Multiplicative Updates in Coordination Games and the Theory of Evolution

21st Century Algorithms, Project Report

Benedikt Bitterli and Zephyr Lucas

November 12, 2018

1 Introduction

> Sex is the queen of problems in evolutionary biology. Perhaps no other natural phenomenon has aroused so much interest; certainly none has sowed as much confusion.
>
> *Graham Bell [1]*

One of the most interesting problems remaining in evolutionary biology today is the role of sex. The Theory of Evolution, originally proposed by Darwin and widely accepted today, predicts that species evolve gradually over multiple generations, and that through the process of natural selection the combination of genes with the highest *fitness* eventually dominate the population.

It is then perplexing that sexual reproduction, i.e. the combination of genomes of two organisms, is so ubiquitous in nature, when it seems to be disadvantageous at every level compared to asexual reproduction, i.e. the replication of the genome of a single organism. Sex requires biological and social processes that are costly and energy-inefficient. Also, producing male and female organisms means that on average only half of all offspring is able to produce more offspring themselves, compared to asexual production, in which all organisms are capable of reproduction. Additionally, combining genomes means that particularly successful gene combinations are broken down, and result in offspring that exhibit lower fitness than their parents.

Breaking down highly favorable gene combinations has been one of the major roadblocks in understanding the role of sex, since it impedes the evolution of the population towards high fitness. A recent paper [4] in evolutionary biology proposes a potential solution to this problem by challenging the assumption that sex must maximize fitness; instead, the authors conjecture that sex optimizes for the *mixability* of genes, i.e. it selects genes that combine well with many others. Conversely, genes that work extremely well in one particular combination

but poorly in all others are negatively selected and ultimately disappear from the gene pool. Although the authors propose a mathematical model for their theory, due to the complexity of the equations they are only able to provide empirical proof using simulations that the theory is correct.

The paper we picked builds on this work, and provides a mathematical proof that mixability is indeed the optimality criterion of sex. Beyond that, it shows that evolution under sex is equivalent to a *cooperation game* between genes, where individual genes represent the players that try to optimize for the average fitness of the population.

Finally, the paper also touches on the problem of diversity. In the simplified theoretical setup considered in this paper, the evolutionary process eventually converges to an equilibrium state. It would be disappointing (and indeed, counter to our observations in nature) if this model of evolution converged to only a few genetic combinations rather than a diverse pool of individuals. To address this, the authors manage to show that there exist a substantial number of equilibria in which the population is diverse, by solving a related problem on stochastic matrices.

The rest of this report is structured as follows: In Section 2, we introduce the theoretical background needed to understand the theory; in Section 3, we explain the main theorems of the paper; and in Section 4, we explain our approach to simulating (and, hence, empirically verifying) the results of the paper. Section 5 concludes the report.

2 Background

The field of *population genetics* is concerned with the distribution of genes in a population over time. For our purposes, a *gene* represents a particular trait of an individual, and *alleles* encode different expressions of that gene. For example, we might have a gene that controls the eye color of a person, and there might be alleles encoding blue, green or brown eye colors. A *genotype* represents one particular combination of alleles for all the genes of an individual. In more abstract terms, we can think of a gene as a set $S_i = \{a_1, \ldots, a_n\}$ of alleles a_j, and a genotype g as an element of the cartesian product $S_1 \times \cdots \times S_k$ for a species with k genes.

In this definition of a genotype, an individual carries exactly one allele for each gene; this is called a *haploid* genotype. However, most species in nature carry two alleles for each gene (one from each parent); the genotype is *diploid*. Models for diploid genotypes are generally more cumbersome and less amenable to analysis, and this paper focuses on the haploid case.

In line with most papers in the field, this paper focuses on the case when $k = 2$, i.e. a species with two genes. This is purely for notational convenience, as this allows us to express all necessary terms as matrices rather than k-dimensional tensors; the results hold in the general case also. A genotype can then be written as ij, where $i \in [1, \ldots, n]$ is the allele of the first gene, and $j \in [1, \ldots, m]$ is the allele of the second gene.

