

האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

SEX AND PORTFOLIO INVESTMENT

By

**OMER EDHAN, ZIV HELLMAN,
and DANA SHERILL-ROFE**

Discussion Paper # 683 (April 2015)

מרכז פדרמן לחקר הרציונליות

**THE FEDERMANN CENTER FOR
THE STUDY OF RATIONALITY**

**Feldman Building, Edmond J. Safra Campus,
Jerusalem 91904, Israel**
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

SEX AND PORTFOLIO INVESTMENT

OMER EDHAN, ZIV HELLMAN, AND DANA SHERILL-ROFE

ABSTRACT. We attempt to answer why sex is nearly ubiquitous when asexual reproduction is ostensibly more efficient than sexual reproduction. From the perspective of a genetic allele, each individual bearing that allele is akin to a stock share yielding dividends equal to that individual's number of offspring, and the totality of individuals bearing the allele is its portfolio investment. Alleles compete over portfolio growth, and evolutionary reproduction strategies are essentially on-line learning algorithms seeking improved portfolio growth, with sexual reproduction a goal-directed algorithmic exploration of genotype space by sampling in each generation. The model assumes a stochastically changing environment but not weak selection. We show that in finite population models the algorithm of sexual reproduction yields, with high probability, higher expected growth than the algorithm of asexual reproduction does, proposing this as an explanation to why a majority of species reproduce sexually.

Keywords: Evolution, Sexual Reproduction, Learning Algorithms

1. INTRODUCTION

Consider evolution from the perspective of a single allele. Both competition and co-operation are significant aspects of its evolutionary process: the allele competes with other alleles within the same locus but must co-operate with the alleles at other loci with whom it forms the genotypes in the cells of organisms.

The ideal situation for an allele is to be part of a genotype that is optimally evolutionarily fit. This touches upon one of the mysteries of the success of sexual reproduction as an evolutionary strategy. If a particularly outstanding combination of alleles comes together in a genotype, what is the point of

School of Social Sciences, University of Manchester; Department of Economics, Bar Ilan University; Federmann Center for the Study of Rationality, Hebrew University and the Department of Management, Bar Ilan University.

The authors thank Uzi Motro, Elad Hazan, Aviv Keren and Yannai Gonczarowski for helpful conversations and comments.

dissolving that combination in successive generations by randomly mixing its alleles with those of a mating partner?

We adopt the view here that each individual organism whose genotype carries a copy of the allele is akin to a share of a stock that the allele has purchased. The growth returned by that stock share is measured by the offspring of that individual that bear copies of the allele into the succeeding generation.

If one individual bearing a copy of the allele is a single share of stock, then the totality of individuals in a given generation who bear that allele form the allele's portfolio investment. The allele's portfolio growth is measured as the weighted average of the growth return (i.e., offspring attaining adult reproductive maturity in the next generation), paid by the components of the portfolio, namely, the individuals with a copy of the allele.

In this view the allele represents a long-term investor in portfolios, rebalancing the portfolio weights from one generation to the next. Alleles, like investors, compete with each other, attempting to surpass each other in posting higher portfolio growth returns.

The picture is, however, significantly more complicated in several ways. Alleles play a double role. They are both investors and elements of investment securities, in the sense that for a pair of alleles, a and b , that share a genotype but are located in different loci, a may regard b as being an element in its investment while b regards a as being an element in *its* investment. As the frequency of alleles in one locus changes over time, it affects the portfolios of alleles in other loci in a complex manner.

Furthermore, the portfolio returns of an allele may depend critically on environmental variables that change stochastically over time. An investment in a genotype that yields brilliant returns in one environment can be rendered toxic if the environment changes. Human investors face similar challenges, but they have the advantage of being able to make use of sophisticated investment strategies with the aid of computational devices and probability distribution estimates to select rational investment channels. The building blocks of genetics are, after all, merely macromolecules. It would seem that, at best, the process of evolution and natural selection can apply simple, adaptive strategies.

A series of papers in recent years have applied to the study of the theory of evolution a class of simple, adaptive strategies, namely on-line learning algorithms based on regret minimisation; see, e.g., Kleinberg, Piliouras, and Tardos (2009), Chastain et al. (2013), Chastain et al. (2014), Mehta, Panageas, and Piliouras (2014) and Meir and Parkes (2015). Roughly speaking, regret is the possible increase in payoff one would have received had

one switched to a different mix of actions in the past. In boosting the frequency of alleles that have contributed to genotypes of high fitness over those that have been in low fitness genotypes, natural selection intuitively seems to be implementing a regret-minimising approach.

In line with these results, we show here that in a model assuming an infinite population, the dynamics of sexual reproduction can be recast simultaneously in different loci as a direct application of a variant of a regret-minimising on-line learning algorithm called the Hedge algorithm (see Freund and Schapire (1997)). Theorem 1 then shows that the long run growth rate of a sexually reproducing haploid species converges to the growth rate that would have been attained if at each locus the only active allele had been the one that in hindsight attained the highest portfolio growth. This is similar to a result in Chastain et al. (2013) (where it is couched in terms of mixability, a concept closely related to our portfolio growth), but we do not assume weak selection¹ and moreover, our result holds for any stochastic sequence of viability fitness landscapes, not only for a deterministic, fixed landscape.

We can re-word this as follows: the dynamic of sexual reproduction acts as if its goal is to attain the cumulative population growth rate of an ‘ideal sexual genotype’, namely the genotype which has at each locus the allele with the highest portfolio growth (which can be identified only in hindsight). Sexual reproduction is an algorithm that reliably attains this ideal goal, no matter what succession of fitness landscapes it faces, whether that be a fixed unchanging landscape at all times or an environment with wild swings from one period to the next.

This brings us back to the question of ‘why sex?’. It has long been noted in the literature that sexual reproduction, in contrast to asexual reproduction, entails several significant costs that are high enough to call into question why sexual reproduction is as ubiquitous across species as it is known to be. To list only a few such costs: sexual reproduction disassociates beneficial gene combinations by way of random shuffling between generations; males are costly, as half the population does not bear offspring,²; and sexual reproduction requires risky and energy-inefficient processes and behaviours, including mate searching, competition for mating partners, and co-ordination between mates (a more complete treatment of the subject may

¹ Independently, Meir and Parkes (2015) have also extended Chastain et al. (2013) to situations without weak selection.

² This is true in nearly all sexually reproducing species in the animal world. Strictly speaking, however, it is not necessarily a cost of sex; hermaphrodite species (which do exist) apparently gain whatever benefits sexual reproduction grants without paying the cost of depending on only half the population to produce female gametes.

be found in Otto (2009)). Indeed, explaining the role of sex is considered a major challenge in the study of evolution.

Can on-line algorithms explain the role of sex in evolution? Barton, Novak, and Paixão (2014) suggest a negative answer by pointing out that in principle asexual reproduction can equally well be characterized as an on-line learning algorithm, which thus leaves unclear why an advantage to sexual reproduction should accrue. That intuition is correct, as we show formally in Theorem 2 and Corollary 1, in the context of an infinite population model.

That is, the analysis of the dynamics of asexual reproduction as an on-line learning algorithm actually leads to the conclusion that asexual reproduction is more efficient than sexual reproduction in cumulative growth rates, the reason being that the asexual reproduction algorithm also strives to attain the cumulative growth rate of an ideal genotype: the genotype with globally optimal hindsight growth. By definition, the ideal sexual genotype can at best match the growth of the optimal growth genotype; it can be shown by example that it can also fall short.

Indeed, this is reminiscent of results reported in Livnat et al. (2008), in which computer simulations indicate greater long-run fitness for asexually reproducing populations than comparable sexually reproducing ones, which select for increased mixability rather than increased fitness. This may be considered a formalisation of one of the classic costs of sex: optimal growth genotypes are broken up by sexual reproduction if their individual alleles lack high mixability, whereas they are boosted in asexual reproduction.

This sharpens the question of ‘why sex’. If sexual reproduction bears significant costs *and* yields less efficient long-run growth rates than asexual reproduction, why is sex such an eminently ubiquitous reproductive strategy?

Our answer, as expressed in Theorems 5 and 6, is that in the finite population model the conclusions are reversed. There, an advantage to sexual reproduction emerges, in a sense, from the very finiteness under consideration.

In an infinite population model, it is possible to suppose that nearly every possible genotype is represented in the population. Yet in a realistic finite population model, this is not feasible, as the space of possible genotypes is combinatorially immense. A finite population in each generation can only be a sample of the fictitious infinite population.

An asexually reproducing population essentially samples only once (in between rare positive mutations). Its growth rate can only be as good as the best genotype in the initial sample. A sexually reproducing population, by

contrast, samples a different portion of genotype space in each generation and furthermore does so in a directed manner.

The picture that emerges is of sexual reproduction as a goal-directed on-line learning algorithmic exploration of genotype space by sampling in each generation. This algorithm is extremely reliable in its outcome; in fact, Theorem 3 shows that the sampling algorithm of a sexually reproducing finite population attains exactly the same ideal genotype population growth that the algorithm of a corresponding sexually reproducing infinite population theoretically achieves. Asexual reproduction in a finite population, in contrast, always marches in place in the same small corner of genotype space.

One might use the following analogy.³ The asexual reproductive strategy is like a poker player adopting the strategy of always playing exactly the hand she is initially dealt from the deck, without discarding and replacing any cards. If the hand that she has been dealt is extremely strong, such as a straight flush, then this strategy will look brilliant in hindsight; but the chances of that happening are slim. Most of the time, this strategy will yield poor results (since, for example, the probability of being dealt a no-pair hand in a five-card deal is over 50 percent).

The sexual reproductive strategy is akin to a player who always, in every circumstance, discards and replaces the lowest face-value cards from the hand he is holding. Sometimes this comes at the irrational cost of breaking up a royal flush, but more often than not the player will improve on his initial position.

2. PRELIMINARIES

2.1. Alleles and Genotypes. The basic unit of our model is the allele. That is, we suppose as a primitive the existence of a finite set of *alleles* \mathcal{A} (which may be thought of as the basic ‘gene pool’). We furthermore suppose the existence of a fundamental partition $\mathcal{L} := \{L_1, \dots, L_\ell\}$ of \mathcal{A} , i.e., $\bigcup_j L_j = \mathcal{A}$ and $L_i \cap L_j = \emptyset$ for every pair $L_i, L_j \in \mathcal{L}$ such that $i \neq j$. Each partition element L_j is a *locus* and \mathcal{L} is the set of *loci*, of cardinality ℓ . For ease of notation, denote $I = (1, \dots, \ell)$.

In principle the analysis here could be conducted assuming populations of any n -ploidy for $n \geq 1$. For concreteness and mathematical simplicity we focus on the haploid (i.e., $n = 1$) case. A (haploid) *genotype* g is a selector function $g : \mathcal{L} \rightarrow \mathcal{A}$; i.e., g satisfies the condition that $g(L_j) \in L_j$ for each $j \in I$.

³ The use of poker as an analogy for genotypic fitness also appears in Otto (2009).

More simply, a generic genotype g can be depicted as a string of alleles,

$$a_1 a_2 \dots a_{\ell-1} a_\ell,$$

such that $a_j \in L_j$ for all j .

The set of all genotypes, which we term the *genotype space* and denote by Γ , has some finite cardinality as determined by the cardinalities of \mathcal{A} and \mathcal{L} . We assume a fixed enumeration of Γ such that each distinct possible genotype can be identified by an integer in $1, \dots, m = |\Gamma|$. The models we construct here exclude the possibility of mutations; we can therefore identify the space Γ of all genotypes with a *species*.

An allele $a \in \mathcal{A}$ *participates in* or *is a member of* a genotype $g \in \Gamma$ if $a \in g(L_j)$ for the unique locus L_j such that $a \in L_j$. Denote the set of all genotypes in which a participates by $G_a \subseteq \Gamma$. We sometimes suggestively term G_a the *portfolio* of allele a .

Alleles in separate loci do not compete; to the contrary, we consider them to have the potential of co-operating. Alleles within each locus, however, do compete with each other, hence we are interested in tracking the relative frequency of alleles at each locus. To that end, for each locus $j \in (1, \dots, \ell)$ we assume an enumeration of the alleles in L_j . A generic allele of L_j is denoted here by $a_{j,i}$, with the initial j indicating the locus and i the relative enumeration of the allele within L_j . To reduce the number of indices, allele $a_{j,i}$'s portfolio is denoted by $G_{j,i}$, rather than $G_{a_{j,i}}$. By definition, then, $g \in G_{j,i}$ iff $a_{j,i} \in g(L_j)$.

