transfer a fitness-increasing resource to each other. Individuals can choose whether
339 to ask for the resource, paying a cost c to signal for help, and whether to donate it, lowering their fitness but increasing
340 that of the receiver. Individuals should take into account their own level of fitness (usually called "thirst", for fidelity to
341 the Sir Philip anecdote), as well as the potential donor's, to decide whether or not to signal. The game asks which
342 conditions of relatedness and cost allow the evolution of completely or partially honest signalling.

343 The mathematical models following Maynard Smith's (*e.g.* (24,28–31)) have tested the problem under different scenarios
344 by varying underlying assumptions; however, evolutionary simulations have challenged the notion that ESSs, obtained
345 through either mathematical analysis or deterministic evolutionary simulations (*e.g.* (32)) and albeit possible under
346 specific conditions, should be the expected evolutionary outcome.

347 2.2. SPS game

We use the SPS game (see **Box 2**) as a testbed problem. We chose to reproduce the study by Hamblin and Hurd (25), who evolved solutions to the SPS scenario investigated by Johnstone and Grafen (24) using a GA. The signalling scenario studied by (24) consists of a population of donors and of two types of signallers, a close relative with relatedness r_1 to the donor (type I) and a distant relative (type II) with relatedness r_2 . Signallers can be in either of two states: thirsty and dying in the absence of the resource (fitness = 0), and healthy, with fitness F_b in the absence of the resource. Resource transfer re-establishes full fitness. Signallers are healthy with probability o and needy with probability $(1 - o)$ and closely related to the potential donor with probability q and distantly related with probability $(1 - q)$. Similarly, donors have fitness of 1 with the resource and of F_d without. The donor can decide whether or not to transfer the resource, based on whether a signal received, and signallers can decide whether to signal or not. Signalling has a cost $d = d_1 \times F$ for close relatives and of $d = d_2 \times F$ for distant relatives, where F is the signaller's fitness after the donor has made its decision. The optimal strategy for the donor will depend on the degree of relatedness to the signaller and on the proportions of relatives of each degree in the population, while the optimal strategy for signallers will depend on relatedness and the cost of signalling.

We searched for solutions for the same SPS game parameter set used by (25) and based on (24) (also accepting (25)'s correction for $d_2 = 0.1$ for the semi-separating equilibrium to be an analytical ESS solution to the game). This parameter set is a point in parameter space at which the analytical ESS is expected. This ESS is a biologically interesting case of partially-honest signalling and is the highest-payoff strategy when only one strategy is present in the population (at this point in parameter space). However, most commonly the population seems to reach a different convergence point, when an evolutionary simulation is used (25): the ES, a two-strategy endpoint that has the highest average payoff in the population when every strategy is equally represented at time zero (see **Supporting File 1**).

2.2.1. Chromosome representation of a strategy

The six loci represent, in order: donor strategy on signal received, donor strategy on no signal (1 = transfer resource, 0 = no transfer), closely-related signaller strategy when thirsty, closely-related healthy signaller strategy when healthy, loosely-related signaller strategy when thirsty, loosely-related healthy signaller strategy when healthy (1 = signal, 0 = no signal) (**Figure 2a**).

Fig 2. How SPS strategies are encoded in the GA. A. Structure of a GA chromosome encoding a strategy used in the SPS game. Each locus (A-F) contains a binary value indicating whether that behavioural response is used by the carrier (1 = yes, 0 = no). The first two loci are behavioural responses from the donor of the resource: transfer the resource when hearing the signal (A) and transfer the resource when *not* hearing the signal (B). The other four loci encode the responses of the signallers to their own health status and to the degree of relatedness to the donor. Loci C-D encode the responses towards closely related donors: emit a signal if thirsty (C) and emit a signal if healthy (D). Loci E-F encode the same

379 responses towards distantly related donors. **B.** The table shows how the ESS and the two strategies within the ES are
380 encoded according to the scheme above.

381 2.2.2. Initialisation

382 At the start of each simulation, a proportion a of the population consists of the putative ancestral strategy [*no-transfer*,
383 *no-signal*]. The remaining strategies are randomly drawn among all possible SPS game solutions.