An important quantity in population dynamics is the *fitness* of an individual. In practice, fitness depends on a variety of factors and can change over time. However, as a simplification, the paper assumes that fitness depends purely on the genotype and is fixed with respect to time. The fitness of an individual with genotype ij is then w_{ij}, and represents the expected number of offspring this individual produces. We collect the fitness values of all possible genotypes in the *fitness matrix* W.

We are interested in how the frequency p_{ij}^t of genotypes in the population evolves over time. p_{ij}^t represents the probability of encountering a particular genotype when picking an individual at random from the population, and we have $\sum_{ij} p_{ij}^t = 1$. Of additional interest is $\bar{w}^t = \sum_{ij} p_{ij}^t w_{ij}$, the "average fitness" of the population.

In addition to the genotype frequency, we can derive the marginals (i.e. the allele frequencies) $x_i^t = \sum_j p_{ij}^t$ and $y_j^t = \sum_i p_{ij}^t$.

2.1 Population Dyamics

In order to reason about the evolution of genotype frequencies p_{ij}^t over time, we require a model of reproduction. Several such models exist; for example, the "selection without recombination" (S) model corresponds to asexual reproduction, in which no new genotypes are introduced in the population:

$$p_{ij}^{S,t+1} = \frac{p_{ij}^t w_{ij}}{\bar{w}^t} \tag{1}$$

$p_{ij}^t w_{ij}$ is the expected number of offspring produced by individuals with genotype ij at time t; dividing by the total number of offspring produced ($\bar{w}^t$) then yields a proper probability distribution at time $t+1$. We can easily see that over time, this reproduction model tends towards a point distribution on the genotype ij with the highest fitness, as we would expect from asexual production.

The reproduction model considered in this paper is called "selection before recombination" (or SR), in the sense that the number of offspring an individual produces depends on the fitness of the individual. In contrast, in the "recombination before selection" (RS) model, all individuals produce an equal amount of offspring, and the survival of the offspring is dependent on their fitness. The difference between the models lies only in when selection is applied - before or after reproduction.

In the SR model, the expected number of descendants of individuals with genotype ij is $p_{ij}^t w_{ij}$. Individuals will mate randomly with other individuals, subject to their frequency in the population and their fitness scores. That is, the probability of choosing to mate with an individual with genotype kl is $p_{kl}^t w_{kl}/\bar{w}^t$; and therefore, the expected number of offspring with parents ij and kl is $p_{ij}^t w_{ij} \cdot p_{kl}^t w_{kl}/\bar{w}^t$.

An individual with genotype ij must have a parent with genotype iX and a parent with genotype Xj. Summing over all possible such combinations and normalizing by the total number of offspring produced yields the population dynamics of the SR model:

$$p_{ij}^{SR,t+1} = \frac{\sum_k \sum_l p_{ik}^t w_{ik} \cdot p_{lj} w_{lj}}{\bar{w}^t \cdot \bar{w}^t} \tag{2}$$

Other papers in this domain typically use a linear mixture of asexual and sexual reproduction models with the "recombination factor" r:

$$p^{t+1} = p^{S,t+1} \cdot (1-r) + p^{SR,t+1} \cdot r \tag{3}$$

This more faithfully captures organisms such as bacteria or certain flora, which are capable of both sexual and asexual reproduction. However, this paper assumes that $r = 1$.

2.2 Wright Manifold and Weak Selection

Theoretical analysis of population dynamics in the general case can be difficult, and this paper relies on a number of assumptions to derive the main results. The main assumption made is that evolution operates in the regime of *weak selection*; that is, the entries of the fitness matrix W are all in range $[1-s, 1+s]$ for some "small" $s > 0$, which is called the *selection strength*.

This assumes that differences in fitness across individuals is small. This is quite a restrictive assumption, but the authors justify this choice by referencing the "neutral theory" of evolution [3]. This theory states that at the molecular level, most mutations are "neutral" in the sense that they do not affect an individual's overall fitness. This assumption comes with a few caveats, and we will revisit the consequences of this assumption later in this report.

