Do Genes Have Fitness?

Tom Westerdale

April 9, 2014

After three days of enjoyable argument with John Holland and Leeann Fu in the summer of 1998, this paper is what I wished I had said. Most of it was written shortly after that summer. Since Holland so enjoys productive academic argument (discussion?), I have written this in the style of an attack on his position. And I think this style helps clarify the issues. Holland's academic viewpoint and the wealth of his ideas have given many of us our research inspiration and direction. In a sense, this paper is a thank you, for the ideas, for the encouragement, and for all the fun. Peter Woolf's observation sums it up: "Do you guys always argue like this?"[1]

1 What Does Fitness Mean?

> "Typically when Darwin is discussing flora and fauna, fitness means reproductive success; however, when he is discussing human populations fitness suddenly no longer means reproductive success, as that would entail recognising the poor with their large families as the fittest. Suddenly 'fitness' becomes suffused with the dominant social values of his time, filled with the ideas of social progress and superiority that elsewhere are given no tolerance in Darwinian theory. This ambiguity around 'fitness' – not to say downright slipperiness – in the original Darwin texts means that this central concept sits there almost asking to be recruited around any political project, typically by the social conservatives but also by social revolutionaries." (Hilary Rose [17])

> "Reproduction is a process in which individual strings are copied according to their objective function values, f (the biologists call this function the fitness function). Intuitively, we can think of the function f as some measure of profit, utility, or goodness that we want to maximise. Copying strings according to their fitness values means that strings with higher value have a higher probability of contributing one or more offspring to the next generation."
> (Goldberg [8, page 10])

In this paper, I shall use two terms, "value" and"reproductive rate". An individual or schema in a population will often have both a value and a reproductive rate. Its *value* is its profit, utility, or goodness, to use Goldberg's words. Value is what Goldberg calls fitness value. In genetic algorithms, the value is usually a number. On the other hand, *reproductive rate* is just what it says, reproductive rate. In a genetic algorithm or in an artificial breeding experiment, we might set reproductive rate equal to the value, or we might not. Usually we set it to some (usually monotonic) function of value.

Does the term "fitness" mean value or reproductive rate? Unfortunately the literature is so contradictory that the term is best avoided. We can see that to Goldberg, fitness means value. But he is wrong in thinking that this is the usual biological usage, and I have never seen the term "fitness function" in the biological literature. To Darwin, the term fitness was evidently slippery enough to mean either.

In a genetic algorithm, an individual's value is how well it succeeds at the assigned tasks in a given environment. If the assigned task is simply to reproduce, then the value is by definition equal to the reproductive rate. This is the situation in much of the population genetics literature, where wild populations are often being discussed. For this reason, it is sometimes hard to tell in this literature whether the term "fitness" is being used (1) to mean value or (2) to mean reproductive rate. A sentence like "Fitness is the reproductive rate." is an assertion of fact in case (1) and merely a definition in case (2). When we look in the biological literature, we see much of the slippery ambiguity that evidently occurs in Darwin.

For example, Fisher's fundamental theorem of natural selection says,

[1] Sections 14-19 were added at a later date and do not pertain to the original argument. They contain an idea that may be new and they are in this paper because they need its notation and results. The rest of the paper is all standard material and the only contribution here is its organization and its viewpoint.

> "The rate of increase in fitness of any organism at any time is equal to its genetic variance in fitness at that time." (Fisher [6, page 37])

But what does Fisher mean by a "fit" organism? He says,

> ".... whenever then, the socially lower occupations are more fertile, we must face the paradox that the biologically successful members of our society are to be found principally among its social failures, and equally that classes of persons who are prosperous and socially successful are, on the whole, the biological failures, the unfit of the struggle for existence, doomed more or less speedily, according to their social distinction, to be eradicated from the human stock."
> (Fisher [6, page 241])

So here, the rapidly reproducing "lower occupations" are the more fit, and fitness means reproductive rate. But ten pages later, he says,

> "Eugenists agree that the rising generation is largely recruited from the less fit. This is attributed partly to the fact that the upper classes marry later and partly to the fact that apart from the question of the postponement of marriage the upper classes are less prolific than the lower."
> (Fisher [6, page 251])

So here the less prolific "upper classes" are the more fit, and fitness means what Fisher thinks of as value. Slippery indeed.

But surely Fisher says clearly how he is using the term "fitness" in the theorem. Surely we need only find the sentence in [6] where he defines fitness. His other terms are all defined when they are first used, in sentences in which the defined word is italicised.

Unfortunately I can find no such definition of fitness. Fisher seems to regard "fitness" as a non-technical term whose meaning is already clear to the reader. On page 26 he says, "let us represent the relative rate of increase by m", and then on page 37, just before the statement of the theorem, he introduces the term "fitness" for the first time with the sentence,

> "Since m measures fitness by the objective fact of representation in future generations, the quantity $\Sigma'(2pa\alpha)$ will represent the contribution of each factor to the genetic variance in fitness." (Fisher [6, page 37])

The term "fitness" is not italicised and this seems to me to be a case of using the sentence, "Fitness is reproductive rate," not as a definition, but as an assertion of fact, an assertion that relies on assuming that the reader knows that the word "fitness" means what I am calling value.

So evidently for Fisher, "fitness" is an informal term that often means value, but sometimes means reproductive rate. It is not surprising to find that later population genetics books have tried to be more careful and formal in their use of the term.

> "Before we can calculate how the frequencies of genes in a population will change in time, we need some measure of the survival and reproduction of the different genotypes. This measure is **fitness** [his emphasis], which has a technical meaning in population genetics. The definition is easiest for populations with separate generations. .... Then the fitness of a particular type, A, is the expected number of offspring contributed by an A individual to the next generation."
> (John Maynard Smith [19, page 36])

> "We have used the word fitness as a synonym for the selective or reproductive value of a genotype. It may be either an absolute value, measured by the number of progeny per parent, or it may be relative to some reference genotype." (Crow and Kimura [2, page 224]

In these two quotes, "fitness" clearly means reproductive rate. It seems to me that this meaning is the common meaning in population genetics.

What is the practice in genetic algorithms?

> "In beginning this exploration we can make good use of a concept from mathematical genetics. The action of the environment $E \in \mathcal{E}$ upon the phenotype (and thereby upon the genotype $A \in \mathcal{A}$) is typically summarized in mathematical studies of genetics by a single performance measure μ_E called *fitness* [his italics]." (Holland [12, page 12])

> "In very general terms, each element of the population is tested against the environment and is ranked according to its fitness – its ability to survive and reproduce." (Holland [12, page 33])

> "First, the number of offspring of each individual A in a finite population $\mathcal{A}(t)$ is determined probabilistically, so that the expected number of offspring of A is proportional to A's observed fitness $\mu_E(A)$." (Holland [12, page 34])

Holland seems as slippery as Darwin and Fisher evidently are.

I think that most, but not all, of the genetic algorithms literature uses "fitness" to mean what I call value, as in the Goldberg quote above. By contrast, most, but not all, of the biological literature uses "fitness" to mean reproductive rate, as in the John Maynard Smith quote above. The early seminal writings on population genetics and genetic algorithms both seem particularly slippery.