384 2.2.3. SPS game parameter values

385 Parameter values used: $o = 0.6$, $q = 0.9$, $r_1 = 0.5$, $r_2 = 0.2$, $F_b = F_d = 0.8$, $d_1 = 0.4$, $d_2 = 0.1$.

386 At this point in parameter space, we find the ESS and the ES detailed in **Figure 2b**.

387 2.3. Deterministic model

388 To have a benchmark for estimating the amount of exploration generated by each selection model, we produced a
389 deterministic model and compared its results with those obtained by the GA. The model assumes a population of fixed
390 size equal to the total possible strategies in the SPS game ($N = 64$). At generation 0, each strategy is equally represented
391 in the population (*i.e.* there is exactly one copy of each strategy). Fitness in the model is given exclusively by the relative
392 quality of each strategy. We define the fitness of strategy i at time t as

$$393 \quad w_i(t) := f_{i,t} \sum_{j=1}^N u_{i,j} f_{j,t},$$

394 where $f_{i,t}$ is the frequency of strategy i in the population at time t and $u_{i,j}$ is given by the average value obtained by i
395 when competing against strategy j in the SPS game, defined as the average between playing as the donor and playing as
396 the beneficiary. These two components of $h_{i,j}$ were calculated as in (24), for the same values of SPS game parameters
397 used in the GA (*i.e.* for the point in the SPS parameter space where the ESS is the strategy detailed in **Figure 2**). The
398 average fitness obtained by all strategies at time t is

$$399 \quad \bar{w}(t) := \sum_{i=1}^N \frac{w_i(t)}{N}.$$

400 The relative fitness of strategy i is thus

$$401 \quad W_i(t) := \frac{w_i}{\bar{w}(t)},$$

402 which becomes the frequency of strategy i at time $t+1$.

403 The model has no mutation nor recombination. We calculate population frequencies for 500 generations, which is the
404 same number of generations used in the GA.

405 The deterministic model was built and run in MATLAB R2020a. The matrix of values resulting from the interaction of
406 all pairs of strategies was built in Python version 2.7.16.

407 2.4. Analysis

408 2.4.1. Fitness landscape exploration

409 To quantify the effect of parameters and methods on landscape exploration, we only considered those simulations in which
410 the ESS or ES reached and maintained a population-level frequency higher than 80% for the last 50 generations of the
411 simulation. We consider those runs as having converged to that solution (as in (25)). For each selection model, we
412 calculated the proportion of runs that converged to each solution type. Within each selection model, we then compared
413 the effect of each parameter value of each evolutionary parameter by looking at the proportion of simulations with that
414 value that had reached each solution type, across all values of other parameters. Mean and standard deviation were
415 calculated as the mean and standard deviation of replicates with the same seed. In total, 72,000 simulations were run for
416 each selection model and 12,000-36,000 simulations were run for each parameter value within selection model (that is,
417 across all values of other parameters). This difference in the number of simulations for each parameter value is an effect
418 of grouping runs along single parameter dimensions: for example, if we calculate statistics for each value of M within one
419 selection model, we analyse $72,000/6 = 12,000$ simulations (where 6 is the number of values of M we tested in total; see
420 **Table 1**), while, if we do the same calculations for N , we look at $72,000/3 = 24,000$ simulations.

421 2.4.2. Statistical analysis

422 We tested for a significant difference in the distribution of number of seeds converged, across selection methods, to three
423 solution types: ESS, ES and all other solutions. We used a Fisher's Exact test with simulated p-values by Monte Carlo
424 simulations, 10,000 iterations. Statistics were implemented in R, version 3.5.1 (33).

425 3. Results

426 3.1. Different selection models lead to different outcomes

427 We reproduced the analysis run by Hamblin and Hurd (25), who used a GA to explore a point in the parameter space of
428 the SPS scenario studied by Johnstone and Grafen (24). We used this case study to analyse the effect that different
429 assumptions made on the evolutionary process have on the trajectory of a simulated evolving population. Specifically, we
430 were interested in the evolutionary outcome under multiple selection models. We also looked at the effect of the
431 evolutionary parameter values listed in **Table 1**. Hamblin and Hurd's algorithm reaches quite different evolutionary
432 endpoints from the original 1993 study of this SPS game variant, with the ESS identified analytically at this point in

parameter space (*[give-on-signal, signal-when-need, always-signal]* for [donors, closely-related signallers, loosely-related signallers]) effectively emerging with very low frequency. More frequent as a solution is an ES of strategies, *[give-when-no-signal, never-signal, never-signal]* and *[always-give, never-signal, never-signal]*.