Assuming weak selection allows the authors to invoke a powerful theorem by Nagylaki [6] (eq. 53), which shows that the genotype frequency of a system in weak selection is equivalent in the limit to the genotype frequency of a system on the *Wright manifold*. The Wright manifold describes all population frequencies that can be written as an outer product, i.e. $p_{ij} = x_i \cdot y_j$, where x and y are the marginals of the genotype frequency. This is the main reason why the authors restrict themselves to weak selection: Population dynamics on the Wright manifold are much easier to reason about, and using Nagylaki's theorem allows us to generalize proofs for systems on the Wright manifold to systems in weak selection.

3 Analysis

With the notation established, we are now ready to look at the core statements of the paper.

3.1 Evolution and Coordination Games

In a two-player coordination game, both players have a set of strategies available that they can play at each step, where the set of strategies might be different for each player. If the first player uses strategy i, and the second player uses strategy j, they both receive a payoff of Δ_{ij}. Because payoffs are shared between players, they must coordinate; this is in contrast to e.g. zero-sum games, where one player's loss is the other player's gain.

The first player will pick their strategy from a *mixed strategy* x^t, which is a probability distribution over the set of actions at round t of the game; analogously for the second player and y^t. The *expected payoff* of strategy i for the first player is $\sum_j y_j \Delta_{ij}$, and the *total expected payoff* $\bar{\Delta}^t$ is trivially $\bar{\Delta}^t = \sum_{ij} x_i y_j \Delta_{ij}$. We can easily verify that the total expected payoff is identical for both players.

How should the mixed strategy profile evolve over time? Consider now a *multiplicative update* scheme: The ith entry of the mixed strategy profile is multiplied by a factor that depends on the expected payoff of playing that strategy. This gives the update rule

$$x_i^{t+1} = \frac{1}{Z^t} \cdot x_i^t \left(1 + \epsilon \cdot \sum_j y_j^t \Delta_{ij}\right) \tag{4}$$

where

$$Z^t = \sum_i x_i^t \left(1 + \epsilon \cdot \sum_j y_j^t \Delta_{ij}\right) = 1 + \epsilon \sum_i \sum_j x_i^t y_j^t \Delta_{ij} = 1 + \epsilon \bar{\Delta}^t \tag{5}$$

is the normalization factor to bring x^{t+1} back to the probability simplex, and ϵ is the learning rate. Inserting into (4) yields

$$x_i^{t+1} = \frac{1}{1 + \epsilon \bar{\Delta}^t} \cdot x_i^t \left(1 + \epsilon \cdot \sum_j y_j^t \Delta_{ij}\right) \tag{6}$$

An analogous equation exists for the second player and y.

So far, we've only utilized standard game theory and multiplicative updates. We will now return to population dynamics, and consider the marginal allele distribution x^{t+1} for a population on the Wright manifold (i.e. $p_{ij} = x_i y_j$):

$$\begin{aligned}
x_i^{t+1} &= \sum_j p_{ij}^{t+1} \qquad \text{(from definition of the marginal)} && (7)\\
&= \frac{1}{\bar{w}^t \cdot \bar{w}^t} \sum_j \sum_k \sum_l p_{ik}^t w_{ik} \cdot p_{lj}^t w_{lj} \qquad \text{(from (2))} && (8)\\
&= \frac{1}{\bar{w}^t \cdot \bar{w}^t} \sum_k p_{ik}^t w_{ik} \left(\sum_l \sum_j p_{lj}^t w_{lj}\right) && (9)\\
&= \frac{1}{\bar{w}^t} \sum_k p_{ik}^t w_{ik} && (10)\\
&= \frac{1}{\bar{w}^t} x_i^t \sum_k y_k^t w_{ik} \qquad \text{(assuming Wright manifold)} && (11)
\end{aligned}$$

Because we are operating in weak selection, we know that $w_{ij} \in [1 - s, s + s]$. Consider now $\Delta_{ij} = (w_{ij} - 1)/s$, which rescales the fitness values to be in $[-1, 1]$. We can easily see that $w_{ij} = 1 + s\Delta_{ij}$ and $\bar{w}^t = 1 + s\bar{\Delta}^t$. Insert into Eq. (11) and obtain

$$x_i^{t+1} = \frac{1}{1 + s\bar{\Delta}^t} \cdot x_i^t \left(1 + s \sum_k y_k^t \Delta_{ik}\right) \tag{12}$$

This equation is identical to Eq. (6)! This means that we can interpret population dynamics under SR-rules and weak selection as a coordination game where the genes are the players. Each gene can play an allele as a strategy, and their total expected payoff is equivalent to the average fitness of the population. The marginal allele frequencies are the mixed strategy, and the learning rate is s, the selection strength.