Note that for each $j \in I$ we have thus defined a partition of the locus L_j ,

$$G_j := \{G_{j,1}, G_{j,2}, \dots, G_{j,|L_j|}\}.$$

Since each genotype g contains some allele $g(L_j)$, it follows that $G_{j,i} \cap G_{j,k} = \emptyset$ if $i \neq k$ and that $\bigcup_{i \in L_j} G_{j,i} = \Gamma$.

From this perspective there are two ways to view the set of genotypes. One way is as the collection Γ , i.e., as a collection of ordered strings of alleles. The other way is via the loci \mathcal{L} and the portfolios; i.e., for each $j \in I$, $\bigcup_{i \in L_j} G_{j,i} = \Gamma$ as a disjoint union.

2.2. Populations and Frequencies. A *population* of a species is a set of individuals, each of whom bears a particular genotype. A finite population of adult individuals of reproductive maturity is represented here as the disjoint union $X = X_1 \cup \dots \cup X_m$, where each X_g is the set of individuals in the population bearing genotype g . The absolute size of a population X is

$$|X| = \sum_{g \in \Gamma} |X_g|.$$

Definition 1. The *genotypic population frequency* (or simply genotypic frequency) of a genotype g is the fractional share that genotype has in the population, denoted by p_g and defined by

$$(1) \quad p_g = \frac{|X_g|}{|X|}.$$

The genotypic population frequency is the m -tuple $p = (p_1, \dots, p_m) \in \Delta^\Gamma$.

◆

3. INFINITE SEXUALLY REPRODUCING POPULATIONS

3.1. Infinite Populations. We adopt the standard fictional construct of an infinite population. As absolute values cannot play a role in such a model, only relative values are used. We continue, for example, to work with a value such as p_g , understanding it as the fractional share of genotype g in the infinite population. In the same vein, one can compare the relative growth rates of two infinite populations; if, for example, the expected number of offspring per individual in population A is 5 in each generation while the expected number of offspring per individual in population B is 3, we may say that population A experiences a higher growth rate than population B , even though both are infinite.

An infinite population is clearly an unrealistic assumption but serves to yield ‘ideal’ results. One of its main advantages is that it enables probabilities in a sense to be treated as actualities, without having to resort to expectations. In Section 6 a finite population model is considered and the results there are compared to those of this and the next section.

3.2. Sexual Reproduction Model. The biological fitness⁴ of an individual in a population is often defined to be the expected number of viable offspring that individual has, i.e., offspring who are born and manage to attain reproductive maturity. We find it more convenient to work with viability, i.e., the expectation of survival from conception to adulthood, as the main environmentally and genotypically dependent variable.

To this end, we work with a sexual reproduction model loosely based on Ewens (2009), but expanded into a two-stage process. In the mating/conception stage, pairs of individuals conceive (haploid) zygotic offspring. After mating, the parents die and their offspring undergo viability selection in the next stage. The offspring that successfully survive viability selection to attain reproductive maturity form the next generation of adults.

⁴ We work here exclusively with genotypic fitness rather than phenotypic fitness, essentially regarding phenotypes as being entirely determined by genotypes.

More formally, we assume a Wright–Haldane model consisting of

- (1) discrete time panmixic reproduction,
- (2) non-overlapping adult generations,
- (3) haploid individuals,
- (4) no mutation,
- (5) and viability selection.

To track dynamic changes in the population, we add a time index. Time proceeds in discrete units throughout the model. The (adult) population at time t is denoted by X^t . The frequency of genotype g at time t is denoted by p_g^t . The m -tuple of the p_g^t 's is then denoted by p^t . It will be assumed that initially, at time $t = 1$, the ideal infinite population contains individuals from each possible genotype, i.e., that $p_g^1 > 0$ for all $g \in \Gamma$.

At each point in time t , each individual in the adult population X^t mates once with another individual in the population, with the identity of that mate being chosen uniformly randomly. For an individual bearing a particular genotype g the probability that s/he will mate with an individual of genotype g_k is given by $p_{g_k}^t$, the frequency of g_k in the population.

Each pair mating uniformly conceives 2ζ zygotes⁵ where ζ is an integer (hence the per capita number of offspring zygotes is ζ). The haploid genotype of each offspring of a mating pair is formed by picking, at each and every locus, an allele from one of the two parent haploid genotypes, independently and with probability one-half each.

Immediately after the mating/conception stage, there is both an adult parent population X^t , with genotypic frequency p^t , and an offspring/zygote population denoted by \tilde{X}^t , with genotypic frequency similarly denoted by $\tilde{p}^t = (\tilde{p}_1^t, \dots, \tilde{p}_n^t)$. The parent population then dies and the offspring population \tilde{X}^t forms the basis for X^{t+1} , the adult population at time $t + 1$, along with its attendant genotypic frequency p^{t+1} .

⁵ The uniformity in the number of conceived zygotes per mating is not as restrictive an assumption as might appear at first glance. Differential numbers of offspring can be accommodated in this model through the viability fitness introduced here in Definition 2. For example, one can model an individual with genotype g as being sterile by setting $v_{g'} = 0$ for all possible offspring of type g' that he or she can conceive. This would mean that formally in the model 2ζ zygotes would be formed from a mating involving a sterile individual but none would attain adult maturity, which from the perspective of the dynamics of the model would be equivalent to having no offspring at all to begin with. Non-uniform mating chances can similarly be accommodated by reducing the fitness v_g of an ‘unattractive’ individual; in other words, we can expand the interpretation of v_g as being the probability that an individual bearing genotype g both attains adult maturity *and* successfully manages to attract a mate.

Not all of the zygotes conceived during mating at time t , however, will necessarily survive to adulthood.

Definition 2. The *viability* of genotype $g \in \Gamma$ at time t , denoted by v_g^t , is the probability that an individual bearing genotype g conceived as a result of the mating stage at time t survives to reproductive maturity at time $t + 1$. The viability values of the genotypes together form an n -dimensional tensor $V^t = \{v_g^t\}_{g \in \Gamma}$, which is the *viability landscape* of the species at time t . \blacklozenge

Regard v_g^t as a time-dependent random variable that cannot deterministically be foreseen. One may consider the underlying randomness to be due to unpredictable environmental variables. We would like to make as few limiting assumptions on the values of v_g^t as possible; in particular, we do not need to assume normality of distributions or even i.i.d. random variables.

We will, however, adopt here the following assumptions:

- (1) The values of v_g^t range as $0 \leq v_g^t \leq 1$;
- (2) For all t , there is at least one $g \in \Gamma$ such that $v_g^t > 0$.
- (3) There is a minimal $v^* > 0$ such that for all t and all g , if $v_g^t \neq 0$ then $v_g^t \geq v^*$.

The first assumption allows for some genetic combinations to be absolutely lethal: for example, it may happen that certain combinations lead to such gross embryonic developmental defects that they necessarily cause fetal or postnatal mortality. On the other hand, it is not the aim of this work to depict the process of extinction of species, hence we require that at each time t the value of $v_g^t > 0$ for at least one genotype g , to rule out a sudden extinction event. But we also need a stronger condition to rule out a ‘sieve’ effect that would kill off each genotype separately over time until none are left. The second assumption suffices to avoid these extreme cases.

The third assumption is mainly needed for technical reasons but it is amenable to a natural interpretation. A genotype with $v_g^t = 0$ has no hope at all to attain maturity. It is essentially conceived with a flaw that guarantees early death; one might as well assume prenatal mortality. Since there are a finite number of possible genotypes, there is some $v^{*t} > 0$ such that for all g at time t , if $v_g^t \neq 0$ then $v_g^t \geq v^{*t}$, i.e., conditional on surviving gestation and being born at time t , each individual has at least some minimal probability v^{*t} of surviving to adulthood. Assumption (3) then states that although the viability landscape can change considerably over time, there is some minimal $v^* > 0$ such that conditional on being born each individual

has at least the probability v^* (which can be an extremely small but non-zero value) of surviving to adulthood, independently of which time period is under consideration.

Definition 3. The *mean viability* at time t is the value

$$(2) \quad \bar{v}^t := \sum_{g \in \Gamma} \tilde{p}_g^t v_g^t.$$

This is a population-level value relevant to the stage at time t between conception and the attainment of reproductive maturity.

As mating produces ζ zygotes per capita, the *population growth rate* at time t is

$$(3) \quad \bar{\rho}^t := \zeta \bar{v}^t,$$

which may also be interpreted as the mean expected number of zygote offspring per capita who survive to maturity. \blacklozenge

3.3. Allelic Frequency Dynamic. Parallel to the genotypic frequency p^t , one can consider the frequency of an allele $a_{j,i}$ at locus j relative to its competitor alleles in the same locus over time.

Definition 4. Given a genotypic frequency $p^t = (p_1^t, \dots, p_n^t)$ at time t , the *allelic frequency* of $a_{j,i}$ at t (in relation to the alleles in locus j) is

$$(4) \quad q_{j,i}^t := \sum_{g \in G_{j,i}} p_g^t.$$

In parallel, define the zygotic allelic frequency of $a_{j,i}$ at t to be $\tilde{q}_{j,i}^t := \sum_{g \in G_{j,i}} \tilde{p}_g^t$. \blacklozenge

Note that $\sum_{i \in L_j} q_{j,i}^t = \sum_{i \in L_j} \sum_{g \in G_{j,i}} p_g^t = \sum_{g \in \Gamma} p_g^t = 1$. Writing

$$q_j^t := (q_{j,1}^t, q_{j,2}^t, \dots, q_{j,|L_j|}^t),$$

we have that $q_j^t \in \Delta^{|L_j|}$.

In a sexually reproducing population, allelic frequency is often a more important variable to track than genotypic frequency, and we regard the allelic frequency as representing the main state of the population as it undergoes dynamic changes. To track changes in both the genotypic and the allelic frequencies from one generation to the next, we rely on our two-stage model of generation formation.

Proposition 1. $\tilde{q}_{j,i}^t = q_{j,i}^t$ for all t and all alleles.

In contrast to the conclusion of Proposition 1, the shuffling inherent in sexual reproduction does change the genotypic population frequency between the adult population and the offspring population. The viability landscape then further recalibrates the genotypic frequencies as the zygotes mature to adulthood.

Proposition 2. *The dynamics of the transition from p^t to p^{t+1} are determined by the viability landscape V^t : \tilde{p}^t is determined by p^t and p^{t+1} is determined for each genotype g by*

$$(5) \quad p_g^{t+1} = \frac{v_g^t \tilde{p}_g^t}{\bar{v}^t},$$

where \bar{v}^t is the mean viability at time t .

3.4. Portfolios and Growth Rates. To track the dynamics of the allelic frequency q^t we need to introduce some more definitions.

If $\tilde{p}^t = \{\tilde{p}_g^t\}_{g \in \Gamma}$ is the offspring genotypic frequency at time t , then for each allele $a_{j,i}$ and genotype g define

$$(6) \quad \tilde{p}_g^t |_{a_{j,i}} := \begin{cases} \frac{\tilde{p}_g^t}{q_{j,i}^t} & \text{if } g \in G_{j,i} \\ 0 & \text{otherwise.} \end{cases}$$

Since \tilde{p}_g^t is the relative weight in the zygote population of individuals bearing genotype g at time t and $q_{j,i}^t$ is the relative weight in that same population of individuals bearing allele $a_{j,i}$, the natural interpretation of $\tilde{p}_g^t |_{a_{j,i}}$ is that it is the population weight of individuals of type g at time t relative to the population *restricted to those bearing allele $a_{j,i}$* .

We have that

$$(7) \quad \sum_{g \in G_{j,i}} \tilde{p}_g^t |_{a_{j,i}} = 1$$

since $q_{j,i}^t = \sum_{g \in G_{j,i}} \tilde{p}_g^t$. Furthermore:

Lemma 1. *For all $g \in G_{j,i}$,*

$$(8) \quad \tilde{p}_g^t |_{a_{j,i}} = \frac{\tilde{p}_g^t}{q_{j,i}^t}$$

Definition 5. The set $\{\tilde{p}_g^t |_{a_{j,i}}\}_{g \in G_{j,i}}$ is the set of *portfolio weights* of allele $a_{j,i}$ at time t . \blacklozenge

Recalling that for an allele $a_{j,i}$ in L_j , the set $G_{j,i}$ was termed the portfolio of $a_{j,i}$, now for each $g \in G_{j,i}$ we interpret $\tilde{p}_g^t |_{a_{j,i}}$ as the weight of element g in the ‘portfolio held by $a_{j,i}$ at time t ’.