We analyse the effect of each selection model on the point in the phenotypic landscape which a population reaches within 500 generations. We compare these results to the convergence point of a fully deterministic (*i.e.* perfectly fitness-proportional) model run for the same number of generations, to better understand the amount and effect of exploration in each of our partially stochastic selection models. For convenience, we define three types of solutions to the SPS game: the Evolutionary Stable Strategy (ESS), the Evolutionary stable Set (ES) and every other solution. We can envisage that exploitation of the best individual will lead to the ES, more commonly than to the ESS, due to the former being the highest average fitness strategy, while the ESS needs high frequency in the population *before* reaching the highest payoff. In accord with this latter statement, we can see in (25) that ESS convergence increases with the frequency at which it is seeded in the initial population.

If we look at the proportion of simulations reaching the two evolutionary stable solutions (ESS and ES) in our study (**Figure 3**), we see that no model reliably reaches the ESS and that there is no difference among the models in the frequency with which they land on this solution (**Figure 3a, bottom panel**). The ESS is thus not an evolutionary attractor. When we shift our attention to the ES, we realize that different selection models reach the highest fitness solution with very different frequencies (**Figure 3a, top panel; note the difference in the y-axis scale with the bottom panel**; Fisher's Exact test, $p\text{-value} < 0.001$).

Most of our selection models reach the ES with very high frequency. If we consider that, in a selection model where fitness equals relative trait quality such as the deterministic model, the ES only reaches a frequency of 0.5 in the population in 500 generations (**Figure 3b**) and that we have set the threshold frequency to consider a strategy as the evolutionary endpoint to 0.8 (maintained over the last 50 generations), these frequencies are considerably higher than expected. What is happening here? Reproductive fitness in our model is dependent, as in most biological scenarios, on an estimate of trait quality. This estimate is obtained from five “events”, each under randomly assigned “conditions” (relatedness, health state, identity of the other strategy and the role played, in this SPS game) – it is a combination of stochasticity and of *median* quality, rather than average, that creates the estimate. Our ES has high estimated quality under most condition (**Supplementary file 1**: strategy 1 in the ES has the highest payoff in almost half of the pairwise competitions with other strategies: 30/64, including itself; strategy 2 has the highest payoff only in 3/64 pairwise competitions, including against strategy 1 and itself – it is an equivalent phenotype to strategy 1 at equilibrium). However, many other strategies have high quality in more than one condition. In many conditions where the ES has maximum quality, one or more other

463 strategies share this property. Essentially, the average landscape is relatively flat, which means that there is a good chance
 464 of error when estimating overall trait quality from only a small number of samples.

465 Competition within small groups (k-tournament) favours strategies with high mean estimated fitness: an individual
 466 carrying that strategy has low chances of competing against an individual with higher estimated fitness and, if they win,
 467 they automatically reproduce. Larger sizes of competition groups slow down the diffusion of the ES (**Figure 4b**). Even
 468 in a relatively flat landscape, assigning fitness in a manner directly proportional to estimated trait quality is enough to
 469 reach quick and reliable convergence to the ES (proportional selection). This is independent of the estimated number of
 470 reproductive slots assigned to the average individual, v , (**Figure 4c**) and we believe this is because drift occurs only
 471 towards high quality, not average, solutions (the average mean payoff from pairwise interactions among all possible
 472 strategies is 1.056; the average mean payoff of strategies that obtain the highest payoff against at least one other strategy
 473 is 1.095). Ensuring that everyone in the population reproduces (ranking selection) makes convergence slower. Even a
 474 reproductive hierarchy that assigns a lot of advantage with every step up in the rank (exponential ranking) reaches the ES
 475 in 500 generations with much lower frequency than with competition within groups (k-tournament) or trait-quality
 476 proportional selection. Increasing the advantage for each step in the rank (c ; **Figure 4e**) increases the speed of convergence.
 477 Linear ranking is the slowest method and seems to work at its best in small populations (N ; **Figure 4d**). However, limiting
 478 reproduction to the individuals with the highest estimated quality (truncation) also slows down convergence to the ES.
 479 The larger the share of the population that is allowed to reproduce, the faster the convergence (l ; **Figure 4a**). Based on
 480 these observations, we think that a strict regime of selective reproductive allocation accentuates the effect of the error in
 481 the quality estimate.