In the paper, the authors use a different proof based on a transformation in a paper by Livnat et al. [4]. The proof is somewhat opaque, and the transformation does not appear to be in the paper they reference (this is possibly a typo). However, we can derive the proof much more easily using algebraic manipulation, as shown above.

3.2 Diminishing Regret

Establishing a link between population dynamics and MWU allows the authors to show two things: First, that evolution in the weak selection regime converges; and second, that evolution maximizes the *mixability* of a gene, which we will define shortly.

The authors use a theorem by Kale [2] for both proofs. However, we already analyzed a very similar form of MWU in class for the experts problem, and we can draw similarities to that analysis here for a better intuition. Pattern matching to the notation in class, we can see that the MWU step for x^t is the same as the experts problem with a loss of

$$l_i(t) = -\sum_j y_j^t \Delta_{ij} \tag{13}$$

for action i. In class, we showed that the average regret after time T is $O(\ln m/T)$; i.e. the regret diminishes as $T \to \infty$. In other words, the marginal distribution over alleles that this algorithm converges to is comparable in loss to the point distribution on the i that minimizes the sum of $l_i(t)$ over time.

Minimizing this loss is equivalent to maximizing $m_x(i) = \sum_j y_j^t \Delta_{ij}$. The authors call this the *mixability*. It represents how well an allele performs in conjunction with all other possible genotypes. Different to asexual reproduction, which converges on the allele that has the highest *maximum* fitness, sexual production in the weak selection case converges on the allele with the highest *average* fitness.

This partially addresses some of the questions about the role of sex in evolution: Breaking apart successful gene combination was seen as a disadvantage of recombination, but it seems like this is more of a feature: Alleles that work well on average are selected, rather than those that work well in only a few cases.

3.3 Temporally Varying Fitness

With little additional effort, the same proofs can also be applied to the case when the fitness varies (randomly) over time. Say that the fitness of genotype ij was actually $w_{ij}^* = w_{ij} + v_{ij}^t$, where v_{ij}^t are i.i.d. random variables with zero mean. The v_{ij}^t must be chosen such that we still operate in weak selection, i.e. it must be that $w_{ij} + v_{ij}^t \in [1-s, 1+s]$.

Consider now ALG* (using the notation from class): It is

$$\begin{aligned}
\text{ALG}^* &= \sum_{t=1}^{T} l_i^*(t) && (14)\\
&= -\sum_{t=1}^{T} \sum_j y_j^t \Delta_{ij}^* && (15)\\
&= -\sum_{t=1}^{T} \sum_j y_j^t \frac{w_{ij} + v_{ij}^t - 1}{s} && (16)\\
&= \sum_{t=1}^{T} l_i(t) - \frac{1}{s} \sum_{t=1}^{T} \sum_j y_j^t v_{ij}^t && (17)
\end{aligned}$$

and similarly for OPT. This is simply the standard loss plus a weighted average of v_{ij} values. This means that the regret of this modified case is the same regret as the case when w_{ij} are fixed, plus a term

$$\frac{1}{s \cdot T} \sum_{t=1}^{T} \sum_{j} y_j^t v_{ij}^t \tag{18}$$

We know that v_{ij} are zero-mean i.i.d. random variables, and that the process converges to a stable distribution of y_j^t. Hence, we can apply the law of large numbers to conclude that this term tends to 0 as T tends to infinity.

This gives us the same regret as in the case when the w_{ij} are fixed, meaning that the same conclusions hold when the fitness values change slightly over time.