I shall avoid the term "fitness" and use the terms "value" and "reproductive rate" instead. If I quote the literature, I will replace the term "fitness" with "value" or "reproductive rate" as is appropriate.[2]

In an evolving population, individuals, genotypes, and schemata necessarily have reproductive rates. And given that we fix our attention on some quality of the individuals, each individual has a value. If the quality is reproductive success, then the value presumably equals reproductive rate.

But do schemata have values in any meaningful sense? That is the question in the title of this paper, and it is the purpose of this paper to answer it. The answer here is not original. In fact, it is the standard answer that can be found (at least for schemata of length one) in most population genetics texts. But the examples in those texts are often not in a form natural to people working with genetic algorithms. In this paper, I have extended the examples in [2] to cover cases more familiar to genetic algorithms people. We will see that many of the narrowing assumptions in genetic algorithms formalism are unnecessary and that, in spite of his abhorrent views of human society, Fisher has much to tell us of how populations evolve. We will see that we need not follow Holland's restricting assumptions, which is pleasing since those assumptions could be used to underpin the sort of views that we glimpsed in the quotes from Fisher. It is also pleasing because it indicates that genetic algorithms may be more useful than those assumptions would imply.

2 Holland's Story

At any time, an evolutionary system has estimates of the values of schemata. It could at that point concentrate on *exploiting* the information it has already and produce lots of copies of what it thinks are the best schemata, or it could concentrate on *exploring* the schemata it knows less about, producing lots of copies of unexplored schemata. Exploitation produces the highest expected short term payoff at the cost of postponing discovery of even better schemata. An appropriate balance between exploration and exploitation is the essence of a good adaptive system, achieving reasonable payoff while exploring for even better opportunities. In Genetic Algorithms, exploration and exploitation mesh gracefully and reinforce each other.

The value of a schema is defined to be the average of the values of the individuals that contain the schema. At least that's what Holland says. Evolution (in biology or in a genetic algorithm) works by reproducing valuable schemata more rapidly than others. To do this, the evolutionary mechanism must estimate schema values. It does this by taking a statistical sample of individuals that contain a schema, and taking the average of the values of the individuals in the sample. It uses this average as an estimate of the schema value. The evolutionary mechanism uses a succession of populations, produced by either synchronous or a asynchronous reproduction, each population containing a sample of individuals that contain the schema. It thus obtains a succession of estimates of the schema value. At each stage, reproductive rate is proportional to the current estimate, so the more valuable schemata, having presumably higher value estimates, reproduce more rapidly and thus their successive samples are larger and larger and hence their value estimates become more and more accurate as exploration becomes concentrated more and more on potentially valuable schemata. At the same time, as the evolutionary process proceeds, the average value of the individuals in the populations tends to increase because the populations tend to have more and more valuable schemata, and since by definition, a schema's value is the average value of the individuals containing it, if the population has valuable schemata, the individuals must, on average, be valuable. Thus the ostensibly good schemata are exploited as the exploration proceeds. The power of the mechanism comes from the large number of schemata that are processed simultaneously. The evolutionary mechanism thus provides a powerful efficient way of increasing the average value of individuals in a population.

Of course schema reproductive rate is not strictly proportional to the current value estimate, because recombination affects the reproductive rate. By definition, the effect is less the more closely linked the

[2] Of course "value" is also a loaded term, but I'm using it to mean whatever we want to measure: size, speed of glycolysis at neutral pH, frequency of attendance at New York Yankees baseball games, whatever. I'm using "value" as in "x has the value 3," not as in "an Oxford degree has more value than a Birkbeck degree." Perhaps I should find another word, but at least "value" is never used in the literature to mean reproductive rate, unless reproductive rate is what we want to measure.

schema is, and for recombination operators like one or two point crossover, the closely linked schemata are those with short defining length. The Schema Theorem places a bound on the effect of this crossover, and quantifies the fact that schema reproductive rate is roughly proportional to the current value estimate for short schemata, and exactly proportional for single alleles. There are still a very large number of these short schemata and these are all processed simultaneously and in parallel. This implicit parallelism produces a population whose short schemata have relatively high value and hence whose individuals must have high value (since if the individuals had low value, the schema values would be low by definition.)

3 The Story is Wrong

A moment's reflection tells us that Holland's story makes no sense. The use of the word "average" in the definition of schema value obscures the fact that what value we say a schema has depends very much on what we are taking the average of.

The definition of schema value is supposed to be a generalization of the usual population genetics definition of allele value. But let us look at what population genetics says about allele value. Population genetics defines the value of an allele to be the weighted average of the values of those genotypes in the population that contain the allele. This average is computed as follows: We take the set of all individuals in the population that contain the allele, and we average their values, except that individuals that contain two copies of the allele are counted twice, those that contain three are counted thrice, etc. Thus the allele value is a weighted average of genotype values, the weights being the relative frequencies of the genotypes in the population. So we see immediately that even if the genotype values are independent of the genotype frequencies (and in many cases they aren't), the allele values are functions of the genotype frequencies.

The same is true for any sensible definition of schema value. The concept of schema value is a population concept, and a schema value changes as a population evolves and the weights in the weighted averages change.

Suppose we follow a particular schema through successive generations in an evolving population. We obtain a sequence of probably different values of that schema. The numbers in this sequence are the values, they are not estimates of a value. Holland seems to think that the schema has a single value, an invariant true value that is being estimated. But there is no such thing. How could there be? It would have to be an average, but what would the weights be?

Some of the genetic algorithms literature discusses what some call the "static value" of a schema. [9] This is the value the schema would have in a population in which the frequency of every genotype is the same. So all genotypes are present, their frequencies are uniform, and all weights are the same. Some of the literature implies that this static value is the true value of the schema and that the genetic algorithm is estimating this static value.

Now it is certainly true that in a genetic algorithm one can define such a static value, but it is rarely a value we are interested in. If we are interested in any one value, it would be a value of the schema in a population in which most of the other schemata are valuable, are the ones we want. After all, this is the sort of population we are trying to achieve.

Luckily, the genetic algorithm is quite obviously not trying to estimate the static value. The successive populations will rarely approach the uniform distribution required to get any estimate of the static value.

The uniform distribution of genotypes assumed in the definition of static value is so artificial that I am surprised anyone attaches significance to static value. To see how artificial it is, consider the following genotypes and values:

genotype	value
ADF	5
ADG	1
AEF	1
AEG	1
BDF	0
BDG	0
BEF	0
BEG	0
CDF	0
CDG	4
CEF	4
CEG	4

The static allele values at the first locus are:

allele	static-value
A	2
B	0
C	3

But notice that schemata DG, EF, and EG seem to have a similar effect on values. Perhaps it is better not to distinguish them. We will call them H. We call the DF schema J. The genotypes and values then are as follows:

genotype	value
AJ	5
AH	1
BJ	0
BH	0
CJ	0
CH	4

The static allele values at the first locus then are as follows:

allele	static-value
A	3
B	0
C	2

What the static allele values are depends on the representation. The static value is artificial.

The value of a schema or allele is a population concept and does not depend on the representation. For example, consider the following population of those same genotypes.

genotype	frequency	value
ADF	1/48	5
ADG	1/16	1
AEF	1/24	1
AEG	1/8	1
BDF	1/24	0
BDG	1/8	0
BEF	1/12	0
BEG	1/4	0
CDF	1/48	0
CDG	1/16	4
CEF	1/24	4
CEG	1/8	4

The allele values at locus 1 are then as follows:

allele	value
A	4/3
B	0
C	11/3

Grouping the schemata as before gives the following population:

genotype	frequency	value
AJ	1/48	5
AH	11/48	1
BJ	1/24	0
BH	11/24	0
CJ	1/48	0
CH	11/48	4

And the allele values at locus 1 are unchanged.