Definition 6. Define

$$(9) \quad \varphi_{j,i}^t := \sum_{g \in \Gamma} v_g^t \tilde{p}_g^t |_{a_{j,i}} = \sum_{g \in G_{j,i}} v_g^t \tilde{p}_g^t |_{a_{j,i}}$$

to be allele $a_{j,i}$'s *mixability*.⁶

Then

$$(10) \quad \rho_{j,i}^t := \zeta \varphi_{j,i}^t$$

can be termed allele $a_{j,i}$'s *portfolio growth* at time t . $\rho_{j,i}^t$ measures the allele's growth rate (an alternative name for this might be the *gross portfolio return* of the allele). \blacklozenge

Given the assumptions specified in Section 3.2 on how the values of $\{v_g^t\}$ range, it follows that $0 \leq \varphi_{j,i}^t \leq 1$ for all t and all alleles, while at the same time at each locus j there is at least one allele $a_{j,i}$ such that $\varphi_{j,i}^t > 0$ (and therefore also $\rho_{j,i}^t > 0$) for all t .

Furthermore, if $\varphi_{j,i}^t \neq 0$ for any $a_{j,i}$ then $\varphi_{j,i}^t \geq v^*$. Based on this, add the following notation: for each t , let $\underline{\varphi}_j^t$ be the value of the minimal non-zero mixability amongst the alleles of locus j at time t , and similarly let $\underline{\rho}_j^t$ be the value of the minimal non-zero portfolio growth amongst the alleles of locus j at time t .

There is also a maximal growth rate: since $v_g^t \leq 1$, the maximal mixability is 1. Hence the maximal $\rho_{j,i}^t$ for any allele at any time is simply ζ .

Lemma 2. For any locus L_j ,

$$(11) \quad \sum_{i \in L_j} q_{j,i}^t \varphi_{j,i}^t = \bar{v}^t.$$

Since by Equation (3) the population growth $\bar{\rho}^t$ at time t is $\zeta \bar{v}^t$, Equations (10) and (11) yield that

$$(12) \quad \bar{\rho}^t = \zeta \bar{v}^t = \zeta \sum_{i \in L_j} q_{j,i}^t \varphi_{j,i}^t = \sum_{i \in L_j} q_{j,i}^t \rho_{j,i}^t,$$

i.e., that the population growth rate is the weighted average of the allelic growth rates within any single locus. We can suggestively think of $\bar{\rho}^t$, under this interpretation, as the average 'market rate' in contrast to the particular portfolio growth rate $\rho_{j,i}^t$ of any single allele $a_{j,i}$.

Note that Lemma 2 holds independently at each and every locus. In particular, since the mean viability \bar{v}^t is a population-level value, this means

⁶ The term is taken from a similar concept appearing in Chastain et al. (2013).

that the weighted averages of the portfolio growth rates at each locus, given by $\sum_{i \in L_j} q_{j,i}^t \rho_{j,i}^t$, are identical for all j .

With all this preparation accomplished, we can finally present the allelic frequency dynamic.

Proposition 3. *For each locus L_j and allele $a_{j,i} \in L_j$,*

$$(13) \quad q_{j,i}^{t+1} = q_{j,i}^t \frac{\varphi_{j,i}^t}{\bar{v}^t}.$$

Since $\rho_{j,i}^t = \zeta \varphi_{j,i}^t$ and $\bar{\rho}^t = \zeta \bar{v}^t$, Equation (13) can be re-written as

$$(14) \quad q_{j,i}^{t+1} = q_{j,i}^t \frac{\rho_{j,i}^t}{\bar{\rho}^t}.$$

Note that Equations (13) and (14) have the form of the standard discrete time replicator equation of the evolutionary game theory literature. Note also that if $\varphi_{j,i}^t = 0$ (equivalently, $\rho_{j,i}^t = 0$) for an allele $a_{j,i}$ for any time t then the allele ceases to exist in the population and $q_{j,i}^\tau = 0$ for all time periods $\tau > t$.

4. NO-REGRET SEXUAL REPRODUCTION

We can recast the dynamics of the previous section in both game theoretic terms and as an on-line algorithm. In this we are inspired by ideas appearing in Livnat et al. (2008), Chastain et al. (2013) and Chastain et al. (2014).

4.1. Sexual Reproduction as a Stochastic Coordination Game. Define a strategic-form game \mathcal{G}^t at time t as follows. Suppose that a viability landscape $\{v_g^t\}_{g \in G}$ is given. The players of the game are $(1, \dots, \ell)$, one player for each locus. The action set of player j is the set L_j ; i.e., each allele $a_{j,i}$ of locus j is an action of player j in the game \mathcal{G}^t .

Players may play mixed strategies. A mixed strategy of player j is a probability distribution q_j^t over his action set. Since the action set is the set of alleles of locus j , the probability distribution q_j^t is equivalently an allelic frequency at locus j .

After each player j has revealed his/her mixed strategy q_j^t at time t , the payoff to each player is, uniformly, the population growth rate $\bar{\rho}^t$, which is determined by the mean viability \bar{v}^t as calculated per Equation (2). Since individual payoffs to all players are identical under all conditions, while the payoff matrix is essentially a random variable that is dependent on the stochastic viability landscape sequence $\{V^t\}$, the sequence $\{\mathcal{G}^t\}$ may be considered to be a stochastic coordination game.

Biological Description	Game \mathcal{G}^t Description
Locus j	Player j
Allele $a_{j,i}$	Action i of Player j
Allelic frequency q_j^t at locus j	(Mixed) strategy q_j^t of Player j
Mean fitness \bar{v}^t	Payoff for strategy profile (q_1^t, \dots, q_ℓ^t)

FIGURE 1. Parallels between the biological description of the previous section and the game of this section.

If we postulate that each player j in $\{\mathcal{G}^t\}$ chooses consecutive strategies according to Equation (13), then the players are implementing an uncoupled dynamic, as defined in Hart and Mas-Colell (2003). We may interpret this dynamic as each player regretting his previous time period mixed strategy and updating in a way that would have counter factually increased his payoff. Given the game structure, the effect of each individual player acting in such a way as to increase his payoff is higher payoffs for all players. This holds even if we regard the ‘true underlying dynamic’ as unbridled competition between the alleles at each locus under replicator-equation conditions: the net effect of the dynamic for the species is increased viability, even without any agent seeking to act for the ‘good of the species’.

4.2. Sexual Reproduction as an On-line Learning Algorithm. In addition to the stochastic coordination game description, the dynamics of Equation (13) at each individual locus can also be recast using the following metaphorical description. Imagine that each locus j is an investor who is uncertain as to which investment portfolio to purchase. The investor hires investment experts to provide investment advice; these ‘experts’ are the alleles $(a_{j,1}, \dots, a_{j,|L_j|})$ in L_j .

At time t , each expert $a_{j,i}$ is given the portion $q_{j,i}^t$ of the investor’s capital, which is then invested in a portfolio $G_{j,i}$ (whose weights depend on $\{q_\gamma^t\}_{\gamma \in \Gamma^{-j}}$). The growth of portfolio $G_{j,i}$ is then revealed as $\rho_{j,i}^t = \zeta \varphi_{j,i}^t$ (depending on the revealed viability landscape V^t , per Equation (10)). The overall growth that the investor sees as a result is given, by Lemma 2, as $\bar{\rho}^t = \zeta \bar{v}^t$.

Looking forward myopically, the investor adopts the strategy in period $t+1$ of increasing the portion $q_{j,k}^{t+1}$ of each expert k who produced better returns than the mean in period t while reducing the portion $q_{j,h}^{t+1}$ of each expert h who produced worse returns than the mean in period t . Furthermore, the amount increased (respectively, decreased) should be proportional to the excess (respectively, deficit) from the mean.

Equation (13) attains this. The dynamics can suggestively be considered to be an expression of the investor's 'regret' after each time period t , in the sense that the investor regrets having placed too much weight on experts whose returns were lower than period t 's 'market average return' and too little weight on experts whose returns were above the market average.

This metaphor, of each locus acting as an investor who divides his wealth amongst experts (the alleles) and at each time period 'boosts' the fraction given to the experts whose portfolios achieve above-average returns, is often used as motivation for the Hedge algorithm, an on line learning algorithm. (The Hedge algorithm was originally due to Freund and Schapire (1997). It has since been significantly generalized and is now a major element in the study of on-line learning algorithms. An extensive recent survey of the subject appears in Arora, Hazan, and Kale (2012).) We present a variant of the Hedge algorithm in detail here as Algorithm 1.

Algorithm 1 Exponential Hedge Algorithm (Freund and Schapire (1997))

Parameters:

number of experts k ,
maximal single-period gain M ,
initial weight vector $q^1 \in \Delta^k$,
a real number $\eta > 0$,
number of trials T

for $t = 1, 2, \dots, T$ **do**

1. Receive reward vector $r^t \in [0, M]^k$
2. Score the weighted mean value $y^t = \sum_{i=1}^k q_i^t r_i^t$
3. Update the weight vector by

$$\hat{q}_i^{t+1} = q_i^t \cdot e^{\eta r_i^t}$$

4. Normalize \hat{q}^{t+1} into an allocation $q^{t+1} \in \Delta^k$ by setting

$$q_i^{t+1} = \frac{\hat{q}_i^{t+1}}{\sum_j \hat{q}_j^{t+1}}$$

for each $i \in 1, \dots, k$

end for

The cumulative gain of the algorithm over all the time periods is

$$\Phi_H := \sum_{t=1}^T y^t = \sum_{t=1}^T \sum_{i=1}^k q_i^t r_i^t.$$

One can ask instead what would have been the cumulative gain if instead of following the algorithm the investor had in every time period given all his

wealth to expert i (i.e., always set the weight $w_i^t = 1$ in the algorithm) and nothing to the other experts, in which case the cumulative gain would have been

$$\Phi_i := \sum_{t=1}^T r^i.$$

The regret at not having selected *ex ante* the expert who was revealed *ex post* to be the optimal expert is the difference

$$\max_i \Phi_i - \Phi_H.$$

One can also ask about the average regret, $(\max_i \Phi_i - \Phi_H)/T$, and its asymptotic value as T increases.

Lemma 3.1 and Theorem 3.2 of Auer et al (1995) show that, depending on the specifics of the problem studied, η can be tuned so as to guarantee that the average regret of the Hedge algorithm decreases at the rate $\mathcal{O}(\sqrt{\ln k/T})$, and hence that it asymptotically vanishes. We denote the optimal η attaining this goal by η^* .

We now recast the allelic frequency dynamics at one locus as a Hedge-styled algorithm. The details are presented in Algorithm 2, which is composed in a way that makes direct comparison to Algorithm 1 as transparent as possible.

The vector dynamics determined by Algorithm 2 come down to $q_i^{t+1} = q_i^t \frac{\rho_i^t}{\underline{\rho}^t}$, exactly mimicking those of Equation (14). The gain scored in each generation for each allele i is $\ln(\rho_i^t/\underline{\rho}^t)$, i.e., the log of the quotient of that allele's portfolio growth rate and the worst performing portfolio (this ensures that the quotient is always greater than or equal to 1 and therefore that the log gain is non-negative). The overall gain y^t scored in each generation is the mean log gain, $\sum_i q_i^t \ln(\rho_i^t/\underline{\rho}^t)$, scaled by $\frac{1}{\eta^*}$ (note that, similarly to ρ^t , the mean log gain is a population-level value even though it is calculated at a single locus).

Ostensibly, tracking the cumulative mean log gain may look wrong, as compared to the more standard log mean return, i.e., $\ln(\sum_i q_i^t \rho_i^t)$. There are two justifications for the mean log gain. Firstly, the mean log gain may be considered to be a proxy for measuring growth rapidity at each time period, in the sense that, all other things being equal, a population with a higher mean log return will experience greater absolute size growth than a comparable population with a lower mean log return.

