

Supporting Information

Chastain et al. 10.1073/pnas.1406556111

SI Text

1. Overview

In this document, we introduce the detailed modeling assumptions and dynamical model we use for recombination under selection (section 2); introduce the experts problem from computer science, and the multiplicative weights update algorithm (MWUA) used to solve it (section 3); point out that the MWUA can be interpreted as an algorithm which at each step optimizes a convex combination of the cumulative sum of gains and of the distribution's entropy (section 4); and prove the equivalence of the dynamics under weak selection to the MWUA in *Theorem 4* with $\varepsilon = s$, the selection strength (section 5).

2. Population Genetics Under Selection and Recombination

We make several (more or less standard) simplifying assumptions, which are generally trusted not to change substantially the essence of the evolutionary dynamics. The population of genotypes is infinite. We assume that the genotypes are haploid (contain only one copy of each gene), and that the organisms mate at random to produce a new generation; further, we assume there is no overlap between generations (as if all mating happens simultaneously and soon before death). Each offspring's genome is formed by picking, for each gene, an allele from one of the two parent genomes, independently and with probability half each.

Our exposition will be for the case of two genes with m and n alleles, respectively, even though our results can be easily seen to extend to any number of genes. Thus, genotypes are pairs ij . Each genotype ij has a fitness value w_{ij} which is the expected number of offspring the genotype produces (by mating randomly). The matrix $W = [w_{ij}]$, often called the fitness landscape of the species, entails the basic genetic parameters of the species (it is a k -dimensional tensor for k genes).

We shall be interested in the statistics of the genotypes in the population. The frequency of the genotype ij will be denoted p_{ij} . The matrix of the p_{ij} s is the state of the dynamical system we shall follow. We denote the value of p_{ij} in generation t by p_{ij}^t .

How do the p_{ij}^t s change from one generation to the next? Each pair of genotypes mates with a probability determined by the frequencies of those genotypes and recombines with probability $r \in [0, 1/2]$ to produce an offspring, which is then selected. Accordingly, the expected frequency of genotype ij at the next generation p_{ij}^{t+1} can be written

$$p_{ij}^{t+1} = \frac{w_{ij}}{\bar{w}_t} \left((1-r)p_{ij}^t + r \sum_l p_{il}^t \sum_k p_{kj}^t \right)$$

where \bar{w}_t is the sum of the numerators for all ij , so that frequencies add up to 1 (1).

Wright Manifold, Weak Selection, and Nagylaki's Theorem

Besides the p_{ij} frequencies, one has the marginal frequencies, one for each allele: $x_i = \sum_j p_{ij}$ and $y_j = \sum_i p_{ij}$. Within the simplex of the p_{ij} s, of particular interest to us is the Wright manifold on which p_{ij} is a product distribution (the matrix p_{ij} has rank 1): $p_{ij} = x_i \cdot y_j$. It turns out that, on the Wright manifold, the population genetic equations take a much simpler form, expressed in terms of the marginal probabilities x_i and y_j (*Lemma 2*).

Life, in general, does not reside on the Wright manifold—that is to say, genotype frequencies do not in general have rank 1. This is called linkage disequilibrium, and is measured by the distance from

the Wright manifold $D_{ij} = p_{ij} - x_i \cdot y_j$. Intuitively, it comes about because differences in the fitness of genotypes distort the allele statistics; just imagine two alleles of two genes whose combination is deleterious. By definition, D_{ij} is zero on the Wright manifold.

Weak selection is an important point of view on evolution, which postulates that the entries of the tensor W are all very close to one another relative to recombination. Differences in fitness are minuscule, and the w_{ij} s all lie within the interval $[1-s, 1+s]$ for some very small $s > 0$ which we call the selection strength.

There is an important connection between the Wright manifold and weak selection, best articulated through Nagylaki's theorem. Consider the evolution of genome frequencies p_{ij}^t (or for more than two genes) in a situation in which the fitness values are within $[1-s, 1+s]$ for some tiny $s > 0$ —that is, weak selection prevails. Consider also the corresponding time series of linkage disequilibria $D_{ij}^t = p_{ij}^t - x_i \cdot y_j$.