Fisher has shown us that we have no need of static values. Let us look at his analysis. All the values in the analysis will depend on the population. So it is a population of genotypes with which we begin.[3]

[3]The analysis here is not taken directly from Fisher. It is a modification of some of what is in Kimura. [2]

4 Asexual Reproduction

All our vectors are row vectors; their transposes are column vectors. The vector $\mathbf{e}$ is the vector all of whose entries are 1. If $\mathbf{u}$ is a vector, we write $\tilde{\mathbf{u}}$ to mean the diagonal matrix whose ii'th entry is u_i. We will sometimes give our numbered equations in two equivalent forms.

Let $\boldsymbol{\varphi}$ be the probability vector of genotype frequencies. We assume that each φ_i is positive. μ is a real valued function from the set of such probability vectors. We call $\mu(\boldsymbol{\varphi})$ the *value* of $\boldsymbol{\varphi}$. We extend μ to the set of all positive vectors (the positive orthant) by defining $\mu(c\boldsymbol{\varphi}) = c\mu(\boldsymbol{\varphi})$, for any positive real scalar c. We assume μ is differentiable. We define several quantities which depend on $\boldsymbol{\varphi}$.

$$\bar{m} = \mu(\boldsymbol{\varphi}) \tag{1}$$

$$m_i = \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i} \tag{2}$$

We call m_i the *value of genotype* i. It is something like the marginal utility of i. The vector $\mathbf{m}$ is the vector whose i'th entry is m_i.

Now think of a fixed positive vector $\boldsymbol{\vartheta}$ and a scalar variable c. Let $\boldsymbol{\varphi} = c\boldsymbol{\vartheta}$. Then

$$\frac{d}{dc}\mu(\boldsymbol{\varphi}) = \sum_i \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i}\frac{d\varphi_i}{dc} \quad .$$

Since $\varphi_i = c\vartheta_i$ and $\mu(\boldsymbol{\varphi}) = c\mu(\boldsymbol{\vartheta})$, this becomes

$$\mu(\boldsymbol{\vartheta}) = \sum_i \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i}\vartheta_i \quad ,$$

and multiplying by c gives

$$\mu(\boldsymbol{\varphi}) = \sum_i \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i}\varphi_i \quad ,$$

or

$$\bar{m} = \sum_i \varphi_i m_i \qquad\qquad \bar{m} = \boldsymbol{\varphi}\mathbf{m}^\top \quad . \tag{3}$$

We call $\bar{m}$ the average value, the average of the genotype values.

We shall be looking at $\boldsymbol{\varphi}$'s for which

$$\sum_i \varphi_i = 1 \qquad\qquad \boldsymbol{\varphi}\mathbf{e}^\top = 1 \quad . \tag{4}$$

These $\boldsymbol{\varphi}$'s form a hyperplane in the positive orthant. We call that portion of the hyperplane the *population domain*. We define

$$a_i = m_i - \bar{m} \qquad\qquad \mathbf{a} = \mathbf{m} - \bar{m}\mathbf{e} \quad . \tag{5}$$

We call a_i the *excess value of genotype* i.

We see that

$$\sum_i \varphi_i a_i = 0 \qquad\qquad \boldsymbol{\varphi}\mathbf{a}^\top = 0 \quad . \tag{6}$$

We define the *total variance* V as follows.

$$V = \sum_i \varphi_i a_i^2 \qquad\qquad V = \mathbf{a}\tilde{\boldsymbol{\varphi}}\mathbf{a}^\top \tag{7}$$

Suppose $\boldsymbol{\varphi}$ is changing continuously and differentiably. We let $'$ represent the derivative with respect to time. We see that

$$(\mu(\boldsymbol{\varphi}))' = \sum_i \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i}\varphi_i' \quad . \tag{8}$$

In other words,

$$\bar{m}' = \sum_i \varphi_i' m_i \quad . \tag{9}$$

Of course from equation (4) we have $\sum_i \varphi_i' = 0$, and so from this and equations (5) and (9) we have

$$\bar{m}' = \sum_i \varphi_i' a_i \quad . \tag{10}$$

If we suppose s_i is the number of copies of genotype i in the population, and $\bar{s} = \sum_i s_i$ is the total population size, then

$$\varphi_i = \frac{s_i}{\bar{s}} \quad .$$

A *natural reward scheme* is one in which the reproductive rate (birth rate minus death rate) of genotype i is $Ka_i + \chi$, where K and χ are the same for all i, but may be functions of $\boldsymbol{\varphi}$ or of time or of anything else, and where K is always non-negative. Then, using $'$ for time derivative, we have

$$s_i{}' = s_i(Ka_i + \chi) = K\bar{s}\varphi_i a_i + \chi s_i \quad . \tag{11}$$

From this and equation (6) we have

$$\bar{s}{}' = \sum_i s_i{}' = \chi\bar{s} \quad , \tag{12}$$

and so

$$\varphi_i{}' = \frac{s_i{}'\bar{s} - \bar{s}{}' s_i}{\bar{s}^2} = \frac{K\bar{s}\varphi_i a_i \bar{s} + \chi s_i \bar{s} - \chi \bar{s} s_i}{\bar{s}^2} = K\varphi_i a_i \quad .$$

In other words,

$$\varphi_i{}' = K\varphi_i a_i \qquad\qquad \boldsymbol{\varphi}' = K\mathbf{a}\tilde{\boldsymbol{\varphi}} \tag{13}$$

In fact, we see that if we differentiate $s_i = \varphi_i \bar{s}$ and use equation (13) we obtain

$$s_i{}' = s_i\left(Ka_i + \frac{\bar{s}{}'}{\bar{s}}\right) \quad ,$$

so if equation (13) holds, the reward scheme is natural. Thus a reward scheme is natural if and only if equation (13) holds.

Substituting equation (13) into equation (10) gives

$$\bar{m}' = K\sum_i \varphi_i a_i^2 \quad . \tag{14}$$

This shows us that $\bar{m}'$ is non-negative. This equation can be written

$$\bar{m}' = KV \quad . \tag{15}$$

This last equation is Fisher's Theorem for the case of no recombination.

Natural reward schemes are common in the literature. In many such schemes, K is constant and $\chi = K\mu(\boldsymbol{\varphi}) = K\bar{m}$. So in these schemes, the population size is changing at rate $\bar{s}{}' = K\bar{m}\bar{s}$, and the reproductive rate of genotype i is Km_i. Here K is a constant of proportionality and we have what we can call *value proportional reproductive rate*.[4]

If we have value proportional reproductive rate, if the reproductive rate of genotype i is Km_i, for all i (K a constant), then the average reproductive rate is $K\bar{m}$. We can ask how the average reproductive rate changes over time, and we see that its rate of change is $K\bar{m}'$. By equation (15), this is K^2V, a non-negative number. If reproductive rate is what we are measuring, if value m_i is reproductive rate of genotype i, then $K = 1$, and V is the variance of reproductive rate, and average reproductive rate is increasing at rate V. Of course the population size is continually changing. This situation appears in the biological literature.