Algorithm 2 Allelic Frequency Algorithm (at a single locus)**Parameters:**

number of alleles k ,
 initial allelic frequency vector $q^1 \in \Delta^k$,
 number of generations T

for $t = 1, 2, \dots, T$ **do**

1. Receive the revealed viability landscape V^t
2. Calculate, for each allele i , the portfolio growth ρ_i^t
3. Calculate the average growth rate $\bar{\rho}^t = \sum_{i=1}^k q_i^t \rho_i^t$ and the minimal non-zero growth rate $\underline{\rho}^t$
4. Score a gain $y^t := \frac{1}{\eta^*} \sum_{i=1}^k q_i^t \ln(\rho_i^t / \bar{\rho}^t)$, with the understanding that the sum does not include indices i such that $\rho_i^t = 0$
5. Update the frequency vector by

$$\hat{q}_i^{t+1} = q_i^t \cdot e^{\eta^* (\frac{1}{\eta^*} \ln(\rho_i^t / \bar{\rho}^t))} = q_i^t (\rho_i^t / \bar{\rho}^t),$$

with the understanding that $\hat{q}_i^{t+1} = 0$ if $\rho_i^t = 0$

6. Normalize \hat{q}^{t+1} into an allocation $q^{t+1} \in \Delta^k$ by setting

$$q_i^{t+1} = \frac{\hat{q}_i^{t+1}}{\sum_j \hat{q}_j^{t+1}} = q_i^t \frac{\rho_i^t}{\bar{\rho}^t}$$

for each $i \in 1, \dots, k$

end for

Secondly and perhaps more convincingly, the result of the Hedge algorithm itself provides a justification ‘in hindsight’. In detail, define the cumulative mean log gain as

$$(15) \quad \Psi_S := \eta^* \sum_{t=1}^T y^t = \sum_{t=1}^T \sum_{i=1}^k q_i^t \ln \rho_i^t$$

and the corresponding cumulative gain in hindsight from sticking to one allele throughout (i.e., always set the weight $q_i^t = 1$ in the algorithm for some particular allele i)

$$(16) \quad \Psi_i := \sum_{t=1}^T \ln \rho_i^t.$$

Furthermore, let the optimal-in-hindsight allele be $i^* := \arg \max_i \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_i$. (Clearly, the optimal allele cannot be one that ever has a portfolio growth rate of zero, even once, because then this last limit goes to zero.)

Then, denoting $\Phi_S := \frac{1}{\eta^*}(\Psi_S - \sum_{t=1}^T \ln \rho^t)$ and $\Phi_i := \frac{1}{\eta^*}(\Psi_i - \sum_{t=1}^T \ln \rho^t)$ we can, as before, inquire about the asymptotic regret between Φ_S and $\max_i \Phi_i$. Making use again of standard results with respect to regret under the Hedge algorithm, one concludes that $\lim_{T \rightarrow \infty} \frac{1}{T} \Phi_S = \lim_{T \rightarrow \infty} \frac{1}{T} \Phi_{i^*}$. This implies that the average cumulative mean log gain, $\frac{1}{T} \Psi_S$, converges to the average log gain of the best allele in hindsight, meaning that we could just as well have been measuring the growth rate $\sum_{t=1}^T \ln \rho_{i^*}^t$ all along.

In the multi-locus setting, we may regard the allelic dynamic described by Equation (14) as implementing Algorithm 2 at every locus simultaneously, in an uncoupled manner. We then understand Ψ_S to refer to the cumulative log growth that results from this dynamic at every locus, and $\Psi_{a_{j,i}}$ to refer to the cumulative log growth that results from when only allele $a_{j,i}$ is present at locus j in the population at every generation while the dynamic at other loci follows Equation (14).

In a similar vein, $\Psi_{a_{j,i}a_{k,m}}$ will mean the cumulative log growth when alleles $a_{j,i}$ and $a_{k,m}$ are respectively fixed at loci j and k , while the other loci follow the dynamic of Equation (14), and so on, until we get to the cumulative log growth attained by supposing that the entire population consists solely of individuals bearing a single genotype g , which we denote by Ψ_g .

Definition 7. The *optimal-in-hindsight allele at locus j* is

$$a_j^* := \arg \max_{a_{j,i} \in L_j} \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{j,i}.$$

The *optimal-in-hindsight sexual genotype* is

$$g^S := a_1^* a_2^* \dots a_{\ell-1}^* a_\ell^*.$$

The term ‘optimal in hindsight’ is used because it is only after the sequence of viability landscapes $\{V^t\}$ has been revealed that the optimal alleles can be identified. \blacklozenge

Theorem 1. *In the infinite population model, the cumulative growth of a sexually reproducing population converges asymptotically to the cumulative growth of a population consisting solely of individuals bearing the optimal-in-hindsight sexual genotype g^S .*

Note that Theorem 1 does *not* state that the population will asymptotically consist solely of individuals bearing genotype g^S ; large stochastic swings in the viability landscape could render this impossible. What is being stated is that the expected cumulative growth rate converges to what *would have* been the growth rate had the population hypothetically consisted only of individuals bearing genotype g^S at each time period.

Theorem 1 is similar in content to Corollary 4 of Chastain et al. (2013), which is also attained using a different no-regret approach (a multiplicative weights update algorithm). There are significant differences between the two results: Chastain et al. (2013) assume weak selection, which is not assumed here (in independent work, Meir and Parkes (2015) have also managed to obtain the result of Chastain et al. (2013) without assuming weak selection), and the dynamic in Chastain et al. (2013) is deterministic, in the sense that the fitness landscape is assumed fixed for all time, while the dynamic here is stochastic, depending on the stochastic revelation of the viability landscape over time.

The statement of Theorem 1 may be considered a formalisation of a long-standing idea in population genetics, as expressed for example in Dawkins (1982): “[T]he sense in which genes may be said to ‘caucus’ and form ‘alliances’ is the following. Selection favours those genes which succeed *in the presence of other genes, which in turn succeed in the presence of them.* Therefore mutually compatible genes arise in gene-pools.”

5. INFINITE ASEXUALLY REPRODUCING POPULATIONS

The dynamics of asexual reproduction are simpler than those of sexual reproduction. In asexual reproduction, the main variable to trace is the population frequency of the genotypes, p^t .

In detail, in the model an individual bearing genotype $g \in \Gamma$ at time t conceives ζ zygotes, each of which also bears the same genotype g . The probability of survival to adulthood of each of these haploid meiotic products of the zygote is given stochastically by v_g^t . As before, the generation at time t dies shortly after reproduction and the offspring born at time t that survive to adulthood form the generation at time $t + 1$.

Denoting again the population frequency of genotype g by p_g^t , the genotypic growth at time t of genotype g is

$$(17) \quad \rho_g^t = \zeta v_g^t.$$

The population growth rate is then

$$(18) \quad \bar{\rho}^t = \sum_{g \in \Gamma} p_g^t \rho_g^t.$$

The population frequency dynamic is a straight-forward replicator:

$$(19) \quad p_g^{t+1} = p_g^t \frac{\rho_g^t}{\bar{\rho}^t} = p_g^t \frac{v_g^t}{\bar{v}^t}.$$

We may regard an asexually reproducing population as implementing a Hedge algorithm, but with significant differences with respect to a sexually

reproducing population. In the asexual setting, not alleles but genotypes directly compete with each other. Furthermore, there is no multi-locus focus in the asexual context, as now the entire population implements the Hedge algorithm.

Algorithm 3 Asexual Reproduction Algorithm

Parameters:

number of genotypes m in genotype space Γ ,
 initial genotypic frequency vector $p^1 \in \Delta^m$,
 number of generations T

for $t = 1, 2, \dots, T$ **do**

1. Receive the revealed viability landscape V^t
2. Calculate, for each genotype g , the genotypic growth ρ_g^t
3. Calculate the average growth rate $\bar{\rho}^t = \sum_{g \in \Gamma} p_g^t \rho_g^t$ and the minimal non-zero growth rate $\underline{\rho}^t$
4. Score a gain $y^t := \frac{1}{\eta^*} \sum_{g \in \Gamma} p_g^t \ln(\rho_g^t / \underline{\rho}^t)$, with the understanding that the sum does not include indices i such that $\rho_i^t = 0$
5. Update the genotypic frequency vector by

$$\hat{p}_g^{t+1} = p_g^t \cdot e^{\eta^* (\frac{1}{\eta^*} \ln(\rho_g^t / \underline{\rho}^t))} = p_g^t \rho_g^t$$

with the understanding that $\hat{p}_g^{t+1} = 0$ if $\rho_g^t = 0$

6. Normalize \hat{p}^{t+1} into an allocation $p^{t+1} \in \Delta^n$ by setting

$$p_g^{t+1} = \frac{\hat{p}_g^{t+1}}{\sum_{g'} \hat{p}_{g'}^{t+1}} = p_g^t \frac{\rho_g^t}{\bar{\rho}^t}$$

for each $g \in \Gamma$

end for

The vector dynamics determined by Algorithm 3 come down to $p_g^{t+1} = p_g^t \frac{\rho_g^t}{\bar{\rho}^t}$, mimicking those of Equation (19). The gain y^t scored in each generation is again the mean log return.

Define the cumulative mean log growth gain as

$$(20) \quad \Psi_A := \eta^* \sum_{t=1}^T y^t = \sum_{t=1}^T \sum_{g \in \Gamma} p_g^t \ln \rho_g^t$$

and the corresponding cumulative gain in hindsight from sticking to one genotype throughout (i.e., always set the weight $p_g^t = 1$ in the algorithm for

one particular $g \in \Gamma$)

$$(21) \quad \Psi_g := \sum_{t=1}^T \ln \rho_g^t.$$

Then, denoting $\Phi_A := \frac{1}{\eta^*} \Psi_A - \sum_{t=1}^T \ln \rho_{-g}^t$ and $\Phi_g := \frac{1}{\eta^*} \Psi_g - \sum_{t=1}^T \ln \rho_{-g}^t$ we can again inquire about the asymptotic regret between Φ_A and $\max_g \Phi_g$.

Definition 8. The *optimal-in-hindsight asexual genotype* is

$$g^A := \arg \max_{g \in \Gamma} \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_g.$$

◆

Theorem 2. *In the infinite population model, the cumulative growth of an asexually reproducing population converges asymptotically to the cumulative growth of a population consisting solely of individuals bearing the optimal-in-hindsight asexual genotype g^A .*

Corollary 1. *In the infinite population model, asymptotic cumulative asexual population growth is greater than or equal to asymptotic cumulative sexual population growth.*

That it is possible for the inequality of Corollary 1 to be strict, i.e., $\Phi_{g^A} > \Phi_{g^S}$, is established by the following example (inspired by an example in Livnat et al. (2008)).

Example. We work with a haploid two-locus model with three alleles per locus. The alleles in locus 1 are denoted by a_1, a_2, a_3 and those in locus 2 by b_1, b_2, b_3 .

The viability landscape is constant for all time. Denoting by v_{kl} the viability of genotype $a_k b_l$, let

$v_{11} = 0.420$	$v_{12} = 0.523$	$v_{13} = 0.550$
$v_{21} = 0.520$	$v_{22} = 0.530$	$v_{23} = 0.500$
$v_{31} = 0.510$	$v_{32} = 0.525$	$v_{33} = 0.410$

Suppose that initially there is a uniform distribution of genotypes. Under asexual reproduction the optimal-in-hindsight asexual genotype $g^A = a_1 b_3$, which consistently has the greatest viability, 0.55.

Under sexual reproduction, consider an initial population with each genotype represented by an equal number of individuals. The alleles with the highest mixability in this example are a_2 and b_2 , and their frequencies are correspondingly boosted in the next generation above those of the other alleles. This results in the increasing of the mixability of these two alleles in the succeeding generation, further boosting their frequencies and so on in

each generation. The asymptotic dynamic yields a population growth rate of the optimal-in-hindsight sexual genotype $g^S = a_2 b_2$, whose 0.53 viability is outpaced by that of g^A . ♦

The intuition behind Corollary 1 is related to the way asexual reproduction and sexual reproduction implement their respective algorithms. Asexual reproduction is akin to a buy-and-hold strategy of investment. The resultant dynamics are a straightforward contest between different genotypes (or stocks, in terms of the metaphor). Genotypes with a better ‘track record’ are boosted, while the below-average ones are penalised.