482 **Fig 3. Proportion of simulation converging to ESS or ES by selection method and deterministic fitness of ESS and**
 483 **ES. A.** Proportion of simulations that converge to each solution type, of those run with the specified selection method
 484 (across all other parameters and parameter values). Note that the top and bottom panel are plotted across different y-axis
 485 values. The standard deviation bars reflect the variation among runs with different random seeds, collapsed across all
 486 parameter values. **B.** In the deterministic model, the ES and the ESS reach these frequencies at generation 500. The
 487 frequency of the ESS is $3.8e^{-47}$. These frequencies are the outcome of a selection process where fitness exactly corresponds
 488 to relative trait quality within the population, with each strategy being present once in the population at time zero and in
 489 the absence of mutation.

490 **Fig 4. Proportion of simulations converging to ESS or ES by parameter value, within each selection method.** Each
 491 group of plots shows the proportion of simulations that converge to each solution type by parameter type and value (*i.e.*
 492 across all other values and all other parameters), within a selection method. The standard deviation bars indicate the

variation found among runs with different random seeds, collapsed across all parameter values. mut = mutation, pc = per chromosome, pl = per locus, rpl = replacement.

3.2. Effect of other parameters and methods

We can identify some general effects of other evolutionary parameters and of the mutation and replacement processes by looking at the general performance across selection models (**Figure 4**).

- In this small fitness landscape, a high mutation regime decreases the speed and reliability of convergence. However, even under very high mutation rate (μ) values, most methods reliably reach the ES in 500 generations.
- A trait-quality dependent death rate (WIFO) increases the speed of convergence across all selection models, compared to an age-dependent death rate (FIFO) (this is in agreement with results found in GA literature (34)).
- A larger population size increases convergence. The effect seems to level off once population size becomes sufficiently larger than the solution space (>64), although we would need to run simulations with intermediate population sizes (e.g. $30 < N < 100$) to confirm this effect. The effect of large N on convergence is true under all selection models, with two exceptions. Under linear ranking, a small population size increases the frequency of convergence; under truncation, population size seems irrelevant – probably because the key parameter here is the proportion of the population that is reproductively active, l .
- Having an intermediate number of reproductive slots available at each generation accelerates convergence, probably by counterbalancing the high mutation rate. The case $M = 1$ shows, nevertheless, a higher frequency of convergence than expected: under linear ranking, this is the only point in the parameter space we analysed where convergence occurs reliably in so many generations. We hypothesise that, when $M = 1$, slots are statistically allocated in the manner that more closely resembles the expected value than other low M values, eliminating some drift. However, further studies and a better statistical insight are necessary to confirm this hypothesis. The effect of M can be interpreted as inversely correlated with the variance between an individual's observed and expected growth curves between generations, according to the GA literature (23).
- A starting population with high frequency (a) of a high-quality but unstable strategy (in this case, the putative ancestor strategy; this strategy has a mean pairwise payoff of 1.102, much higher than the average, but only achieves the highest possible payoff against six other strategies) does not slow down convergence to the ES in this size of fitness landscape.

4. Discussion

We can think of three measures of the impact of an evolutionary factor on the evolutionary trajectory. We can define, assuming all other conditions fixed: the *speed of convergence* to the global optimum, that is, the average time with which

a population converges to it; the *frequency of convergence*, which is the frequency with which populations reach the global optimum, given infinite evolutionary time; and the *accuracy of convergence*, that is, the amount of time, on average, that a population spends with one or more solutions present at very high frequency (e.g. >80%) in its pool, independently from whether it will ultimately converge to the global optimum or not. Our results show that different assumptions made on the selection model affect the convergence speed. We have also endeavoured to find common patterns in the effect of components of the evolutionary process other than the selection model (e.g. population size, mutation rate).