On the other hand, if we set $\chi = 0$ then the population size is constant, as we can see by equation (12). One common way of doing this in a genetic algorithm is for genotype i to have birth rate Km_i, and for the population size to be held constant by artificially eliminating (killing) individuals without regard to their genotype. Since each individual is equally likely to be killed, the death rate from this of all genotypes is the same. Let's call it ν. (Of course ν may be changing with time.) So the reproductive rate of i is $Km_i - \nu$, and $s_i{}' = s_i(Km_i - \nu) = \bar{s}\varphi_i(Ka_i + K\bar{m} - \nu) = K\bar{s}\varphi_i a_i + \bar{s}\varphi_i(K\bar{m} - \nu)$. If we sum this equation over all i and use equations (6) and (4), we obtain $\bar{s}{}' = \sum_i s_i{}' = \bar{s}(K\bar{m} - \nu)$. Since our killing is keeping the population size constant, $\bar{s}{}' = 0$, so $\nu = K\bar{m}$, and the reproductive rate is $Km_i - \nu = Ka_i$. In other words, $\chi = 0$.

The derivation of equation (15) was independent of χ, so equation (15) holds happily with $\chi = 0$. In this case, m_i is of course the value of i, but it cannot be the reproductive rate of i (unless $\bar{m} = 0$), since the reproductive rate is Ka_i. Equation (15) says how the average value $\bar{m}$ is changing, but not how average reproductive rate is changing. Average reproductive rate isn't changing, it's zero. The population size is constant.

[4] Holland calls this "fitness proportional selection".

5 Steepest Ascent

Remember that we are calling the set of probability vectors (the set of stochastic vector values that $\boldsymbol{\varphi}$ can take) the population domain. So $\mathbf{v}$ is in the population domain if and only if $\mathbf{v}\,\mathbf{e}^\top = 1$ and every v_i is positive. We are thinking of $\boldsymbol{\varphi}$ as traveling differentiably over the population domain, and thinking of μ as a function from the population domain (even though it's actually from the whole positive orthant). We want $\boldsymbol{\varphi}$ to climb uphill on the function μ. Note that $\bar{m}'$ is the rate of climb. We see from equation (15) that if we use a natural reward scheme then $\boldsymbol{\varphi}$ never goes downhill (since $V \geq 0$), and it climbs uphill unless $V = 0$.

Note that the following four statements are equivalent: (1) $V = 0$; (2) every a_i is zero; (3) the m_i's are all the same; (4) the a_i's are all the same. If these statements hold, I shall say *the ground is level*. Note that the ground is level at a local maximum of μ. (Local maximum of μ restricted to the population domain.)

Although natural reward schemes send $\boldsymbol{\varphi}$ uphill, they do not in general send it by the path of steepest ascent. But by transforming axes, we can make the natural reward schemes into steepest ascent schemes.

Define $\psi_i = \sqrt{\varphi_i}$, so $\varphi_i = \psi_i^2$. As $\boldsymbol{\varphi}$ travels over the population domain in our original vector space, its image $\boldsymbol{\psi}$ travels over the transformed space, which I shall call the square root space. The population domain in the original space corresponds to a portion of the unit sphere in the square root space, and $\boldsymbol{\psi}$ travels over that portion.

Corresponding to the value function μ over the positive orthant of the original space, we define the value function f over the positive orthant of the square root space by $f(\boldsymbol{\psi}) = \mu(\boldsymbol{\varphi}) = \bar{m}$. For each genotype i, we define the function f_i by

$$f_i(\boldsymbol{\psi}) = \frac{\partial f(\boldsymbol{\psi})}{\partial \psi_i} \quad .$$

We want $\boldsymbol{\psi}$ to move uphill on the f function in the square root space.

Suppose $\boldsymbol{\psi}$ is at some point on the unit sphere in the square root space, and that $\boldsymbol{\psi}$ is in the positive orthant. We want $\boldsymbol{\psi}'$ to aim in the steepest ascent direction. We assume the ground is not level, that is, we assume $V \neq 0$. To find the steepest ascent direction, we need to find the unit vector $\mathbf{x}$ that maximizes $f(\boldsymbol{\psi}) + \sum_i f_i(\boldsymbol{\psi}) x_i$, subject to the constraints $\|\mathbf{x}\| = 1$, and $\mathbf{x}$ must be a vector tangent to the unit sphere at point $\boldsymbol{\psi}$. That is, $\boldsymbol{\psi} \cdot \mathbf{x} = 0$. We use Lagrange multipliers. We define

$$F(\mathbf{x}) = f(\boldsymbol{\psi}) + \sum_i f_i(\boldsymbol{\psi}) x_i + A \sum_i x_i^2 + B \sum_i \psi_i x_i \quad .$$

Now

$$f_i(\boldsymbol{\psi}) = \frac{\partial f(\boldsymbol{\psi})}{\partial \psi_i} = \frac{\partial \mu(\boldsymbol{\varphi})}{\partial \varphi_i} \frac{d\varphi_i}{d\psi_i} = m_i 2\psi_i \quad .$$

Let

$$F_i(\mathbf{x}) = \frac{\partial F(\mathbf{x})}{\partial x_i} \quad ,$$

so

$$F_i(\mathbf{x}) = 2m_i\psi_i + A2x_i + B\psi_i \quad .$$

If this is zero for all i then $\mathbf{x}$ will point either in the steepest ascent or steepest descent direction, so we set $F_i(\mathbf{x})$ to zero and obtain

$$2m_i\psi_i + A2x_i + B\psi_i = 0 \tag{16}$$

Multiplying by ψ_i gives $2m_i\varphi_i + A2\psi_i x_i + B\psi_i^2 = 0$.
Summing over i then gives $2\bar{m} + A2(\boldsymbol{\psi} \cdot \mathbf{x}) + B\|\boldsymbol{\psi}\|^2 = 0$, which is simply $2\bar{m} + B = 0$. Therefore,
$B = -2\bar{m}$. Substituting this back into equation (16) gives $2m_i\psi_i + A2x_i - 2\bar{m}\psi_i = 0$. Dividing by 2 gives $m_i\psi_i + Ax_i - \bar{m}\psi_i = 0$, which is $Ax_i + a_i\psi_i = 0$. Thus we obtain

$$Ax_i = -\psi_i a_i \quad . \tag{17}$$

Squaring this gives $A^2 x_i^2 = \varphi_i a_i^2$, and then summing over i gives $A^2\|\mathbf{x}\|^2 = V$, so we see that $A^2 = V$, and A is either $\sqrt{V}$ or $-\sqrt{V}$. Equation (17) gives us $x_i = -\frac{1}{A}\psi_i a_i$. The $\mathbf{x}$ so defined points in either the steepest ascent or steepest descent direction, so we want $\boldsymbol{\psi}'$ to equal this $\mathbf{x}$. That is, ${\psi_i}' = -\frac{1}{A}\psi_i a_i$, and so ${\varphi_i}' = 2\psi_i{\psi_i}' = -\frac{2}{A}\varphi_i a_i$. Then using equation (10) gives $\bar{m}' = \sum_i {\varphi_i}' a_i = -\frac{2}{A}\sum_i \varphi_i a_i^2 = -\frac{2}{A}V$. We said that A is either $\sqrt{V}$ or $-\sqrt{V}$, and we can now see that $-\sqrt{V}$ gives steepest ascent and $\sqrt{V}$ gives steepest descent. Thus for steepest ascent, we need to have ${\varphi_i}' = \frac{2}{\sqrt{V}}\varphi_i a_i$. Any positive multiple of these values will do, so we see that we get steepest ascent if and only if we have ${\varphi_i}' = K\varphi_i a_i$ (equation (13)), for some positive number K. This is the natural reward scheme.