The sexual reproduction approach, at each locus, is akin to partitioning genotype space into genotypes bearing allele a_1 , genotypes bearing allele a_2 , and so on. The algorithm then boosts the genotypes belonging to partition elements with high *average* performance, while penalising genotypes in partition elements with low average growth. This may lead to situations in which an optimal genotype has an asymptotically low representation in the population because its partition element exhibits consistently poor average growth.

6. THE FINITE POPULATION MODEL

The result of Corollary 1 sharpens the question ‘why sexual reproduction’. It is well known that sexual reproduction as a rule induces significant costs relative to asexual reproduction. If, in addition, sexual reproduction is also less efficient in long-term fitness, one could easily imagine that a majority of biological species would be using the asexual reproductive strategy, and yet exactly the opposite is true.

Corollary 1, however, relates to a model that includes the two assumptions of an infinite population and initial full genotypic population support. We show in this section that removing these assumptions significantly changes the picture.

Even if one interprets the infinite population model as referring to a ‘sufficiently large’ finite population, in most cases only a ‘relatively small’ finite population model can be realistic. For most species, $|\Gamma|$ is a far bigger number than the number of organisms that have ever existed in the species. Every generation can only contain a tiny fraction of the immense collection of all possible genotypes.

6.1. Finite Sexually Reproducing Populations. As in the infinite population model, in the model here random pair-wise mating occurs at each discrete time period t . The adult population is denoted by \hat{X}^t , and the individuals bearing genotype g by \hat{X}_g^t . The total size of the adult population is

denoted by $|\overset{\circ}{X}^t|$. The genotypic population frequency of the adult population at time t is given by $\overset{\circ}{p}_g^t := |\overset{\circ}{X}_g^t|/|\overset{\circ}{X}^t|$.

Again, each mating conceives 2ζ zygotes. The haploid genotype of each offspring of a mating pair is formed by picking, at each and every locus, an allele from one of the two parent haploid genotypes, independently and with probability one-half each. The mating of the adult parent population $\overset{\circ}{X}^t$, with genotypic population frequency $\overset{\circ}{p}^t$ produces an offspring/zygote population $\overset{\circ}{\tilde{X}}^t$ with population genotypic frequency $\overset{\circ}{\tilde{p}}^t$, which forms the basis for the adult population at $t + 1$ after taking into account viability selection as determined by the viability landscape V^t .

However, the genotypic frequency distribution here, $\overset{\circ}{p}^t$, will typically be ‘sparse’ in the finite model. That is, $\overset{\circ}{X}_g^t = 0$ for nearly all $g \in \Gamma$, resulting in $\overset{\circ}{p}_g^t = 0$ for any such ‘un-represented’ genotype. Even when $\overset{\circ}{X}_g^t$ is non-zero, we may presume that it is a small integer, perhaps even 1.

In contrast, we will suppose that the finite-population allelic frequency distribution $\overset{\circ}{q}_j^t$ at each locus j has full support at time $t = 1$. An allele $a_{j,i}$ may disappear in later generations if $\overset{\circ}{p}_{j,i}^t = 0$ at some time t ; by assumption, at each locus j there is at least one allele for which this does not occur.

The allelic frequency dynamic in the finite population model proceeds as in the infinite population model. However, we now attach a significant new interpretation to the portfolio weights, $\{\overset{\circ}{p}_g^t|_{a_{j,i}}\}_{g \in G_{j,i}}$ of a locus j , to the mixability $\overset{\circ}{\varphi}_{j,i}^t$ of an allele $a_{j,i}$ and to its portfolio return $\overset{\circ}{\rho}_{j,i}^t$, motivated by the following metaphor.

If the allelic frequency dynamic cannot be implemented over the entire genotype space Γ , the next best thing is to do what election polling entities do when they cannot poll an entire population, namely, use of a small random sample. In this view, the genotypic population $\overset{\circ}{X}^t$ in each generation of a finite sexually reproducing species is in effect a random sample of the entire genotype space.⁷

We then interpret the portfolio weights $\{\overset{\circ}{p}_g^t|_{a_{j,i}}\}_{g \in G_{j,i}}$ of an allele $a_{j,i}$ to represent a random sample of the portfolio and similarly

$$(22) \quad \overset{\circ}{\varphi}_{j,i}^t := \sum_{g \in G_{j,i}} v_g^t \overset{\circ}{p}_g^t|_{a_{j,i}}$$

⁷ Some care should be taken with regard to this metaphor. In election polls each random sample ideally satisfies the i.i.d. conditions. In contrast, the genotype space sampling of sexual reproduction is not i.i.d. since the probability of the formation of a particular genotype at time t does depend on the genotypic frequency at time $t - 1$.

to be an estimator of the mixability $\varphi_{j,i}^t = \sum_{g \in G_{j,i}} v_g^t \tilde{p}_g^t |_{a_{j,i}}$ that would be attained in an imaginary corresponding infinite population. The same holds for the finite population portfolio growth $\rho_{j,i}^t := \zeta \varphi_{j,i}^t$ of allele $a_{j,i}$, which becomes an estimator of $\rho_{j,i}^t = \zeta \varphi_{j,i}^t$.

As a weighted average estimator using random sampling, $\rho_{j,i}^t$ is an unbiased estimator, that is, $E(\rho_{j,i}^t) = \rho_{j,i}^t$. The error of $\rho_{j,i}^t$ is of order $\mathcal{O}(1/|\hat{X}^t|^{\frac{1}{2}})$.

Proposition 4. *If the Hedge algorithm (Algorithm 1 above) is implemented with an unbiased estimator of the reward vector r^t replacing the true reward vector r^t in each time period, then the long-term weight vector \hat{q}_i^T resulting from using the estimator a.s. converges to the true long-term weight vector q_i^T .*

Theorem 3. *In the finite population model, the cumulative growth of a sexually reproductive population almost surely converges asymptotically to Ψ_S , the cumulative growth of a population consisting solely of individuals bearing the optimal-in-hindsight sexual genotype g^S .*

Theorem 3 shows, perhaps surprisingly, that even without the strong assumptions of an infinite population and full support over all genotypes, by the use of sampling a sexually reproducing population attains the same ideal asymptotic growth rate as in Theorem 1.

As in the infinite population model, this result does not state that the population will asymptotically consist solely of individuals bearing genotype g^S , only that the expected cumulative growth rate converges to what would have been the growth rate had the population hypothetically consisted only of individuals bearing genotype g^S at each time period. However, the result here is even stronger: the growth rate of g^S is attained even if the finite population never actually contains *any* individuals bearing the genotype g^S .

6.2. Finite Asexually Reproducing Populations. The dynamic of the population frequency of a finite asexually reproducing population is exactly the same as that of an infinite population except that the genotypic population frequency vectors p^t are sparse. In detail, an initial population sample of individuals exists at time 1. Denoting by Ω_0 the support of the initial population frequency distribution p^1 , we suppose that $\Omega_0 \subset \Gamma$ and that $|\Omega_0| \ll |\Gamma|$.

An individual bearing genotype $g \in \Omega_0$ at each time t conceives 2ζ zygotes bearing genotype g . The probability of survival to adulthood of such a zygote is v_g^t . The population frequency dynamic is exactly the same as that of Equation (19), namely, $p_g^{t+1} = p_g^t \frac{v_g^t}{\bar{v}^t}$. Significantly, this implies that

the support of p^t for all t will always be a subset of Ω_0 ; i.e., no genotype that is not positively represented in the initial population sample is created.⁸

As before, the dynamic can be considered to be an application of the Hedge algorithm, as in Algorithm 3, with cumulative mean log growth gain $\Psi_A := \sum_{t=1}^T \sum_{g \in \Omega_0} p_g^t \ln \rho_g^t$.

Definition 9. The *optimal-in-hindsight asexual genotype within Ω_0* is

$$g_{\Omega_0}^A := \arg \max_{g \in \Omega_0} \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_g.$$

◆

Theorem 4. *In the finite population model, the cumulative growth of an asexually reproducing population converges asymptotically to $\Psi_A^{\Omega_0}$, the cumulative growth of a population consisting solely of individuals bearing the optimal-in-hindsight asexual genotype $g_{\Omega_0}^A$.*

The proof of Theorem 4 is immediate from the implementation of the algorithm and the restriction of the genotypes to Ω_0 .

6.3. Why Sex.

The previous subsections show that in the finite population model, sexual reproduction reliably attains the asymptotic growth rate Ψ_S of the optimal-in-hindsight sexual genotype. Asexual reproduction in contrast attains the asymptotic growth rate $\Psi_A^{\Omega_0}$ of the optimal-in-hindsight asexual genotype within Ω_0 (as before, we denote by Ψ_g the growth rate of a particular genotype g).

Hence the question of whether sexual reproduction is more advantageous than asexual reproduction or that the opposite holds comes down to whether or not Ψ_S overtakes the expected value of $\Psi_A^{\Omega_0}$. That in turn depends on the expected genotypes included in the initial sampling Ω_0 .

The asexual genotype sampling depends on mutations to provide genotype variation. In other words, exploration of genotype space by asexual reproduction depends on mutations. We have not included mutations formally in the model because they affect both sexual and asexual populations at similar rates. Moreover, beneficial mutations occur rarely over many generations, alongside many deleterious mutations, while reproduction occurs much more rapidly, in each generation by definition. What we are concentrating on is the competition between sexual and asexual reproduction in the relatively lengthy time periods between the appearance of beneficial mutations.

⁸ No mutations are presumed in this model.

The distinct genotypes in Ω_0 are a random sampling with uniform distribution of the genotype space Γ . Denoting the number of distinct genotypes in Ω_0 by k and the size of genotype space by $m = |\Gamma|$, there are $\binom{m}{k}$ different possibilities for Ω_0 (where two initial samples that contain the same number of distinct genotypes are entirely equivalent from the perspective of the asymptotic asexual reproduction growth rate).

For fixed Γ , k and a stochastic sequence of viability landscapes, let r denote the ratio of genotypes $g \in \Gamma$ such that $\Psi_g \leq \Psi_S$, i.e.,

$$(23) \quad r := \frac{|\{g \in \Gamma : \Psi_g \leq \Psi_S\}|}{|\Gamma|}.$$

Then the number of possible samples of size k that contain only genotypes g such that $\Psi_g \leq \Psi_S$ is $\binom{rm}{k}$. All the other samples, $\binom{m}{k} - \binom{rm}{k}$ in number, contain at least one genotype g such that $\Psi_g > \Psi_S$, which is sufficient to give the asexual population an advantage over the sexual population. It follows that if $\binom{m}{k} - \binom{rm}{k} \leq \frac{1}{2} \binom{m}{k}$, equivalently if $\binom{rm}{k} \geq \frac{1}{2} \binom{m}{k}$, then $\Psi_S \geq \mathbb{E}(\Psi_A^{\Omega_0})$, where the expectation is w.r.t. the uniform distribution over all samples of size k .

Definition 10. The minimal r such that $\binom{rm}{k} \geq \frac{1}{2} \binom{m}{k}$, where $m = |\Gamma|$, is the *threshold value for sexual reproduction advantage* over asexual reproduction. Denote this threshold value by r^* .

Whether or not the threshold value is attained depends on the details of the sequence of viability landscapes. We have, however, the following result:

Theorem 5. *As the size of the genotype space $|\Gamma|$ increases the threshold value for sexual reproduction advantage r^* monotonically decreases.*

Theorem 5 shows that as $|\Gamma|$ increases, the required threshold r^* decreases, making it increasingly likely that sexual reproduction will be advantageous.

We present the existence of the threshold r^* for an advantage in expectation for sexual reproduction as one possible answer to the question of ‘why sex’. Furthermore, since the size $|\Gamma|$ increases as the number of loci ℓ increases (along with the number of alleles), we hypothesise Theorem 5 is a possible explanation of the observed fact that sexual reproduction is particularly prevalent in species with longer lengths of genotypes, since the larger $|\Gamma|$ grows, the easier it becomes for sexual reproduction to out-compete asexual reproduction.