The proof in the paper follows a different approach, again referencing Kale [2]. The key part of the proof is to show that the randomness becomes a simple additive term in the regret and apply the law of large numbers. We therefore opted to reference the proof from class instead to focus only on this key part of the derivation and avoid the other mathematical details.

3.4 Diversity

The final corollary of the paper provides some evidence that evolution under recombination can converge to an equilibrium that maintains diversity.

In the asexual case, we can easily see that the evolutionary process will converge on the single genotype with the highest fitness scores. However, this is clearly not the case in nature: Diversity is ubiquitous, and if we can show that sexual reproduction is likely to lead to a diverse distribution of genes, it would significantly strengthen the role of sex in evolution.

The authors of this paper show that in the sexual reproduction case, there are exponentially many equilibria where the equilibrium distributions x and y are not point distributions. The proof in the paper is somewhat tedious, but we paraphrase parts of it and (hopefully) provide more intuition for the important steps.

Remember that the mixability of gene x is defined as $m_x(i) = \sum_j y_j w_{ij}$, or in other words, $m_x = Wy$, and $m_y = W^T x$.

At the equilibrium, all non-zero entries of x (and similarly for y) must have the same value; otherwise, x would change after one additional round of the update rule (Eq. (12)), and it would therefore not be an equilibrium distribution (and similarly y).

This means that the probability distributions x and y are a valid equilibrium only if $Wx = a\mathbf{1}$ and $W^T y = b\mathbf{1}$ for some non-negative constants a and b. Because some components of x and y can be 0, we can write equivalent equations $Ax' = a\mathbf{1}$ and $A^T y' = b\mathbf{1}$ for some submatrix A of W, as long as the subvectors x' and y' induced by the submatrix A sum to 1 (i.e. the rows and columns we discard from W are not part of the equilibrium) and are positive (i.e. we are not retaining rows/columns that don't contribute).

We want to know how many equilibria there are where x and y are not point distributions. For a submatrix A, the linear equations above tell us whether a valid equilibrium exists for some x', y', a and b. Therefore, we only need to count for how many of the (non-trivial) submatrices of W these linear equations are satisfied.

This is difficult to answer in general, and the authors instead do a probabilistic proof: They assume that the entries of W are *random* and drawn i.i.d. from $[1-s, 1+s]$. Then the problem reduces to computing the *probability* that a submatrix A fulfills these two equations.

The authors also assume that $a > 1$ and $b > 1$. This is because values smaller than one lead to extinction: If $\sum_j y_j w_{ij} < 1$, then certainly $\sum_j x_i y_j w_{ij} < 1$ for any i. However, this is the expected number of offspring that carry allele i. This means that the number of individuals carrying any allele shrinks at every reproduction cycle, and the population eventually goes extinct. This is not a useful equilibrium, and the assumption is reasonable.

Because we operate in weak selection, we can write $A = U + sB$, where U is a matrix of all 1s. Remember that we assumed the fitness matrix to be random, i.e. the entries of B are i.i.d. in $[-1, 1]$. Then, A must be square with almost certainty, because otherwise one of the equations $Ax' = a\mathbf{1}$ or $A^T y' = b\mathbf{1}$ is overdetermined, and the probability that a linear system of equations holds for a random matrix is almost zero. Hence, assume that $A \in \mathbb{R}^{k \times k}$ for some $k > 1$.

If we expand the linear equations, we get

$$Ax' = a\mathbf{1} \tag{19}$$

$$Ux' + sBx' = a\mathbf{1} \tag{20}$$

$$\mathbf{1} + sBx' = a\mathbf{1} \qquad \text{(because } x \text{ is a probability distribution)} \tag{21}$$

$$sBx' = (a-1)\mathbf{1} \tag{22}$$

$$x' = \frac{a-1}{s} B^{-1}\mathbf{1} \tag{23}$$

First, if an equilibrium exists for submatrix A, then B must be invertible. Furthermore, because every element of x' is positive and we assumed $a > 1$, it must be that the sum of rows of B^{-1} is positive. A similar equation exists for y', and hence the column sums of B^{-1} must also be positive.