That argument using Lagrange multipliers is the usual approach to these matters, as I shall discuss in section 19, but I must admit that the conclusion of the above argument feels counterintuitive to me. We are saying that as long as the ground is not level, as long as the a_i's are not all the same, then there is exactly one steepest ascent direction. But what if for example there are four genotypes, the value function μ is given by $\mu(\boldsymbol{\varphi}) = \varphi_1 + \varphi_2 - \varphi_3 - \varphi_4$, and $\boldsymbol{\varphi} = \langle \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \rangle$. Then $\mathbf{a} = \langle 1, 1, -1, -1 \rangle$. How could it possibly make any difference to the rate of climb in the square root space if we raise the rate of increase of φ_1 and compensate by lowering the rate of increase of φ_2 by the same amount? After all, the two genotypes are worth exactly the same. ($a_1 = a_2 = 1$) True, but it turns out that that sort of compensation is inappropriate for finding steepest ascent in the square root space. I found it all a bit confusing. Lagrange multipliers seem too much like hocus pocus to me and I don't really trust my ability to use them properly.

Luckily, I've found another more geometric proof of the above result that makes it clear there really is just one steepest ascent direction. Since I will want to use this result later, I now give the more geometric proof.

Define the vector $\boldsymbol{\psi}$ as above, and define the vector $\mathbf{v} = 2\mathbf{a}\tilde{\boldsymbol{\psi}}$. We assume the ground is not level, so $\mathbf{v}$ is not the zero vector. Then $v_i = 2\psi_i a_i$ and $v_i^2 = 4\varphi_i a_i^2$. Summing over i gives $\|\mathbf{v}\|^2 = 4V$, so $\|\mathbf{v}\| = 2\sqrt{V}$. Note that $\mathbf{v} \cdot \boldsymbol{\psi} = \sum_i 2\psi_i a_i \psi_i = 2\sum_i \varphi_i a_i = 0$.

As before, we look for a unit vector $\mathbf{x}$ that points in the steepest ascent direction at point $\boldsymbol{\psi}$, which is a point on the unit sphere. But as before, $\mathbf{x}$ must point tangent to the unit sphere at point $\boldsymbol{\psi}$, that is, $\mathbf{x}$ must be orthogonal to the vector $\boldsymbol{\psi}$. So we want $\mathbf{x}$ to point in the steepest ascent direction at point $\boldsymbol{\psi}$ subject to the constraint that $\mathbf{x}$ must be orthogonal to $\boldsymbol{\psi}$.

Now suppose we point $\mathbf{x}$ in some direction. What is the rate of ascent in that direction? It is what the rate of climb $\bar{m}'$ would be if we set $\boldsymbol{\psi}'$ equal to $\mathbf{x}$. Now notice that if we set $\boldsymbol{\psi}'$ equal to $\mathbf{x}$, then by $\varphi_i' = 2\psi_i\psi_i'$ and equation (10) we have $\mathbf{x} \cdot \mathbf{v} = \boldsymbol{\psi}' \cdot \mathbf{v} = \sum_i \psi_i' 2\psi_i a_i = \sum_i \varphi_i' a_i = \bar{m}'$. So $\mathbf{x} \cdot \mathbf{v}$ is the rate of ascent in direction $\mathbf{x}$.

So in which direction should we point $\mathbf{x}$? Remembering that we must keep the unit vector $\mathbf{x}$ orthogonal to $\boldsymbol{\psi}$, we want to point it so as to maximize $\mathbf{x} \cdot \mathbf{v}$. So we want to point it so as to minimize the angle between $\mathbf{x}$ and $\mathbf{v}$. How small can we make that angle and still have $\mathbf{x}$ orthogonal to $\boldsymbol{\psi}$?

It's easy. We already noted that $\mathbf{v} \cdot \boldsymbol{\psi} = 0$, so $\mathbf{v}$ is itself orthogonal to $\boldsymbol{\psi}$. We can get the angle down to zero degrees by simply pointing $\mathbf{x}$ in the same direction as $\mathbf{v}$. And we see easily that this gives the direction of steepest ascent and that this is unique. There is exactly one direction of steepest ascent unless the ground is level. Any other direction gives a lower rate of ascent.

To tidy up, let's point the unit vector $\mathbf{x}$ in the same direction as $\mathbf{v}$ and look at the reward scheme we get. In that case, $\mathbf{x} = \frac{1}{\|\mathbf{v}\|}\mathbf{v}$. We set $\boldsymbol{\psi}' = \mathbf{x}$ and so we have $\psi_i' = x_i = \frac{1}{\|\mathbf{v}\|} v_i = \frac{1}{2\sqrt{V}} 2\psi_i a_i = \frac{1}{\sqrt{V}} \psi_i a_i$. Then $\varphi_i' = 2\psi_i\psi_i' = \frac{2}{\sqrt{V}} \varphi_i a_i$, which is exactly the result we got using Lagrange multipliers.

In the square root space, as long as the ground is not level there is one unique steepest ascent direction. We follow that steepest ascent direction if and only if we use a natural reward scheme.

6 Epistasis

We let N be the ploidy, and L be the number of loci. Associated with each locus is its *locus set*, the indices of those alleles that can be found at that locus. (We will use the letter $\mathcal{L}$ to refer to a locus set.) The locus sets are disjoint.

For example, suppose $N = 3$ (triploid) and $L = 4$, and suppose the four locus sets are
$\{1, 2, 3\}, \{4, 5\}, \{6, 7\}, \{8, 9\}$.
Let's look at one of these triploid genotypes:

1 3 3
4 4 5
7 7 7
8 9 8

We can think of the genotype as consisting of 3 chromosomes. The chromosomes are ordered, so the genotype

3 1 3
4 4 5
7 7 7
9 8 8

is a different genotype. Of course in many applications these two genotypes will have the same value.[5] In humans, $N = 2$, and so each genotype has two chromosomes, except that each of these chromosomes is chopped up into 23 separate segments (called chromosomes). Such chopping will be irrelevant in this essay since we will not be discussing linkage details.

We let c_{ij} be the number of occurrences of allele j in genotype i. The matrix C is the matrix whose ij'th entry is c_{ij}.

The probability of allele j in the population is p_j. That is,

$$p_j = \frac{1}{N} \sum_i \varphi_i c_{ij} \qquad\qquad \mathbf{p} = \tfrac{1}{N} \boldsymbol{\varphi} C \quad . \tag{18}$$

The vector $\mathbf{e}$ has the same dimensionality as $\boldsymbol{\varphi}$. The "shorter" vector of all ones, which has the same dimensionality as $\mathbf{p}$, we write as $\check{\mathbf{e}}$.