6.3.1. *Asymptotic Threshold Value.* We calculate here the value of $r^* = r^*(m, k)$ for a fixed sampling size $k \geq 1$ as $m \rightarrow \infty$. By definition, r^* must satisfy $\binom{rm}{k} \geq \frac{1}{2} \binom{m}{k}$. By Stirling's formula, if n is much larger than k then

$$(24) \quad \ln \binom{n}{k} \approx k \ln \left(\frac{n}{k} - \frac{1}{2} \right) + k - \frac{1}{2} \ln(2\pi k).$$

Fix a large k . Notice that $r^*(m)m \rightarrow \infty$ as $m \rightarrow \infty$, hence by increasing m we may write as $m \rightarrow \infty$

$$(25) \quad -\ln 2 \leq \ln \frac{\binom{r^*m}{k}}{\binom{m}{k}} \approx k \left[\ln \left(\frac{r^*m}{k} - \frac{1}{2} \right) - \ln \left(\frac{m}{k} - \frac{1}{2} \right) \right].$$

After applying some algebraic manipulation and taking into account the minimality of r^* (amongst values of r such that $\binom{rm}{k} \geq \frac{1}{2} \binom{m}{k}$), we have that

$$(26) \quad r^*(m, k) \approx \frac{k}{2m} \left(1 + \left(\frac{m}{k} - \frac{1}{2} \right) e^{-k} \right)$$

as $m \rightarrow \infty$.

6.4. Conditions for Absolute Advantage for Sexual Reproduction.

As noted, the value of r^* in Theorem 5 depends on the sequence of viability landscapes over time; without further assumptions, not much more can be said. In addition, r^* only gives a relative advantage, that is, when the fraction of genotypes g such that $\Psi_g \leq \Psi_S$ is r^* then sexual reproduction overtakes asexual reproduction in expectation.

In this section we present conditions that give sexual reproduction an *absolute* advantage, in the sense that they guarantee that Ψ_S overtakes the ‘‘average’’ asexual growth rate.

As before, genomes contain ℓ loci. Suppose that every allele is chosen with uniform probability, hence we may regard the allele at locus i as a random variable g_i . In this view, a genotype g itself becomes a random variable. The process of viability landscape selection over the time periods is, as before, a stochastic process; we assume that the viability process is independent of the random variable g .

Define a linear order \preceq on \mathbb{Z}_+ to be of *Shapley type* if and only if for every $t \in \mathbb{Z}_+$ and $i \in \{1, \dots, \ell\} \equiv [\ell]$ we have $t\ell \preceq t\ell + i$. For convenience, we denote $[a, b]_{\preceq}$ the interval w.r.t. \preceq . For $n = t\ell + i$, $i \in [\ell]$, consider the random variables

$$(27) \quad R_{\preceq}^n = \mathbb{E}(\ln \rho_g^t | (g_j)_{t\ell+j \in [t\ell+1, n]_{\preceq}})$$

and

$$(28) \quad S_{\preceq}^n = R_{\preceq}^n - R_{\preceq}^{n-1}.$$

R_{\preceq}^n may be thought of as the expected growth rate at time t of a random genotype with loci j , $t\ell + j \in [t\ell + 1, n]_{\preceq}$ being kept fixed. In a sense, this is the market value at time t of a portfolio for which loci j_1, \dots, j_i have already been selected to hold alleles g_{j_1}, \dots, g_{j_i} . S_{\preceq}^n can now be thought of as the selective pressure at time t on g_i conditioned on the other alleles being kept fixed. The environment is assumed to have *asymptotically vanishing selection variation* if and only if for each T

$$(29) \quad \min_{\preceq} \sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2 = o(T^2)$$

as $\ell \rightarrow \infty$, with \preceq being a linear order of Shapley type. We denote by Ψ_g the limit of the time average growth rate for a genotype g and $\Psi_S = \Psi_{g^S}$. The genotype g^S *beats the market* if and only if

$$(30) \quad \mathbb{E}_{g,\rho} \left[\liminf_{T \rightarrow \infty} \left(\Psi_{g^S} - \frac{1}{T} \sum_{t=1}^T \ln \rho_g^t \right) \right] > 0.$$

Theorem 6. *Under asymptotically vanishing selection variation, if g^S beats the market then as the genotype length ℓ grows to ∞ , the likelihood that an asexually reproducing finite population will eventually exhibit greater cumulative growth than the sexually reproducing finite population diminishes to 0.*

Theorem 6 depends on the genotype g^S beating the market, namely on the fact that the inequality

$$(31) \quad \mathbb{E}_{g,\rho} \left[\liminf_{T \rightarrow \infty} \left(\Psi_{g^S} - \frac{1}{T} \sum_{t=1}^T \ln \rho_g^t \right) \right] > 0$$

holds. The following examples show that this condition is indeed satisfied in some general settings.

Example 1. Consider a deterministic viability landscape, with the property that if the initial population is chosen randomly via the uniform distribution at each locus, then for each locus j the initial mixability of the different alleles is not constant. We call this property *uniform $\neg U$* . Thus, for each locus j there is an allele a_j^* whose mixability attains the maximum, and it is strictly larger than the average mixability. Now, repeatedly applying Algorithm 2 at each locus will increase the mixability of the alleles a_j^* for each locus j , and we may conclude that $\Psi_S > \mathbb{E}_g(\Psi_g)$ which implies that g^S beats the market.

Example 2. Consider a viability landscape with a.s. stationary distribution, namely, for each $g \in \Gamma$ the empirical measure of the process V_g converges

a.s. to a distribution on \mathbb{R}^Γ that may depend on the realisation. The growth of genotypes may be recast as a Bayesian game whose underlying state space is the set Ω of realisations of $(V_g)_{g \in \Gamma}$, prior P (which is the process measure), and payoff $\pi(g|\omega) = \int \ln(\zeta V) d\mu_g(V|\omega)$, given the realisation $\omega \in \Omega$, and the stationary distribution $\mu_g(\cdot|\omega)$ on \mathbb{R}^Γ . Next, suppose that $\pi(\cdot|\omega)$ has the uniform $\neg U$ property a.s. (w.r.t. P). Given each realisation $\omega \in \Omega$ of the process, we know by the previous example that

$$(32) \quad \pi(g^{S,\omega}|\omega) > \mathbb{E}_g(\pi(g|\omega)),$$

P -a.s., with $g^{S,\omega}$ being the optimal-in-hindsight sexual genotype for the realization ω . Using our assumption of P -a.s. convergence of empirical measure of the process, we have, assuming the realization has the form $(V_g^t)_{t=1}^\infty$, for each $g \in \Gamma$,

$$(33) \quad \mathbb{E}_g(\pi(g|\omega)) = \mathbb{E}_g \left(\int \ln(\zeta V) d\mu_g(V|\omega) \right) =$$

$$(34) \quad \mathbb{E}_g \left(\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ln(\zeta V_g^t) \right) = \mathbb{E}_g(\Psi_g)$$

P -a.s., and

$$(35) \quad \pi(g^{S,\omega}|\omega) = \int \ln(\zeta V) d\mu_{g^{S,\omega}}(V|\omega) =$$

$$(36) \quad \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ln(\zeta V_{g^{S,\omega}}^t) = \Psi_S$$

P -a.s., and we conclude, by Equation (32), that $\Psi_S > \mathbb{E}_g(\Psi_g)$ whenever the uniform $\neg U$ property holds, namely P -a.s., which implies that g^S beats the market.

6.4.1. Asymptotic Genome Length Calculation. What can be said regarding the threshold ℓ^* such that for each $\ell \geq \ell^*$, sexual reproduction has an absolute advantage? Denote by β the probability of some genotype in a sample Ω_0 out-growing the optimal sexual growth by time T . Denote $k = |\Omega_0|$. Since $P(\ell, T) = e^{-\delta^2 h(\ell)}$ for some function h satisfying $h(\ell) \rightarrow \infty$ as $\ell \rightarrow \infty$, one has that

$$(37) \quad \beta = 1 - (1 - P(\ell, T))^k \approx - (1 - \exp(-\delta^2 h(\ell^*)))^k.$$

Hence, using the approximation $\ln(1+x) \approx x$ one obtains

$$(38) \quad h(\ell^*) \approx -\frac{1}{\delta^2} \ln \left| \frac{\ln(1-\beta)}{k} \right|$$

yielding the threshold value ℓ^* .

6.5. Conclusion.

In the view presented here, evolution via selection applied to reproducing populations is essentially the implementation of on-line learning algorithms to genotype spaces via finite population sampling.

Sexually and asexually reproducing populations use different algorithms, applied to finite and relatively small samples from the available genotype space. The sexual reproduction algorithm, even without mutations, constitutes a directed exploration of genotype space, reliably sampling different samples of genotype space in successive generations to seek out the optimal sexual genotype.

The asexual reproduction algorithm can only identify the optimal genotype in the initial sample with which it starts; it must rely on relatively rare mutations to accomplish exploration of genotype space. Theorems 5 and 6 indicate that the larger the genotype space, the less likely that a small random population sample will contain an optimal genotype. Sexual reproduction, of course, can also use mutations for genotype space exploration, but at the same time it makes use of the time gaps between beneficial mutations for efficient exploration. Furthermore, the asexual population must contain the optimal genotype to attain its optimal asymptotic growth rate, while in contrast the sexual population need not ever actually contain its optimal genotype, even as it reliably attains its optimal growth rate.

Two final comments: The use of game theory and learning algorithms to describe evolutionary processes may be tricky. We do not assume a locus has a will or an agenda, nor is there a master planner optimizing the learning mechanism. This model describes a competition between two different learning algorithms which are a result of the genetically determined reproduction mechanism evolving in each population, to describe the conditions for which each will be favorable. Furthermore, the main viewpoint here is that of the allele, for which a better genotypic environment, and a more efficient learning algorithm, will increase its probability of increasing its copy number, i.e. its fitness.

The results here hold for any initial population, in the sense that given an initial finite population with any particular genotypic composition at time 0, Theorems 5 and 6 give conditions that would make it more efficient for that population to reproduce sexually from that point on. We do not specify by what process that population composition was attained, as it makes no difference for the results. One may, for example, consider a population arriving at a new niche due to migration or environmental changes. The model can also include the case of a population arising from a single cell,

thus almost genetically identical at time 0. As time proceeds, that population may include mutations, but that is not treated in our model. After the occurrence of mutations the population is, strictly speaking, again at time 0 in the model and the analysis proceeds from there; our model captures the time periods between mutations. We leave it to future work to incorporate mutations formally into the model.

7. POSSIBLE DIRECTIONS FOR FURTHER RESEARCH

- (1) Extend the model to diploid organisms.
- (2) Add density-dependent (also known as frequency-dependent) viability selection to the model. In the present model, v_g^t is a function solely of the time t and the identity of the genotype g . There is no dependency of v_g^t on p^t , the relative frequency of the genotypes in the population at time t .
- (3) Formally include the possibility of mutation in the model.
- (4) Extend the model to overlapping generations.
- (5) Extend the model to a continuous-time model.
- (6) Add fertility differences to the model (i.e., take into account gametic fitness). In the present model, each mating produces a uniform expected per capita number ζ of zygotes, independently of the identities of the mating pair.
- (7) Consider what happens if sexual selection is taken into account, i.e., if mating between individuals in the population does not occur under a uniform probability distribution over the population.
- (8) Expanding on the previous point, one might ask why two separate sexes exist even in species where there is no anisogamy (e.g. α and α in baking yeast). The results here indicate that sexual reproduction yields an advantage irrespective of the number of distinct sexes. It might then seem that a population consisting of sexually reproducing hermaphrodites would gain all the benefits of sexual reproduction without paying the well known ‘double cost’ of sex.

8. APPENDIX

8.1. Proofs of Propositions and Theorems.

Proof of Proposition 1. In the model, each mating produces 2ζ offspring. Consider an allele $a_{j,i}$ and a mating individual bearing $a_{j,i}$. In expectation, ζ of that individual’s offspring will also bear $a_{j,i}$. As this is true of every such individual, and $q_{j,i}^t$ is the fraction of the population bearing $a_{j,i}$, it follows that the fraction of offspring bearing $a_{j,i}$ will also be $q_{j,i}^t$, i.e., that $\tilde{q}_{j,i}^t = q_{j,i}^t$. In other words, alleles may be shuffled around into different

genetic combinations during mating, but they maintain their frequencies relative to each other at each locus. ■

Proof of Proposition 2. First consider how \tilde{p}^t depends on p^t . We illustrate the dependence here for the special case of a three-loci species; the more general case can readily be accomplished but rapidly becomes a messy complication of indices.