What is the probability that a random matrix B is invertible and that its inverse has positive row and column sums? Unless the distribution of fitnesses is degenerate (i.e. discrete), the probability that B is invertible is almost 1. However, it could be that the inverse of B has negative row or column sums. But if this is the case, then we can "fix" it by modifying B until it has the desired properties.

Say B^{-1} had a row/column with negative sum. We can make the sum positive by "flipping" (i.e. inverting the signs of all values of) the row/column in B^{-1} (or conversely, flipping the column/row in B). It may be that this has created more rows or columns with negative sum. However, we know that the sum of the row/column before flipping was $-\sigma$, for some $\sigma > 0$ (otherwise, we would not have flipped it). Therefore, flipping the row/column increased the sum of all elements of B by 2σ. We can do this as long as there is a row or column with negative sum.

Through a simple counting argument, we can see that there exist exactly 2^{2k-1} distinct matrices we can obtain from flipping rows and columns of B. One of these matrices must have the largest sum of elements. Then, this matrix also has row- and column sums that are all positive: If not, we could flip a row or column and increase the sum of all elements - but this violates the maximum property.

Hence, out of all 2^{2k-1} possible permutations, at least one has positive row and column sums. This is true for any B. If we assume that the probability distribution of elements of B is symmetric around 0, then each of those permutations has equal likelihood; hence, the probability of a random B having positive row and column sums is at least $2^{-(2k-1)}$.

Putting all of this together, we can say the following: Assuming a *random fitness matrix* W with values distributed i.i.d. in $[1-s, 1+s]$ and symmetrically around 1, the probability that there exists an equilibrium on a submatrix $A \in \mathbb{R}^{k\times k}$ of W is $2^{-(2k-1)}$.

How many such submatrices are there? There are $\binom{n}{k}\binom{m}{k}$ ways to pick a submatrix of size $k \times k$ from W. Each of these has a probability of $2^{-(2k-1)}$ of having an equilibrium; then, the expected number of such matrices is

$$\binom{n}{k}\binom{m}{k}2^{-(2k-1)} \geq \left(\frac{n}{k}\right)^k \left(\frac{m}{k}\right)^k 2 \cdot 4^{-k} = 2\left(\frac{n\cdot m}{4k^2}\right)^k \tag{24}$$

In comparison, there are at most $n \cdot m$ equilibria where x and y are a point distribution. This means that the number of possible equilibria that predict diversity (i.e. $k > 1$) substantially outnumber sparse equilibria.

4 Simulation

For the purposes of experimentally verifying the results of the paper, we wrote a simulation which numerically calculates many years of evolution of a particular instance of a population and fitness tensor.

The main process for the simulation is:

- Set the initial conditions for the starting population, and fitness tensor. Depending on the type of experiment we want to run, we do this by filling them i.i.d. on some interval (like $[1-s, 1+s]$, or by systematically setting them as we go iterate through all possible values (within some small δ).
- Using equation 2 we iterate through some large number of generation (200) to see what the long term effects of evolution are on the population
- Depending on the type of experiment we are running we print out population information each year or only after a long period of time has occurred.

4.1 Experiment 1: How much does the initial population matter?

In section 2.2 we stated that population dynamics is much easier to study in the "Wright manifold" (recall this is when the population frequency can be written as the tensor product of the marginals of the genes: $p_{ij} = x_i \otimes y_j$). How do these dynamics change when we violate these conditions?

Consider the fitness tensor given by:

$$W = \begin{pmatrix} 2 & .5 \\ .5 & 2 \end{pmatrix} \tag{25}$$

Note that this tensor does not fit into the idea of "weak selection." Hence we have no reason to believe Nagylaki's theorem holds, and therefore we cannot claim that studying the dynamics on the Wright manifold is sufficient. Through extensive testing of marginals ranging from 0 (not present) to 1 (completely present) we find that there are only five equilibrium population frequencies.