We want to stress the limitations of generalising other details of our results to other fitness landscapes. Some fitness landscape features are independent from simplifications that are made at the simulation building stage when studying a biological system: the small size of this landscape and the relatively small difference in fitness between many of its strategies (the ‘flatness’ of the landscape). They nevertheless affect the trajectory of the population over time. Other features we have chosen not to include for simplicity or inapplicability to this game theory landscape: sexual reproduction and hybrid fitness are two of them. They are likely to be major factors influencing the evolutionary trajectory. We believe it to be an interesting future direction to investigate how much of these findings holds true in a larger and differently contoured fitness landscape and under a more complicated relationship between phenotype and genotype. Among the characteristics we suggest testing there are also the assumption of a continuous landscape and varying mutation sizes.

How important is exploration in reaching the global optimum? Our results show that, the larger the gap in reproductive advantage between the best quality individual and the second best (that is, the greater the exploitation), the faster the convergence, even in a relatively flat landscape. This performance is likely linked to landscape size, relatively to mutation rate and size. This fitness landscape is a high connectivity scenario, because of its small size, meaning that all areas are relatively close together. Therefore, we expect the landscape to have been thoroughly explored by mutation within the 500 generations of our model. High connectivity also means that exploitation is enough to reach the global optimum. In larger landscapes, exploration might play a more important role on the evolutionary endpoint reached, including whether the population can find the global optimum.

The ability of exploitation to find the global optimum is also limited to cases where the optimum is an evolutionary attractor: in the fitness landscape used here, the location with the highest average fitness is also stable (ES), but there are likely to be biological scenarios in which highest average fitness leads to instability and exploration might be essential to reach a stable outcome.

Our investigation focuses on the fitness optimum. However, natural populations are likely to stay in area of high fitness close to but not exactly matching the optimum, due to other constraints including the trade-off between different traits. New analyses should investigate differences between methods in converging to these extended “optimality zones”.

552 Finally, our results also show that the accuracy of trait quality estimation affects the amount of drift observed. Accordingly,
553 the details of our results are likely to be sensitive to the number of rounds we used in trait evaluation

554 We have provided, in this article, an overview of the methodology of GAs. This modelling framework offers a useful tool
555 to evolutionary biologists interested in the effect of evolutionary process components on a population's evolutionary
556 trajectory. The advantage of this modelling architecture is that it is already set up to mimic the evolutionary process, thus
557 limiting the degrees of freedom compared to *ad hoc* modifications of a more universal tool like agent-based modelling.
558 This enhances comparability and replicability. GAs, on an equal level to agent-based models, have the flexibility needed
559 for modelling the effect of more complex dynamics, such as the relationship between phenotype and genotype, on the
560 evolutionary outcome. For example, as mentioned, a key feature of our simplistic case study is the absence of phenotypic
561 plasticity. This can easily be integrated into the simulation.

562 **5. Conclusion**

563 We have highlighted here the use of genetic algorithms as a simulation framework for studying the effect of all
564 components of the evolutionary process, by presenting a case study of five selection models and their effect on the
565 evolutionary trajectory of a population. Although we emphasise that this outcome cannot be generalised to all possible
566 correlated fitness landscapes, the results outlined here are nevertheless interesting: they highlight substantial differences
567 in the way different assumptions about the evolutionary process influence the expected evolutionary trajectory. We hope
568 that this modelling framework will be used to integrate both known and predicted characteristics of fitness landscapes in
569 the study of evolutionary trajectories. The resulting findings will be useful both to evolutionary theorists and ultimately,
570 if used to generate more accurate models of trajectories for existing populations, to empirical scientists.

571

572 **Acknowledgements**

573 Funding: This work was funded by the John Templeton Foundation as part of the research collaboration grant “Putting
574 the extended evolutionary synthesis to the test” (grant no. 60501) and by the University of St Andrews. Our thanks go to
575 Dr Barbara Trubenova for the keen feedback that has helped reshape this manuscript into a readable form and to Petri
576 Rautiala for essential discussions on the mathematical interpretation of the SPS game landscape.