Theorem 1 [Nagylaki (2, 3)]. (1) For any $t \geq t_0 = 3 \log(1/s)$ and any i, j , $D_{ij}^t = O(s)$; and furthermore (2) for $t \geq t_0$ there is a corresponding process $\{\hat{p}_{ij}^t\}$ on the Wright manifold such that (a) $|\hat{p}_{ij}^t - p_{ij}^t| = O(s)$; and (b) both processes converge and there is one-to-one correspondence between the equilibria of \hat{p}_{ij}^t and the equilibria of p_{ij}^t .

Nagylaki's theorem states essentially that, to understand a genotype frequency process in the weak selection regime, one can instead follow a closely related process on the Wright manifold. As we shall see next, it turns out that this brings about some unexpected connections.

3. Experts Problem

We now discuss a seemingly completely unrelated problem from computer science, and an important algorithm used to solve it.

Imagine that every day you receive advice from n financial experts, and then you must select one of them and follow his advice for that day. Following the advice of expert i in day t results in a net gain (or loss) of g_i^t , a number between -1 and 1 . The g_i^t s are arbitrary numbers in this range, and are not known in advance. This process is repeated for a large number T of days. In the end of the T days, the optimum expert is the one with the largest cumulative gain $G_i = \sum_{t=1}^T g_i^t$; let i^* be this expert, and G^* be this maximum cumulative gain. We wish to come up with an algorithm—possibly randomized—for selecting an expert on each day so that in the end of the T days our total gain is in expectation very close to G^* . In other words, we want to achieve, in the end of the T days, a performance very close to the performance of the expert who is best in retrospect, even though the g_i^t s are unknown and arbitrary—for example, they could be chosen by an adversary striving to deteriorate the performance of the algorithm.

This ambitious goal (which, some would argue, seems intuitively impossible to achieve) can be attained by a very simple method called MWUA. This method was first discovered by the economist Hannan in connection with repeated games (4), then rediscovered by Cover in relation to portfolio analysis (5); later, it was used in artificial intelligence under the name “Boosting” (6), and earlier in a version called “Winnow” (7), until it was recognized as the common idea underlying several simple and curiously effective optimization algorithms developed by computer scientists to solve linear and convex programming problems and network congestion problems, among many others, and codified as MWUA (8).

The MWUA assigns each day t weights, or probabilities, $p_i^t > 0$ to the experts, and each day selects an expert at random among the n with these probabilities. Initially all weights are, say, equal

$p_i^1 = 1/n$ for all i . Then each day, the weight of the i th expert is modified as follows:

$$p_i^{t+1} = \frac{1}{Z^{t+1}} p_i^t (1 + \epsilon g_i^t) \quad (\text{MWUA}),$$

where the normalization $Z^{t+1} = \sum_{i=1}^n p_i^t (1 + \epsilon g_i^t)$ keeps the weights probabilities, and $\epsilon > 0$ is a small constant chosen to balance long-term risk with short-term gains (in the experts problem, a good choice of ϵ turns out to be $\sqrt{(\ln n)/T}$). That is, the probability of selecting expert i is “boosted” by a small multiple of the expert’s gain that day (decreased if that gain is negative).

The following result now captures the surprisingly favorable performance of this simple algorithm:

Theorem 2. *The total gain achieved by the MWUA is in expectation at least $(1 - \epsilon) \cdot G^* - (\ln n)/\epsilon$.*

To see how favorable the performance of MWUA is as stated by this result, notice that it comes ϵ close to the optimum, minus a quantity that does not depend on T . To put it differently, if we choose $\epsilon = \sqrt{(\ln n)/T}$, on an average day this algorithm does only $O(1/\sqrt{T})$ worse than the ex post facto best expert.

4. Optimization Interpretation of MWUA

In this section we point out that the MWUA can be thought of as a multistep optimization algorithm which, at each step, strives to optimize a convex combination of (i) the expected cumulative gain; and (ii) the entropy of the experts’ distribution. This interpretation is implicit in the literature (9, 10).