$x = [0, 1]$ $y = [1, 0]$	$x = [1, 0]$ $y = [0, 1]$	$x = [\alpha, 1-\alpha]$ $y = [1-\alpha, \alpha]$ $0 < \alpha < 1$	$x = [\alpha, 1-\alpha]$ $y = [1-\beta, \beta]$ $0 \leq \alpha < \beta \leq 1$	$x = [\alpha, 1-\alpha]$ $y = [1-\beta, \beta]$ $0 \leq \beta < \alpha \leq 1$
$\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$	$\begin{pmatrix} .25 & .25 \\ .25 & .25 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$

Table 1: Table showing the equilibrium populations based on a fitness tensor given by 25 and the initial populations given by the first row. Notice that the first two equilibria are "unstable" meaning that if the starting state is even slightly off from those conditions then over time the system will tend towards on of the other equilibrium states. and that the last two are "stable equilibria" meaning if the system is slightly off of that equilibria it will correct itself.

What happens if we choose a initial population that is not in the Wright manifold? It turns out out of millions of runs with randomly chosen starting states, every equilibria falls into one of those five categories. While more equilibria could exist, if they do. they are most likely unstable equilibrium with very specific starting conditions. One direction to push this work further is to ask ourselves does this always hold?

4.2 Experiment 2: Diversity or Survival of only the Fittest?

The main result of this paper is that most instances of weak selection on the wright manifold will converge to a diverse equilibria rather than a sparse one. This experiment seeks to verify that result and expand upon it. Here we randomly choose an initial condition, and unlike in the previous analysis we do not require that the initial state is in the wright manifold. We chose a fitness tensor by randomly filling it's entries i.i.d. on the interval $[1-s, 1+s]$ for some varying parameter s. After allowing evolution to pass after some large number of years 200 (we actually scale this number according to the size of the tensors we are dealing with so that larger tensors will have more time to approach their end behavior) , we examine how many entries in the population frequency tensor are above some small value (.001 divided by the number of entries in the tensor). If this number is large then our population is fairly diverse confirming the results of the paper, however if this count is small, then our population is sparse and only filled with a few of the fittest gene combinations.

After a large number of tests we found that when s is large (not during weak selection) the diversity was small, and as s got smaller the diversity blossomed. In fact, for 14×14 matrices (one of our larger test conditions) and $s = \frac{1}{16}$ we found over half of our equilibrium

states had more non-zero entries than close to zero entries! Figures 1 and 2 summarize these results.

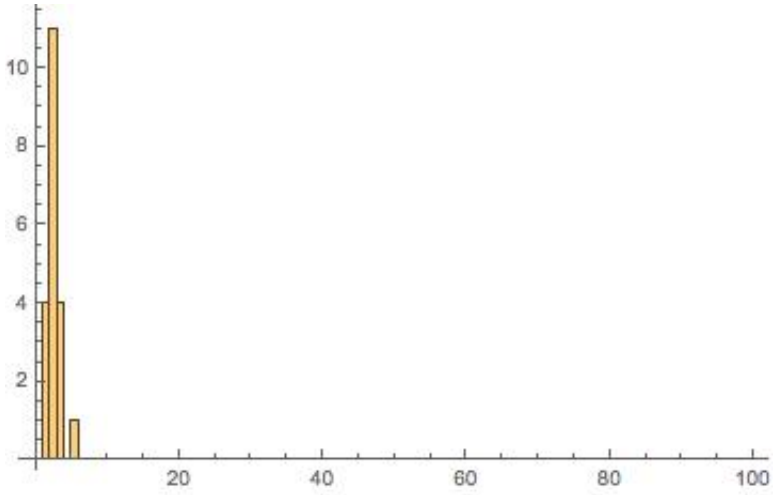

Figure 1: Histogram of final population with $s = \frac{1}{2}$ and dimensions of 10×10. For this value of the selection strength (s), we find that most of the end equilibria are very sparse, and almost all counts of surviving gene combinations are clustered close to 0.