577

578 6. References

- 579 1. Wright S. The roles of mutation, inbreeding, crossbreeding and selection in evolution. Proc Sixth Int Congr
580 Genet [Internet]. 1932;356–66. Available from:
581 <https://www.blackwellpublishing.com/ridley/classictexts/wright.pdf>
- 582 2. Fragata I, Blanckaert A, Dias Louro MA, Liberles DA, Bank C. Evolution in the light of fitness landscape theory.
583 Vol. 34, Trends in Ecology and Evolution. Elsevier Ltd; 2019. p. 69–82.
- 584 3. Orr HA. The genetic theory of adaptation: A brief history. Vol. 6, Nature Reviews Genetics. 2005. p. 119–27.
- 585 4. Fisher RA. The genetical theory of natural selection. Oxford: Oxford University Press; 1930.
- 586 5. Kimura M. The neutral theory of molecular evolution. Cambridge: Cambridge University Press; 1983.
- 587 6. Orr HA. The population genetics of adaptation: the distribution of factors fixed during adaptive evolution.
588 Evolution (N Y). 1998 Aug;52(4):935–49.
- 589 7. Sarkisyan KS, Bolotin DA, Meer M V., Usmanova DR, Mishin AS, Sharonov G V., et al. Local fitness landscape of
590 the green fluorescent protein. Nature. 2016 May 11;533:397–401.
- 591 8. Steinberg B, Ostermeier M. Environmental changes bridge evolutionary valleys. Sci Adv. 2016 Jan 1;2(1).
- 592 9. Holland JH. Adaptation in natural and artificial systems : an introductory analysis with applications to biology,
593 control, and artificial intelligence. MIT Press; 1975. 211 p.
- 594 10. Holland JH. Outline for a Logical Theory of Adaptive Systems. J ACM [Internet]. 1962 Jul 1 [cited 2018 Dec
595 4];9(3):297–314. Available from: <http://portal.acm.org/citation.cfm?doid=321127.321128>
- 596 11. Montana DJ, Davis L. Training Feedforward Neural Networks Using Genetic Algorithms. In: Proceedings of the
597 Eleventh International Joint Conference on Artificial Intelligence [Internet]. 1989 [cited 2018 Dec 15]. p. 762–
598 7. Available from: <https://www.ijcai.org/Proceedings/89-1/Papers/122.pdf>
- 599 12. Axelrod R. The Evolution of Strategies in the Iterated Prisoner’s Dilemma. In: Davis L, editor. Genetic
600 Algorithms and Simulated Annealing. Los Altos, CA: Morgan Kaufman; 1987. p. 32–41.
- 601 13. Hamblin S, Hurd PL. Genetic algorithms and non-ESS solutions to game theory models. Anim Behav.
602 2007;74(4):1005–18.

- 603 14. Stockwell D, Peters D. The GARP modelling system: problems and solutions to automated spatial prediction.
604 Int J Geogr Inf Sci. 1999;13(2):143–58.
- 605 15. Barabash Katz G, Benbassat A, Diesendruck L, Sipper M, Henik A. From Size Perception to Counting: An
606 Evolutionary Computation Point of View. In: Proceedings of the 15th annual conference companion on
607 Genetic and evolutionary computation [Internet]. 2013 [cited 2018 Dec 15]. p. 1675–8. Available from:
608 <https://www.cs.bgu.ac.il/~sipper/publications/From Size Perception to Counting.pdf>
- 609 16. Geritz SAH, Metz JAJ, Kisdi É, Meszéna G. Dynamics of adaptation and evolutionary branching. Phys Rev Lett.
610 1997;78(10):2024–7.
- 611 17. Dercole F, Rinaldi S. Analysis of evolutionary processes: The Adaptive Dynamics Approach and Its Applications.
612 Princeton, NJ: Princeton University Press; 2008.
- 613 18. De Visser JAGM, Krug J. Empirical fitness landscapes and the predictability of evolution. Vol. 15, Nature
614 Reviews Genetics. Nature Publishing Group; 2014. p. 480–90.
- 615 19. Kauffman SA. The origins of order. New York: Oxford University Press; 1993.
- 616 20. Gavrillets. S. High-dimensional fitness landscapes and speciation. In: Pigliucci M, Muller GB, editors. Evolution:
617 The Extended Synthesis. MIT Press; 2009. p. 45–80.
- 618 21. Maynard Smith J. Honest signalling: The Philip Sidney game. Anim Behav. 1991;42(6):1034–5.
- 619 22. Hamblin S. On the practical usage of genetic algorithms in ecology and evolution. Methods Ecol Evol.
620 2013;4(2):184–94.
- 621 23. De Jong KA, Sarma J. Generation Gaps Revisited. Found Genet Algorithms [Internet]. 1993 Jan 1 [cited 2018
622 Dec 4];2:19–28. Available from: <https://www.sciencedirect.com/science/article/pii/B9780080948324500076>
- 623 24. Johnstone RA, Grafen A. Dishonesty and the handicap principle. Anim Behav [Internet]. 1993 Oct 1 [cited 2018
624 Dec 4];46(4):759–64. Available from: <https://www.sciencedirect.com/science/article/pii/S0003347283712534>
- 625 25. Hamblin S, Hurd PL. When will evolution lead to deceptive signaling in the Sir Philip Sidney game? Theor Popul
626 Biol [Internet]. 2009;75(2–3):176–82. Available from: <http://dx.doi.org/10.1016/j.tpb.2009.02.002>
- 627 26. Mitchell M. An Introduction to Genetic Algorithms. 6th ed. Cambridge, Massachusetts: MIT Press; 1996. 201 p.