Let $\bar{g} = a_{k_1} a_{k_2} a_{k_3}$ be a specific ‘target’ genotype. Consider a mating adult bearing genotype $g = a_{i_1} a_{i_2} a_{i_3}$ at time t mating with another adult with genotype $g' = a_{j_1} a_{j_2} a_{j_3}$. If $a_{k_1} \notin \{a_{i_1}, a_{j_1}\}$, $a_{k_2} \notin \{a_{i_2}, a_{j_2}\}$ or $a_{k_3} \notin \{a_{i_3}, a_{j_3}\}$, then there is no possibility for the mating of g and g' to produce an offspring with genotype \bar{g} . Hence in that case define $\beta_{g,g'}^{\bar{g}} = 0$.

If $a_{i_1} = a_{j_1} = a_{k_1}$, $a_{i_2} = a_{k_2} \neq a_{j_2}$, and $a_{j_3} = a_{k_3} \neq a_{i_3}$, then an offspring of a mating of g and g' will definitely contain allele a_{k_1} at the first locus, will contain allele a_{k_2} at the second locus with probability $1/2$, and will contain allele a_{k_3} at the third locus with probability $1/2$. Hence the probability of such a mating producing an offspring with genotype \bar{g} is $1/4$, leading to $\beta_{g,g'}^{\bar{g}} = 1/4$ in this case.

The general pattern for the values of $\beta_{g,g'}^{\bar{g}}$ should be clear from this. Since a single individual of type g at time t randomly encounters with one of type g' with probability $p_{g'}^t$, the total probability that such an individual will produce an offspring of type \bar{g} is $\sum_{g' \in \Gamma} \beta_{g,g'}^{\bar{g}} p_{g'}^t$. Since the fraction of individuals of type g is p_g^t , the probability that an offspring of type \bar{g} is produced by such individuals is $\sum_{g' \in \Gamma} \beta_{g,g'}^{\bar{g}} p_g^t p_{g'}^t$. Summing over all the genotypes leads to

$$\tilde{p}_{\bar{g}}^t = \sum_{g \in \Gamma} \sum_{g' \in \Gamma} \beta_{g,g'}^{\bar{g}} p_g^t p_{g'}^t.$$

Next we temporarily switch to considering a large finite population reproducing under these conditions. If $|X^t|$ is the absolute size of the mating parent generation at time t , the expected absolute size their offspring zygote population bearing genotype g is $\tilde{p}_g^t \zeta |X^t|$. Taking into account the viability v_g^t , which is the probability of the survival of an individual bearing genotype g to reproductive maturity, the expected size of the surviving population at time $t + 1$ bearing genotype g is then $v_g^t \tilde{p}_g^t \zeta |X^t|$.

The size of the total surviving population at $t + 1$ is $\sum_{g' \in \Gamma} v_{g'}^t \tilde{p}_{g'}^t \zeta |X^t|$. Hence

$$(39) \quad p_g^{t+1} = \frac{v_g^t \tilde{p}_g^t \zeta |X^t|}{\sum_{g' \in \Gamma} v_{g'}^t \tilde{p}_{g'}^t \zeta |X^t|} = \frac{v_g^t \tilde{p}_g^t}{\bar{v}^t}$$

using \bar{v}^t as defined in Equation (2).

Since the r.h.s. of Equation (39) does not involve absolute population size, it holds as the size of the population increases to infinity, hence it holds in the infinite population model. ■

Proof of Lemma 1. This follows straightforwardly from the statement of Proposition 1, which enables replacing $\tilde{q}_{j,i}^t$ by $q_{j,i}^t$ in Equation 6. ■

Proof of Lemma 2.

$$\begin{aligned}
\sum_{i \in L_j} q_{j,i}^t \varphi_{j,i}^t &= \sum_{i \in L_j} q_{j,i}^t \sum_{g \in \Gamma} v_g^t \tilde{p}_g^t |_{a_{j,i}} \\
&= \sum_{i \in L_j} q_{j,i}^t \sum_{g \in G_{j,i}} \frac{\tilde{p}_g^t}{q_{j,i}^t} v_g^t \\
&= \sum_{i \in L_j} \sum_{g \in G_{j,i}} \tilde{p}_g^t v_g^t \\
&= \sum_{g \in \Gamma} \tilde{p}_g^t v_g^t \\
&= \bar{v}^t.
\end{aligned}$$

■

Proof of Proposition 3. Recall that

$$\varphi_{j,i}^t = \sum_{g \in \Gamma} v_g^t \tilde{p}_g^t |_{a_{j,i}} = \sum_{g \in G_{j,i}} \frac{v_g^t \tilde{p}_g^t}{q_{j,i}^t}$$

making use of Equations (8) and (9).

From this deduce

$$q_{j,i}^t \varphi_{j,i}^t = \sum_{g \in G_{j,i}} v_g^t \tilde{p}_g^t,$$

hence

$$\begin{aligned}
q_{j,i}^t \frac{\varphi_{j,i}^t}{\bar{v}^t} &= \sum_{g \in G_{j,i}} \frac{v_g^t \tilde{p}_g^t}{\bar{v}^t} \\
&= \sum_{g \in G_{j,i}} p_g^{t+1} \\
&= q^{t+1}
\end{aligned}$$

making use of Equations (4) and (5). ■

Proof of Theorem 1. As Algorithm 2, an instantiation of the Hedge algorithm, is being implemented at each locus, the asymptotic average regret between the cumulative logarithmic population growth achieved under Φ_S and the corresponding growth attained in the counter-factual situation in

which at locus 1 the optimal-in-hindsight allele $a_{1,i}$ alone is fixed while Algorithm 2 is implemented at all other loci, $\lim_{T \rightarrow \infty} \frac{1}{T}(\Phi_S - \Phi_{1,i})$, vanishes.

This means that asymptotically in expectation

$$\frac{1}{\eta^*} \sum_{t=1}^T \sum_{i=1}^{|L_1|} q_i^t \ln \rho_{1,i}^t \xrightarrow{T \rightarrow \infty} \max_{a_{1,i} \in L_1} \frac{1}{\eta^*} \sum_{t=1}^T \ln \rho_{1,i}^t$$

which in turn implies that

$$\sum_{t=1}^T \sum_{i=1}^{|L_1|} q_i^t \ln \rho_{1,i}^t \xrightarrow{T \rightarrow \infty} \max_{a_{1,i} \in L_1} \sum_{t=1}^T \ln \rho_{1,i}^t,$$

hence $\lim_{T \rightarrow \infty} \frac{1}{T} \Psi_S = \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{a_1^*}$.

Next we suppose that a_1^* is fixed at locus 1 while the other loci follow Algorithm 2. By the same reasoning, this time concentrating on fixing the optimal allele at locus 2, we can conclude that $\lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{a_1^*} = \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{a_1^* a_2^*}$, and hence $\lim_{T \rightarrow \infty} \frac{1}{T} \Psi_S = \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{a_1^* a_2^*}$.

Continuing inductively to all loci, the conclusion is that $\lim_{T \rightarrow \infty} \frac{1}{T} \Psi_S = \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{a_1^* a_2^* \dots a_\ell^*} = \lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{g^S}$. ■

Proof of Theorem 2. The proof follows the same reasoning as that of Theorem 1, namely applying the no-regret result of the Hedge algorithm represented by Algorithm 3, but it is even simpler here since there is no need to keep track of multiple loci. The result is straightforward by construction of the algorithm. ■

Proof of Corollary 1. Given Theorems 1 and 2, the statement here is tantamount to showing that

$$(40) \quad \Phi_{g^A} \geq \Phi_{g^S},$$

where g^A is the optimal-in-hindsight asexual genotype and g^S is the optimal-in-hindsight sexual genotype.

The weak inequality in Equation (40) follows trivially from the fact that g^A is by definition an optimal growth genotype in Γ , hence its cumulative growth must be greater than or equal to that of g^S . ■

Proof of Proposition 4. The presentation above of Algorithm 1, for expositional clarity, divides the updating of the weight vector q^t into two steps, first calculating $\widehat{q}_i^{t+1} = q_i^t \cdot e^{\eta r_i^t}$, and then normalising into a distribution by setting $q_i^{t+1} = \frac{\widehat{q}_i^{t+1}}{\sum_j \widehat{q}_j^{t+1}}$. This can be reduced to one step by multiplying each r_i^t by a factor β_t that ensures that $q_i^{t+1} = q_i^t e^{\eta \beta_t r_i^t}$. We will suppose here w.l.o.g. that r^t already incorporates this normalisation factor and therefore that $q_i^{t+1} = q_i^t e^{\eta r_i^t}$ for each t .

It follows that the long-term weight vector element q_i^T is determined as

$$(41) \quad q_i^T = q_i^1 \cdot e^{\eta \sum_{t=1}^T r_i^t}.$$

If instead of implementing the algorithm with the true reward vector r^t one inputs \hat{r}_i^t an estimator for r_i^t at each time t (and the process starts from the initial distribution q^1), then the long-term weight vector element \hat{q}_i^T is determined as

$$(42) \quad \hat{q}_i^T = q_i^1 \cdot e^{\eta \sum_{t=1}^T \hat{r}_i^t}.$$

Hence the long-term difference between q_i^T and \hat{q}_i^T as T increases comes down to the difference between $\sum_{t=1}^T r_i^t$ and $\sum_{t=1}^T \hat{r}_i^t$.

Define a sequence of random variables $\{Y^\tau\}_{\tau \geq 1}$ by

$$(43) \quad Y^\tau = \sum_{t=1}^{\tau} r_i^t - E\left(\sum_{t=1}^{\tau} \hat{r}_i^t\right).$$

Since \hat{r}_i^t is an unbiased estimator of r_i^t and $r_i^t \in [0, 1]$, the sequence $\{Y^\tau\}$ is a martingale with bounded differences.

We can therefore apply the Azuma–Hoeffding inequality to conclude that the long-term expectation $E\left(\sum_{t=1}^T \hat{r}_i^t\right)$ a.s. converges to $\sum_{t=1}^T r_i^t$, and in expectation \hat{q}_i^T a.s. converges to q_i^T . ■

Proof of Theorem 3. By Theorem 1, the infinite population attains asymptotic cumulative growth equal to that of g^S , and does so by implementing the Hedge algorithm at each locus. The finite population also implements the Hedge algorithm, by using an unbiased estimator for the portfolio growth of each allele.

By Proposition 4, the finite population long-term allelic frequency that results from using estimators almost surely converges to the infinite population allelic frequency, with very rapid convergence as guaranteed by the Azuma–Hoeffding inequality. The result follows. ■

Proof of Theorem 5. It is implicit in the statement of Theorem 5 and the definition immediately preceding it that r is such that $rm \in \mathbb{N}$. We will find it more convenient to let $rm \in \mathbb{R}_+$ in general, and then return to the special case of $rm \in \mathbb{N}$.

However, the notation for combinations, namely $\binom{m}{k}$, is defined only for $m, k \in \mathbb{N}$. For the purpose of this proof, we will abuse this notation and extend it in the following way. Let $r \in \mathbb{R}_+$, while $m, k \in \mathbb{N}$. Define

$$(44) \quad \binom{rm}{k} := \frac{rm(rm-1)(rm-2)\dots(rm-k+1)}{k!}$$

which coincides with the standard notation when $rm \in \mathbb{Z}$.

Fix $r \in (0, 1)$, $k \in \mathbb{N}$ and $m \in \mathbb{N}$, with $k \ll m$. Suppose that $\binom{rm}{k} = \alpha \binom{m}{k}$, for some α , equivalently that $\binom{rm}{k} / \binom{m}{k} = \alpha$. By Equation (44), this becomes

$$(45) \quad \binom{rm}{k} / \binom{m}{k} = \frac{rm(rm-1)(rm-2)\dots(rm-k+1)}{m(m-1)(m-2)\dots(m-k+1)} = \alpha.$$

We next want to consider what happens when we increase m by one, to $m+1$, and to compare

$$(46) \quad \frac{rm(rm-1)(rm-2)\dots(rm-k+1)}{m(m-1)(m-2)\dots(m-k+1)}$$

to the value of $\binom{r(m+1)}{k} / \binom{m+1}{k}$, which is

$$(47) \quad \frac{r(m+1)(r(m+1)-1)\dots(r(m+1)-k+2)}{(m+1)m(m-1)\dots(m-k+2)}.$$

Compare these expression by expression. The first expression in (46) is $rm/m = r$; the first expression in (47) is $r(m+1)/(m+1) = r$. However, once we move to the second expression there is a difference.