Let $G_i^t = \sum_{\tau=1}^t g_i^\tau$ be the cumulative gain of expert i in the first t days; thus, $G^t = \max\{G_1^t, \dots, G_n^t\}$. Suppose now that at step t we wish to choose the probabilities p_i^t so as to maximize the sum of two quantities: The expected cumulative gain so far, and (some positive constant $1/\epsilon$ times) the entropy of the distribution p_i^t ’s, $-\sum_{i=1}^n p_i^t \ln p_i^t$. That is,

$$\max_{\{p_i^t\}} \left(\sum_{i=1}^n p_i^t G_i^t - \frac{1}{\epsilon} \sum_{i=1}^n p_i^t \ln p_i^t \right),$$

subject to $\sum_i p_i^t = 1$. Now this is a strictly convex optimization problem, as one term is linear and the other strictly concave, and thus it has a unique optimum, which can be found through the Karush–Kuhn–Tucker conditions (11):

$$G_i^t - (1 + \ln p_i^t) + \mu^t = 0, i = 1, \dots, n,$$

where μ^t is the Lagrange multiplier corresponding to the equality constraint. Similarly, we can write the same equation for the next generation, with t replaced by $t + 1$:

$$G_i^{t+1} - (1 + \ln p_i^{t+1}) + \mu^{t+1} = 0, i = 1, \dots, n.$$

Subtracting these two equations and solving for p_i^{t+1} , and recalling that $G_i^{t+1} - G_i^t = g_i^{t+1}$, we obtain precisely equation (MWUA), where the normalization $Z^{t+1} = \exp(\mu^{t+1} - \mu^t)$.

5. Coordination Games Between Genes

We now introduce the basic formalism of game theory. In a game, each of finitely many players has a set of strategies, and a payoff function mapping the Cartesian product of the strategy sets to the reals. A game in which all payoff functions are identical is called a coordination game. In a coordination game the interests of all players are perfectly aligned, and, intuitively, they all strive to hit the same high value of the common payoff function. In terms of equilibrium calculations, they are trivial.

Fix a game, and a mixed strategy profile, that is, for each player p a distribution x^p over her strategies. For each player p and each

strategy $a \in S_p$, one can calculate the expected payoff of this strategy, call it $q(a)$. How does the player’s strategy change in time? One possibility is inspired by the MWUA of the previous section. The multiplicative weight update dynamics of the game transforms the mixed strategy profile $\{x^p\}$ as follows: For each player p and each strategy $a \in S_p$, the probability $x^p(a)$ of player p playing a becomes

$$\frac{x^p(a) \cdot (1 + \epsilon \cdot q(a))}{1 + \epsilon \cdot \sum_{b \in S_p} x^p(b) q(b)} = \frac{x^p(a) \cdot (1 + \epsilon \cdot q(a))}{1 + \epsilon \cdot \bar{q}},$$

where by \bar{q} we denote the expected payoff to p (in a coordination game, to all players). That is, the probability of playing a is boosted by an amount proportional to its expected payoff, and then renormalized. It is known that two players following the multiplicative update dynamics attain the Nash equilibrium in zero-sum games (this has been rediscovered many times over the past 50 years; see for example ref. 8), but not in general games. It follows directly from the results below that it also converges to the Nash equilibrium in coordination games. Beyond games and portfolio management, the multiplicative updates dynamics lies at the foundations of a very simple, intuitive, robust, and powerful algorithmic idea of very broad applicability (8).

Going back now to population genetics dynamics, let w_{ij} be a fitness landscape (matrix for two genes, tensor for more) in the weak selection regime, that is, each entry is in the interval $[1 - s, 1 + s]$. Define the differential fitness landscape to be the tensor with entries $\Delta_{ij} = (w_{ij} - 1)/s$.

We next point out a useful way to express the important analytical simplification afforded by the Wright manifold:

Lemma 3. *On the Wright manifold, the population genetics dynamics becomes*

$$p_{ij}^{t+1} = \frac{1}{\bar{w}_t} x_i^t \cdot y_j^t \cdot w_{ij},$$

and similarly for more genes.