- 628 27. Baker JE. Adaptive Selection Methods for Genetic Algorithms. In: Grefenstette JJ, editor. Proceedings of The
629 First International Conference on Genetic Algorithms and Their Applications. Hillsdale, NJ, USA: Erlbaum; 1985.
630 p. 101–11.
- 631 28. Bergstrom CT, Lachmann M. Signalling among relatives. I. Is costly signalling too costly? Philos Trans R Soc B
632 Biol Sci. 1997;352(1353):609–17.
- 633 29. Bergstrom CT, Lachmann M. Signaling among relatives. III. Talk is cheap. Proc Natl Acad Sci U S A [Internet].
634 1998;95(9):5100–5. Available from:
635 <http://www.ncbi.nlm.nih.gov/pubmed/9560235>
636 <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC20220>
- 637 30. Johnstone RA. Honest signalling, perceptual error and the evolution of ‘all-or-nothing’ displays. Proc R Soc
638 London Ser B Biol Sci [Internet]. 1994 May 23 [cited 2018 Dec 4];256(1346):169–75. Available from:
639 <http://rspb.royalsocietypublishing.org/lookup/doi/10.1098/rspb.1994.0066>
- 640 31. Lachmann M, Bergstrom CT. Signalling among Relatives. Theor Popul Biol. 1998;160:146–60.
- 641 32. Huttegger SM, Zollman KJS. Methodology in biological game theory. Br J Philos Sci. 2013;64(3):637–58.
- 642 33. R Core Team. R: A Language and Environment for Statistical Computing [Internet]. Vienna, Austria: R
643 Foundation for Statistical Computing; 2018. Available from: <https://www.r-project.org/>
- 644 34. Syswerda G. A Study of Reproduction in Generational and Steady-State Genetic Algorithms. Found Genet
645 Algorithms [Internet]. 1991 Jan 1 [cited 2018 Dec 4];1:94–101. Available from:
646 <https://www.sciencedirect.com/science/article/pii/B9780080506845500094>

647

648

649 **Supplementary Information captions**

650 **S1. Derivation of z**

651 **Supporting File 1. Matrix of payoff resulting from pairwise interaction between all possible strategy pairs.**

Proportion of runs converging to the solution

A Truncation

B k-tournament

C Proportional selection

D Linear ranking

E Exponential ranking

Legend

population

trait evaluation

selection

parental pool

reproduction

parental copies

mutation and crossover

offspring

replacement

new population

A	B	C	D	E	F
---	---	---	---	---	---

	Transfer resource on signal	Transfer resource on no-signal	Signal if thirsty	Signal if healthy	Signal if thirsty	Signal if healthy
ESS	1	0	1	0	1	1
ES-1	0	1	0	0	0	0
ES-2	1	1	0	0	0	0