For any genotype i and any locus set $\mathcal{L}$, we have

$$\sum_{j \in \mathcal{L}} c_{ij} = N \quad . \tag{19}$$

Now

$$\sum_{j \in \mathcal{L}} \frac{1}{N} \sum_i \varphi_i c_{ij} = \frac{1}{N} \sum_i \varphi_i \sum_{j \in \mathcal{L}} c_{ij} \quad ,$$

so

$$\sum_{j \in \mathcal{L}} p_j = 1 \quad . \tag{20}$$

Since there are L loci,

$$\sum_j p_j = L \quad . \tag{21}$$

We let $\bar{\varphi}_i$ be the probability that genotype i would have if we did complete recombination. That is,

$$\bar{\varphi}_i = \prod_j p_j^{c_{ij}} \quad . \tag{22}$$

The population $\boldsymbol{\varphi}$ is said to be at linkage equilibrium, or at Hardy-Weinberg equilibrium, if $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$.

We define the *excess value of allele* j as

$$\alpha_j = \frac{1}{N p_j} \sum_i \varphi_i c_{ij} a_i \qquad\qquad \boldsymbol{\alpha} = \tfrac{1}{N} \, \mathbf{a} \, \tilde{\boldsymbol{\varphi}} \, C \, \tilde{\mathbf{p}}^{-1} \quad . \tag{23}$$

In other words, the excess value α_j of allele j is the weighted average of the excess values of the genotypes in which j occurs, counting a genotype twice if j occurs in it twice, thrice if j occurs in it thrice, etc.

Note that

$$\sum_{j \in \mathcal{L}} p_j \frac{1}{N p_j} \sum_i \varphi_i c_{ij} a_i = \sum_i \varphi_i a_i \frac{1}{N} \sum_{j \in \mathcal{L}} c_{ij} = \sum_i \varphi_i a_i = 0 \quad .$$

Then from this and equation (23), we obtain the following two equations.

$$\sum_{j \in \mathcal{L}} p_j \alpha_j = 0 \tag{24}$$

$$\sum_j p_j \alpha_j = 0 \qquad\qquad \mathbf{p} \, \boldsymbol{\alpha}^\top = 0 \tag{25}$$

We now define the *excess genic value* $\hat{a}_i$ of genotype i.

$$\hat{a}_i = \sum_j c_{ij} \alpha_j \qquad\qquad \hat{\mathbf{a}} = \boldsymbol{\alpha} \, C^\top \tag{26}$$

[5] Suppose $N = 6$ and $L = 1$, and the locus set is $\{1, 2, 3\}$. Then 132133 and 331231 are distinct genotypes. They might have different values. The system might have an architecture that learns which the better orderings are, or it might not. See [23] and [20] for a discussion of why we call this situation "payoff only" and how to use this as a basis for studying adaptive rule based systems.

From equations (26), (18), and (25), we have $\boldsymbol{\varphi}\,\hat{\mathbf{a}}^\top = \boldsymbol{\varphi} C \boldsymbol{\alpha}^\top = N\mathbf{p}\boldsymbol{\alpha}^\top = 0$, so

$$\boldsymbol{\varphi}\,\hat{\mathbf{a}}^\top = 0 \qquad\qquad \sum_i \varphi_i \hat{a}_i = 0 \quad . \tag{27}$$

The quantity $a_i - \hat{a}_i$ is called the *epistatic (and dominance) value* of the genotype i. Note that the value of a genotype is the average population value, plus the excess genic value, plus the epistatic value. That is, $m_i = \bar{m} + \hat{a}_i + (a_i - \hat{a}_i)$.

We now define the *genic variance* V_g and the *epistatic variance* V_e.

$$V_g = \sum_i \varphi_i \hat{a}_i^2 \qquad\qquad V_g = \hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top \tag{28}$$

$$V_e = \sum_i \varphi_i (a_i - \hat{a}_i)^2 \qquad\qquad V_e = (\mathbf{a} - \hat{\mathbf{a}})\,\tilde{\boldsymbol{\varphi}}\,(\mathbf{a} - \hat{\mathbf{a}})^\top \tag{29}$$

7 An Example

Here is an example: $N = 1$ and $L = 2$. There are 4 alleles, which we shall call A, B, C, and D. Locus 1 is either A or B, and locus 2 is either C or D. So the genotypes are AC, AD, BC, and BD, which we number in that order. Their frequencies are $\frac{1}{2}, 0, 0, \frac{1}{2}$, respectively, and their values are $6, 100, 3, 0$, respectively. So for example, genotype 2 is AD, $\varphi_2 = 0$, and $m_2 = 100$. Thus $\bar{m} = 3$ and we have the following quantities.

i	genotype	φ_i	m_i	a_i	$\hat{a}_i$	$a_i - \hat{a}_i$	$\bar{\varphi}_i$
1	AC	$\frac{1}{2}$	6	3	6	-3	$\frac{1}{4}$
2	AD	0	100	97	0	97	$\frac{1}{4}$
3	BC	0	3	0	0	0	$\frac{1}{4}$
4	BD	$\frac{1}{2}$	0	-3	-6	3	$\frac{1}{4}$

j	allele	p_j	α_j
1	A	$\frac{1}{2}$	3
2	B	$\frac{1}{2}$	-3
3	C	$\frac{1}{2}$	3
4	D	$\frac{1}{2}$	-3

$\hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}\hat{\mathbf{a}}^\top = 18$
$\hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}\mathbf{a}^\top = 18$
$V_g = \hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}\hat{\mathbf{a}}^\top = 36$
$V = \mathbf{a}\tilde{\boldsymbol{\varphi}}\mathbf{a}^\top = 9$
$V_e = (\mathbf{a} - \hat{\mathbf{a}})\tilde{\boldsymbol{\varphi}}(\mathbf{a} - \hat{\mathbf{a}})^\top = 9$
$N\boldsymbol{\alpha}\tilde{\mathbf{p}}\boldsymbol{\alpha}^\top = 18$
$\hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}(\mathbf{a} - \hat{\mathbf{a}})^\top = -18$

8 Some Useful Equations

Consider the following useful equation.

$$\prod_{\mathcal{L}} \left(\sum_{j \in \mathcal{L}} z_j \right)^N = \sum_i \prod_j z_j^{c_{ij}} \tag{30}$$

where $\mathcal{L}$ ranges over locus sets and the z_j's are any numbers.

It is really no more than the distributive law. If it's not obvious, consider a haploid example. (So in the example, $N = 1$, and each c_{ij} is either 1 or 0.) Look at the twelve genotypes in the first chart in section 3 and ignore the values listed there. Pretend A, B, C, D, E, F, and G are numbers, the z_j's. Now each of the genotypes looks like a product, and if we sum the products we are doing what the right side of equation (30) tells us to do. On the other hand, that sum is clearly $(A + B + C)(D + E)(F + G)$, which is the left side of equation (30).

Now if we set $z_j = p_j$ in equation (30) and use equation (20), we have

$$1 = \sum_i \prod_j p_j^{c_{ij}} \quad ,$$

which by equation (22) is

$$\sum_i \bar{\varphi}_i = 1 \quad . \tag{31}$$

We use a similar trick in lemmas 1 and 3 below.