Proof: As is shown in ref. 1, we can rewrite the population genetics dynamics as

$$p_{ij}^{t+1} = \frac{1}{\bar{w}_t} w_{ij} (p_{ij}^t - r D_{ij}^t),$$

where $D_{ij}^t = p_{ij}^t - x_i^t y_j^t$ is the linkage disequilibrium. Now, because $D_{ij} = 0$, we have

$$p_{ij}^{t+1} = \frac{1}{\bar{w}_t} p_{ij}^t w_{ij}.$$

Finally, because at $D_{ij}^t = 0$, $p_{ij}^t = x_i^t y_j^t$. The result follows.

We are now ready for the main result of this section:

Theorem 4. *Under weak selection with selection strength s , the population genetic dynamics is precisely the multiplicative update dynamics of a coordination game whose payoff matrix is the differential fitness landscape and $\epsilon = s$.*

Proof: We only show the derivation for two genes, the general case being a straightforward generalization.

$$\begin{aligned} x_i^{t+1} &= \sum_j p_{ij}^{t+1} = \frac{1}{\bar{w}_t} \sum_j x_i^t y_j^t w_{ij} = \frac{x_i^t}{\bar{w}_t} \left(1 + s \sum_j y_j^t \Delta_{ij} \right) \\ &= \frac{x_i^t \cdot (1 + s \sum_j y_j^t \Delta_{ij})}{1 + s \cdot \bar{\Delta}}. \end{aligned}$$

Here the first equation is the definition of marginal frequencies, the second is the lemma, the third uses the definition of Δ_{ij} , and the last one follows from the expectation of w_{ij} being 1

plus s times the expectation of Δ_{ij} . The last expression is precisely the multiplicative update dynamics, completing the proof.

Finally, we can connect our result to the optimization interpretation of MWUA:

Corollary 5. *Under weak selection with selection strength s , the population genetics dynamics is tantamount to each gene optimizing at generation t a quantity equal to the cumulative expected fit-*

ness over all generations up to t , plus $1/s$ times the negative entropy of the allele distribution of the gene at time t .

One interpretation is this: If the optimization of cumulative expected fitness is sought, then it makes sense at each step, and in view of the uncertainties of future steps, to balance off cumulative expected fitness so far against the distribution's entropy (a well-known measure of dispersion).

1. Livnat A, Papadimitriou C, Dushoff J, Feldman MW (2008) A mixability theory for the role of sex in evolution. *Proc Natl Acad Sci USA* 105(50):19803–19808.
2. Nagylaki T (1993) The evolution of multilocus systems under weak selection. *Genetics* 134(2):627–647.
3. Nagylaki T, Hofbauer J, Brunovsky P (1999) Convergence of multilocus systems under weak epistasis or weak selection. *J Math Biol* 38(2):103–133.
4. Hannan J (1957) Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, eds Dresher M, Tucker AW, Wolfe P (Princeton University Press, Princeton), Vol 3, pp 97–139.
5. Thomas M (1991) Cover. Universal portfolios. *Math Finance* 1(1):1–29.
6. Freund Y, Schapire RE (1995) A decision-theoretic generalization of on-line learning and an application to boosting. *Computational Learning Theory* (Springer, Berlin), pp 23–37.
7. Littlestone N (1988) Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Mach Learn* 2(4):285–318.
8. Arora S, Hazan E, Kale S (2012) The multiplicative weights update method: A meta-algorithm and applications. *Theory Comput* 8(1):121–164.
9. Mahoney M, Orecchia L (2011) Implementing regularization implicitly via approximate eigenvector computation. *Proc 28th Intl Conf Machine Learning* (Morgan Kaufmann, San Mateo, CA), pp 121–128.
10. Tsuda K, Rätsch G, Manfred K (2005) Warmuth. Matrix exponentiated gradient updates for on-line learning and Bregman projections. *J Mach Learn Res* 6: 995–1018.
11. Boyd SP, Vandenberghe L (2004) *Convex Optimization* (Cambridge Univ Press, Cambridge, UK).