Specifically, we want to compare $\frac{rm-1}{m-1}$ to $\frac{r(m+1)-1}{m} = \frac{rm-1+r}{m}$. This in turn depends on comparing

$$(48) \quad m(rm-1) \text{ versus } (m-1)(rm-1+r)$$

$$(49) \quad rm^2 - m \text{ versus } rm^2 - rm - m + 1 + rm - r$$

$$(50) \quad rm^2 - m \text{ versus } rm^2 - m + (1-r),$$

which, by dint of $r \in (0, 1)$, yields the conclusion that the right-hand side is greater than the left-hand side. From this one concludes that

$$\frac{rm-1}{m-1} < \frac{r(m+1)-1}{m}.$$

A similar calculation holds for comparing each of the successive terms respectively in (46) and (47).

The conclusion is that if $\binom{rm}{k} / \binom{m}{k} = \alpha$ then $\binom{r(m+1)}{k} / \binom{m+1}{k} > \alpha$. From this result, the statement of Theorem 5 follows immediately: if the goal is attaining $\binom{rm}{k} \geq \frac{1}{2} \binom{m}{k}$ then the greater m is, the less demanding it is to attain the threshold value r^* . ■

Proof of Theorem 6. Let $g \in \Gamma$ be chosen randomly using the uniform distribution over Γ and consider realizations $\rho = (\rho^1, \dots, \rho^T, \dots) \in (\mathbb{R}_{++}^\Gamma)^\infty$. Naturally, $g \mapsto \rho_g^t$ is a random variable for each $g \in \Gamma$ and t , thus each ρ^t may be considered to be a random variable. By abusing notation, we denote

by Ψ_S the limit $\lim_{T \rightarrow \infty} \frac{1}{T} \Psi_{g^S}$. Denote $\Psi_S - \mathbb{E}_{g,\rho} \left[\frac{1}{T} \sum_{t=1}^T \ln \rho_g^t \right] = \delta_T$. Hence

$$(51) \quad \liminf_{T \rightarrow \infty} \delta_T \geq \mathbb{E}_{g,\rho} \left[\liminf_{T \rightarrow \infty} \left(\Psi_S - \frac{1}{T} \sum_{t=1}^T \ln \rho_g^t \right) \right] > 0,$$

with the inequality following from Fatou's Lemma. Hence, for sufficiently large T we have $\delta_T \geq \delta > 0$ for some fixed $\delta > 0$. Fix a linear order \preceq of Shapley type on \mathbb{Z}_+ . Recall that, given $n = t\ell + i$, we defined

$$(52) \quad X_{\preceq}^n = \mathbb{E} \left(\sum_{t'=1}^t \ln \rho_{g'}^{t'} | (g_j)_{t\ell+j \in [t\ell+1, n]_{\preceq}} \right),$$

along with $R_{\preceq}^n = \mathbb{E}(\ln \rho_g^t | (g_j)_{t\ell+j \in [t\ell+1, n]_{\preceq}})$ and $S_{\preceq}^n = R_{\preceq}^n - R_{\preceq}^{n-1}$. We may think of S_{\preceq}^n as the marginal selective pressure at time t on the random allele g_i within genotypes containing alleles g_j with $t\ell + j \in [t\ell + 1, n]_{\preceq}$.

For sufficiently large T and $s > 0$ we have

$$(53) \quad P(\ell, T) \equiv \mathbb{P} \left(\frac{1}{T} \sum_{t=1}^T \ln \rho_g^t - \Psi_S \geq 0 \right) \leq$$

$$(54) \quad \mathbb{P} \left(e^{s \left(\sum_{t=1}^T \ln \rho_g^t - \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{g,\rho}[\ln \rho_g^t] \right)} \geq e^{\delta T s} \right),$$

where the inequality in Equality (53) following from the inequality in Equation (51).

Consider the σ -fields \mathcal{F}_{\preceq}^n generated, for $n = t\ell + i$, by the random variables g_j for $t\ell + j \in [t\ell + 1, n]_{\preceq}$, and ρ^1, \dots, ρ^t . The martingales $(S_{\preceq}^n)_{n=1}^{T\ell}$ together with the filtration $(\mathcal{F}_{\preceq}^n)_{n=1}^{T\ell}$ define a martingale decomposition of

$X_{\preceq}^{T\ell}$. Thus

$$(55) \quad \mathbb{P}\left(e^{s\left(\sum_{t=1}^T \ln \rho_g^t - \sum_{t=1}^T \mathbb{E}_{g,\rho}[\ln \rho_g^t]\right)} \geq e^{\delta T s}\right) \leq$$

$$(56) \quad e^{-\delta T s} \mathbb{E} \left(e^{s\left(\sum_{t=1}^T \ln \rho_g^t - \mathbb{E}_{g,\rho}(\ln \rho_g^t)\right)} \right) =$$

$$(57) \quad e^{-\delta T s} \mathbb{E} \left(\mathbb{E} \left(e^{s\left(\sum_{t=1}^T \ln \rho_g^t - \mathbb{E}_{g,\rho}(\ln \rho_g^t)\right)} \mid \mathcal{F}_{\preceq}^{T\ell-1} \right) \right) =$$

$$(58) \quad e^{-\delta T s} \mathbb{E} \left(\mathbb{E} \left(e^{s(X_{\preceq}^{T\ell-1} - X_{\preceq}^0 + S_{\preceq}^{T\ell})} \mid \mathcal{F}_{\preceq}^{T\ell-1} \right) \right) =$$

$$(59) \quad e^{-\delta T s} \mathbb{E} \left(e^{s(X_{\preceq}^{T\ell-1} - X_{\preceq}^0)} \mathbb{E} \left(e^{s S_{\preceq}^{T\ell}} \mid \mathcal{F}_{\preceq}^{T\ell-1} \right) \right) \leq$$

$$(60) \quad e^{-\delta T s + \frac{1}{2} s \|S_{\preceq}^{T\ell}\|_{\infty}^2} \mathbb{E} \left(e^{s(X_{\preceq}^{T\ell-1} - X_{\preceq}^0)} \right).$$

Repeating the calculation recursively, we obtain

$$(61) \quad \mathbb{P}\left(e^{s\left(\sum_{t=1}^T \ln \rho_g^t - \sum_{t=1}^T \mathbb{E}_{g,\rho}[\ln \rho_g^t]\right)} \geq e^{\delta T s}\right) \leq e^{-\delta T s + \frac{1}{2} s^2 \sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2}.$$

We may choose s such that the probability is optimized, which happens for $s^* = \frac{2\delta T}{\sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2}$ hence

$$(62) \quad \mathbb{P}\left(e^{s\left(\sum_{t=1}^T \ln \rho_g^t - \sum_{t=1}^T \mathbb{E}_{g,\rho}[\ln \rho_g^t]\right)} \geq e^{\delta T s}\right) \leq e^{-\delta^2 T^2 / (2 \sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2)}.$$

Optimizing the choice of the Shapley type linear order \preceq , we in fact have the following

$$(63) \quad \mathbb{P}\left(\sum_{t=1}^T \ln \rho_g^t - \sum_{t=1}^T \mathbb{E}_{g,\rho}[\ln \rho_g^t] \geq \delta T\right) \leq e^{-\delta^2 T^2 / (2 \min_{\preceq} \sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2)}.$$

By asymptotic weak selection, for any fixed large enough T , $\min_{\preceq} \sum_{n=1}^{T\ell} \|S_{\preceq}^n\|_{\infty}^2 = o(T^2)$ as $\ell \rightarrow \infty$. By combining that with the inequality in Equation (63) we finally deduce that for sufficiently large T

$$(64) \quad \lim_{\ell \rightarrow \infty} P(\ell, T) = 0.$$

One can ask, if one independently chooses a bounded (w.r.t. ℓ) random sample Ω_0 , what is the probability that it hits the set of genotypes (and viability landscapes) for which the population carrying genotype g grows at

a greater average rate than that of g^S in the limit? It is $1 - (1 - P(\ell, T))^{|\Omega_0|}$, which diminishes to 0 as $\ell \rightarrow \infty$. ■

8.2. Reproductive Fitness. Fitness is most often interpreted as meaning the reproductive fitness of a genotype, i.e., the expected number of offspring reaching reproductive maturity, as opposed to the viability fitness that is the focus of the model here. We can relate our viability fitness to a measure of reproductive fitness.

Let $g_a = a_1 a_2 \dots a_\ell$ and $g_b = b_1 b_2 \dots b_\ell$ be a pair of genotypes. Denote by $C(g_a, g_b)$ the collection of all 2^ℓ genotypes $g : \mathcal{L} \rightarrow A$ such that for each j , $g(L_j) = a_j$ or $g(L_j) = b_j$. Then the expected number of viable offspring produced by a mating between two individuals bearing genotypes g_a and g_b is $\sum_{g \in C(g_a, g_b)} v_g \frac{\zeta}{2^\ell}$.

From this, we define the fitness of a genotype g_h at time t to be the expected number of viable offspring produced by that genotype under uniformly random probability of mating with any other possible genotype to be:

$$(65) \quad w_{g_h}^t := \frac{\zeta}{2^\ell} \sum_{g \in \Gamma} p_g^t \left(\sum_{k \in C(g_h, g)} v_k^t \right).$$

The fitness values of all the genotypes together form an n -dimensional tensor $W^t = \{w_g^t\}_{g \in \Gamma}$, which is the fitness landscape of the species at time t . One can use this to define the mean fitness

$$(66) \quad \bar{w}^t := \sum_{g \in \Gamma} w_g^t p_g^t$$

If the total population at time t is $|X^t|$, then at time $t+1$ the total population is $|X^{t+1}| = |X^t| \bar{w}^t$.

REFERENCES

- Arora, S., E. Hazan, and S. Kale (2012), The Multiplicative Weights Update Method: A Meta-Algorithm and Applications, *Theory of Computing*, 8, 121–164.
- Auer, P., N. Cesa-Bianchi, Y. Freund, and R. Schapire (1995), Gambling in a Rigged Casino: The Adversarial Multi-Armed Bandit Problem, *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 322–331.
- Barton, N., S. Novak, and T. Paixão (2014), Diverse Forms of Selection in Evolution and Computer Science. *Proceedings of the National Academy of Sciences of the United States of America*, 111 (29), 10398–10399.

- Chastain, E., A. Livnat, C. Papadimitriou, and U. Vazirani (2013), Multiplicative Updates in Coordination Games and the Theory of Evolution, *Proceedings of ITCS 2013*.
- Chastain, E., A. Livnat, C. Papadimitriou, and U. Vazirani (2014), Algorithms, Games, and Evolution, *Proceedings of the National Academy of Sciences of the United States of America*, 111 (29), 10620–10623.
- Dawkins, R. (1982), *The Extended Genotype*. Oxford University Press, Oxford.
- Ewens, J. (2009), Mathematical Population Genetics: Introduction to the Stochastic Theory, *Lecture Notes, Guanajuato, March 2009*.
- Freund, Y., and R. Schapire (1997), A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting, *Journal of Computer and System Sciences*, 55, 119–139.
- Hart, S., and A. Mas-Colell (2003), Uncoupled Dynamics Do Not Lead to Nash Equilibrium, *American Economic Review*, 93 (5), 1830–1836.
- Kleinberg, R., G. Piliouras, and E. Tardos (2009), Multiplicative Updates Outperform Generic No-Regret Learning in Congestion Games, *ACM Symposium on the Theory of Computing, 2009*.
- Livnat, A., C. Papadimitriou, J. Dushoff, and M. Feldman (2008), A Mixability Theory for the Role of Sex in Evolution., *Proceedings of the National Academy of Sciences of the United States of America*, 105 (50), 19803–19808.
- Mehta, R., I. Panageas, and G. Piliouras (2014), Natural Selection as an Inhibitor of Genetic Diversity, *working paper*.
- Meir, R., and D. Parkes (2015), On Sex, Evolution, and the Multiplicative Weights Update Algorithm, *working paper*.
- Otto, S. (2009), The Evolutionary Enigma of Sex, *American Naturalist*, 174, S1–S14.