Lemma 1 $N\mathbf{p} = \bar{\boldsymbol{\varphi}} C$

Proof: We start with equation (30):

$$\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} z_k \right)^N = \sum_i \prod_k z_k^{c_{ik}} \quad . \tag{32}$$

Select any allele j and let $\hat{\mathcal{L}}$ be the locus set of which j is a member. We set $z_k = x^{\delta_{jk}} p_k$, where x is a variable and δ_{jk} is the Kronecker delta. Then $z_k - p_k = (x^{\delta_{jk}} - 1)p_k$ and
$(\sum_{k\in\hat{\mathcal{L}}} z_k) - (\sum_{k\in\hat{\mathcal{L}}} p_k) = \sum_{k\in\hat{\mathcal{L}}}(z_k - p_k) = (x-1)p_j$. Then by equation (20), this becomes
$\sum_{k\in\hat{\mathcal{L}}} z_k = xp_j - p_j + 1$. Equation (20) also tells us that if $\mathcal{L} \neq \hat{\mathcal{L}}$, then $\sum_{k\in\mathcal{L}} z_k = \sum_{k\in\mathcal{L}} p_k = 1$. These last two equations give this:

$$\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} z_k \right)^N = (xp_j - p_j + 1)^N \tag{33}$$

Now $j = k \Rightarrow x^{\delta_{jk}c_{ik}} = x^{c_{ij}}$, whereas $j \neq k \Rightarrow x^{\delta_{jk}c_{ik}} = 1$. Therefore $\prod_k x^{\delta_{jk}c_{ik}} = x^{c_{ij}}$. If we use this and equation (22), we obtain $\prod_k z_k^{c_{ik}} = (\prod_k x^{\delta_{jk}c_{ik}})(\prod_k p_k^{c_{ik}}) = x^{c_{ij}}\bar{\varphi}_i$.
If we use this and equation (33), then equation (32) becomes $(xp_j - p_j + 1)^N = \sum_i \bar{\varphi}_i x^{c_{ij}}$.
Taking derivative with respect to x gives $N(xp_j - p_j + 1)^{N-1} p_j = \sum_i \bar{\varphi}_i c_{ij} x^{c_{ij}-1}$, and setting $x = 1$ turns this into $Np_j = \sum_i \bar{\varphi}_i c_{ij}$. □

From lemma 1 and equation (18), we have

$$\bar{\boldsymbol{\varphi}} C = \boldsymbol{\varphi} C \quad . \tag{34}$$

Lemma 2 $\bar{\bar{\boldsymbol{\varphi}}} = \bar{\boldsymbol{\varphi}}$

Proof: define $\boldsymbol{\vartheta} = \bar{\boldsymbol{\varphi}}$ and forget for the moment how $\boldsymbol{\vartheta}$ was defined and think of it just as a distribution vector over genotypes. We ask what $\bar{\boldsymbol{\vartheta}}$ is.

Well, $\boldsymbol{\vartheta}$ induces a distribution vector $\mathbf{q}$ over alleles, and $\mathbf{q}$ determines $\bar{\boldsymbol{\vartheta}}$. The equations are analogous to equations (18) and (22), namely

$$\mathbf{q} = \frac{1}{N} \boldsymbol{\vartheta} C \tag{35}$$

and

$$\bar{\vartheta}_i = \prod_j q_j^{c_{ij}} \quad . \tag{36}$$

We now remember that $\boldsymbol{\vartheta} = \bar{\boldsymbol{\varphi}}$, so equation (35) becomes $\mathbf{q} = \frac{1}{N}\bar{\boldsymbol{\varphi}} C$, which by equation (34) becomes $\mathbf{q} = \frac{1}{N}\boldsymbol{\varphi} C$. Comparing this with equation (18) tells us that $\mathbf{q} = \mathbf{p}$, so equation (36) becomes $\bar{\vartheta}_i = \prod_j p_j^{c_{ij}}$. This and equation (22) gives $\bar{\vartheta}_i = \bar{\varphi}_i$. Thus $\bar{\boldsymbol{\vartheta}} = \bar{\boldsymbol{\varphi}}$. But we also have $\bar{\boldsymbol{\vartheta}} = \bar{\bar{\boldsymbol{\varphi}}}$ (since $\boldsymbol{\vartheta} = \bar{\boldsymbol{\varphi}}$). □

Lemma 3 *If $\boldsymbol{\beta}$ is a real vector of the same dimensionality as $\mathbf{p}$, and if $\sum_{j\in\mathcal{L}} p_j \beta_j = 0$ for all locus sets $\mathcal{L}$, then*

$$\boldsymbol{\beta} = \tfrac{1}{N}\, \boldsymbol{\beta}\, C^\top \tilde{\bar{\boldsymbol{\varphi}}}\, C\, \tilde{\mathbf{p}}^{-1} \quad .$$

Proof: First we define $\hat{\mathbf{b}} = \boldsymbol{\beta} C^\top$, so we have $\hat{b}_i = \sum_j c_{ij}\beta_j$ for all i. We now take equation (30).

$$\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} z_k \right)^N = \sum_i \prod_k z_k^{c_{ik}}$$

We substitute $y_k x^{\beta_k}$ for z_k, where the y_k's and the x are independent real variables. We obtain

$$\begin{aligned}
\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k x^{\beta_k} \right)^N &= \sum_i \prod_k y_k^{c_{ik}} x^{c_{ik}\beta_k} \\
&= \sum_i \left(\prod_k y_k^{c_{ik}} \right) \left(\prod_k x^{c_{ik}\beta_k} \right) \\
&= \sum_i \left(\prod_k y_k^{c_{ik}} \right) x^{\sum_k c_{ik}\beta_k} \\
&= \sum_i \left(\prod_k y_k^{c_{ik}} \right) x^{\hat{b}_i} \quad .
\end{aligned}$$

We can write this as

$$exp\left(\sum_{\mathcal{L}} N\, log\left(\sum_{k \in \mathcal{L}} y_k x^{\beta_k} \right) \right) = \sum_i \left(\prod_k y_k^{c_{ik}} \right) x^{\hat{b}_i} \quad .$$

Taking derivative with respect to x gives

$$exp\left(\sum_{\mathcal{L}} N\, log\left(\sum_{k \in \mathcal{L}} y_k x^{\beta_k} \right) \right) \sum_{\mathcal{L}} N \frac{\sum_{k \in \mathcal{L}} y_k \beta_k x^{\beta_k - 1}}{\sum_{k \in \mathcal{L}} y_k x^{\beta_k}} = \sum_i \left(\prod_k y_k^{c_{ik}} \right) \hat{b}_i x^{\hat{b}_i - 1} \quad .$$

Setting $\mathbf{x} = 1$ gives

$$exp\left(\sum_{\mathcal{L}} N\, log\left(\sum_{k \in \mathcal{L}} y_k \right) \right) N \sum_{\mathcal{L}} \frac{\sum_{k \in \mathcal{L}} y_k \beta_k}{\sum_{k \in \mathcal{L}} y_k} = \sum_i \left(\prod_k y_k^{c_{ik}} \right) \hat{b}_i \quad . \tag{37}$$

In the proof of this lemma only, let $'$ mean derivative with respect to y_j. Let $\hat{\mathcal{L}}$ be the locus set of which j is a member. Then using $exp(\sum_{\mathcal{L}} N\, log(\sum_{k \in \mathcal{L}} y_k)) = \prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k\right)^N$ gives the following two equations.