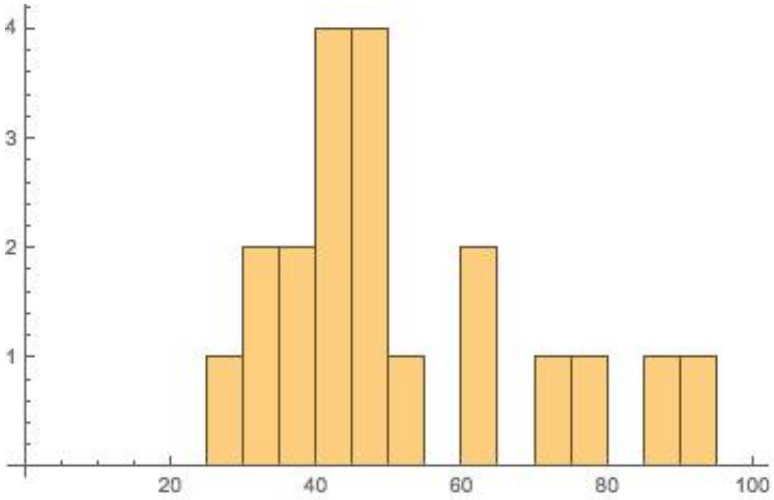

Figure 2: Histogram of final population with $s = \frac{1}{16}$ and dimensions of 10×10. For this value of the selection strength (s), we are operating within the bounds of weak selection and the conditions for the previous analysis hold (with the exception of the starting conditions) Unsurprisingly, we find the end equilibria are far more diverse with many gene combinations still being viable even after many years of evolution.

5 Conclusion

We closely investigated and explained the results of the work of Chastain et al. We've seen how evolution under recombination and cooperative games with multiplicative updates are identical if we assume weak selection. Furthermore, we have seen that in evolution under recombination, the potential equilibria in which the population is diverse vastly outnumber those where the population only has a sparse number of genotypes.

However, there are a number of assumptions that limit the scope of these results. The first and most restrictive assumption is that evolution operates in the *weak selection* regime. The authors base this on the "neutral theory" of evolution [3]. The issue with this assumption is that the neutral theory states that the majority of variation in genome between existing individuals of a species is explained by mutations in alleles that are selectively neutral, i.e. do not change the fitness scores by much. However, this is not the same as saying that all possible alleles (or mutations of alleles) have similar fitness scores; in fact, most mutations will be deleterious and lead to immediate extinction of that mutation. Similarly, although Kimura et al. claim that most changes in alleles are neutral, some may be advantageous (or disadvantageous). Therefore, assuming that the fitness matrix is in some small interval $[1-s, 1+s]$ significantly restricts the scope of the results, since we would expect most entries to be close to zero, some to be close to 1, and a few to significantly exceed 1.

Fortunately, follow-up work by Meir et al. [5] lifts this restriction, which significantly broadens the applicability of these results.

Finally, although the proof of diverse equilibria is encouraging, it only holds under quite significant assumptions: The fitness values must be random, i.i.d., distributed symmetrically around zero; it only holds under weak selection and only for two genes; and finally, although we can prove that many diverse equilibria exist, it is not clear how often these equilibria are actually reached. Knowing that a large number of diverse equilibria exist is not sufficient - it may be that the vast majority of initial conditions will converge to sparse equilibria only, though our experiments have hinted that this is not the case.

In conclusion, this was a very interesting paper to read and although the derivations were terse and sometimes hard to follow, I think the results are sound, especially in view of the followup work that lifts the major restrictions. We believe we were able to explain the derivations more clearly in the context of the class, and (hopefully) added some additional value through research in related works.

References

[1] Graham Bell. *The Masterpiece of Nature: The Evolution and Genetics of Sexuality*. Croom Helm, 1982.

[2] Satyen Kale. *Efficient Algorithms Using The Multiplicative Weights Update Method*. PhD thesis, Princeton University, 2007.

[3] Motoo Kimura. *The neutral theory of molecular evolution*. Cambridge University Press, 1983.

[4] Adi Livnat, Christos Papadimitriou, Jonathan Dushoff, and Marcus W. Feldman. A mixability theory for the role of sex in evolution. *Proceedings of the National Academy of Sciences*, 105(50):19803–19808, 2008.

[5] Reshef Meir and David Parkes. On sex, evolution, and the multiplicative weights update algorithm. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 929–937, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[6] T Nagylaki. The evolution of multilocus systems under weak selection. *Genetics*, 134(2):627–647, 1993.