$$\begin{aligned}
\left(exp\left(\sum_{\mathcal{L}} N\, log\left(\sum_{k \in \mathcal{L}} y_k \right) \right) \right)' &= \left(exp\left(\sum_{\mathcal{L}} N\, log\left(\sum_{k \in \mathcal{L}} y_k \right) \right) \right) \left(N\, log\left(\sum_{k \in \hat{\mathcal{L}}} y_k \right) \right)' \\
&= \left(\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k \right)^N \right) \left(\frac{N}{\sum_{k \in \hat{\mathcal{L}}} y_k} \right)
\end{aligned} \tag{38}$$

$$\left(N \sum_{\mathcal{L}} \frac{\sum_{k \in \mathcal{L}} y_k \beta_k}{\sum_{k \in \mathcal{L}} y_k} \right)' = N \left(\frac{\sum_{k \in \hat{\mathcal{L}}} y_k \beta_k}{\sum_{k \in \hat{\mathcal{L}}} y_k} \right)' = N \frac{\beta_j \sum_{k \in \hat{\mathcal{L}}} y_k - \sum_{k \in \hat{\mathcal{L}}} y_k \beta_k}{\left(\sum_{k \in \hat{\mathcal{L}}} y_k \right)^2} \tag{39}$$

Taking derivatives of both sides of equation (37) with respect to y_j and using the product rule and equations (38) and (39) and $exp(\sum_{\mathcal{L}} N\, log(\sum_{k \in \mathcal{L}} y_k)) = \prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k\right)^N$ gives

$$\begin{aligned}
&\left(\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k \right)^N \right) \left(\frac{N}{\sum_{k \in \hat{\mathcal{L}}} y_k} \right) N \sum_{\mathcal{L}} \frac{\sum_{k \in \mathcal{L}} y_k \beta_k}{\sum_{k \in \mathcal{L}} y_k} + \left(\prod_{\mathcal{L}} \left(\sum_{k \in \mathcal{L}} y_k \right)^N \right) N \frac{\beta_j \sum_{k \in \hat{\mathcal{L}}} y_k - \sum_{k \in \hat{\mathcal{L}}} y_k \beta_k}{\left(\sum_{k \in \hat{\mathcal{L}}} y_k \right)^2} \\
&\qquad = \sum_i \left(\prod_{k \neq j} y_k^{c_{ik}} \right) c_{ij} y_j^{c_{ij}-1} \hat{b}_i = \frac{1}{y_j} \sum_i \left(\prod_k y_k^{c_{ik}} \right) c_{ij} \hat{b}_i \quad .
\end{aligned}$$

Now let $y_k = p_k$ and use $\sum_{j\in\mathcal{L}} p_j\beta_j = 0$ and equations (20) and (22), and we get

$$N\beta_j = \frac{1}{p_j}\sum_i \bar{\varphi}_i c_{ij}\hat{b}_i \qquad \beta_j = \frac{1}{Np_j}\sum_i \bar{\varphi}_i c_{ij}\hat{b}_i \qquad \boldsymbol{\beta} = \tfrac{1}{N}\,\hat{\mathbf{b}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\tilde{\mathbf{p}}^{-1} \quad . \ \square$$

If we substitute $\boldsymbol{\alpha}$ for $\boldsymbol{\beta}$ in lemma 3, then equation (24) allows us to use the lemma to obtain

$$\boldsymbol{\alpha} = \tfrac{1}{N}\,\boldsymbol{\alpha}\,C^\top\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\tilde{\mathbf{p}}^{-1} \quad . \tag{40}$$

Then using equation (26) gives

$$\alpha_j = \frac{1}{Np_j}\sum_i \bar{\varphi}_i c_{ij}\hat{a}_i \qquad \boldsymbol{\alpha} = \tfrac{1}{N}\,\hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\tilde{\mathbf{p}}^{-1} \quad . \tag{41}$$

Compare this with equation (23).

9 Variances with Complete Recombination

By equations (23) and (41), we have

$$\mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,C = N\boldsymbol{\alpha}\,\tilde{\mathbf{p}} = \hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,C \quad . \tag{42}$$

Multiplying on the right by $\boldsymbol{\alpha}^\top$ gives $\mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,C\,\boldsymbol{\alpha}^\top = \hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\boldsymbol{\alpha}^\top$, which by equation (26) is $\mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top = \hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,\hat{\mathbf{a}}^\top$. Its transpose is $\hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\mathbf{a}^\top = \hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,\hat{\mathbf{a}}^\top$. So

$$\hat{\mathbf{a}}\tilde{\bar{\boldsymbol{\varphi}}}\hat{\mathbf{a}}^\top = \mathbf{a}\tilde{\boldsymbol{\varphi}}\hat{\mathbf{a}}^\top = \hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}\mathbf{a}^\top \quad . \tag{43}$$

If $\boldsymbol{\varphi}$ is at linkage equilibrium, we have $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$ and so

$$V_g = \hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top = \mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top = \hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\mathbf{a}^\top \quad , \tag{44}$$

(at linkage equilibrium) so

$$V_e = (\mathbf{a}-\hat{\mathbf{a}})\,\tilde{\boldsymbol{\varphi}}\,(\mathbf{a}-\hat{\mathbf{a}})^\top = \mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,\mathbf{a}^\top - \mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top - \hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\mathbf{a}^\top + \hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\hat{\mathbf{a}}^\top = V - V_g - V_g + V_g = V - V_g \quad .$$

Thus

$$V = V_g + V_e \tag{45}$$

at linkage equilibrium. We note from equations (26) and (42) that $\hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,\hat{\mathbf{a}}^\top = \hat{\mathbf{a}}\,\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\boldsymbol{\alpha}^\top = N\boldsymbol{\alpha}\,\tilde{\mathbf{p}}\boldsymbol{\alpha}^\top$. Thus

$$V_g = N\boldsymbol{\alpha}\,\tilde{\mathbf{p}}\,\boldsymbol{\alpha}^\top = N\sum_j p_j\alpha_j^2 \tag{46}$$

if we are at linkage equilibrium.

Notice that equations (44), (45), and (46) are true at linkage equilibrium, but are not true in general, as the example in section 7 shows.

10 Approximating Genotype Values

We can show a sense in which the genic values are the least mean square approximation to the values. We shall do the one locus case.

In this section, we assume one locus ($L = 1$) and linkage equilibrium ($\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$). We define

$$T = \tfrac{1}{N}\,C^\top\tilde{\bar{\boldsymbol{\varphi}}}\,C\,\tilde{\mathbf{p}}^{-1} = \tfrac{1}{N}\,C^\top\tilde{\boldsymbol{\varphi}}\,C\,\tilde{\mathbf{p}}^{-1} \quad .$$

Lemma 4 *If $L = 1$ and $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$, then $\check{\mathbf{e}}T = N\check{\mathbf{e}}$.*

Proof: $\check{\mathbf{e}}\,C^\top = N\mathbf{e}$, so using this and equation (18), we have $\check{\mathbf{e}}T = \boldsymbol{\varphi}\,C\,\tilde{\mathbf{p}}^{-1} = N\,\mathbf{p}\,\tilde{\mathbf{p}}^{-1} = N\check{\mathbf{e}}$. $\square$

Lemma 5 *If $L = 1$ and $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$, then T is nonsingular.*

Proof: Take any "short" real vector $\mathbf{u}$ of the same dimensionality as $\mathbf{p}$. We will show that $\mathbf{u}$ is in the range of the transformation T. Define $\bar{u} = \mathbf{p}\,\mathbf{u}^\top$ and $\boldsymbol{\beta} = \mathbf{u} - \bar{u}\check{\mathbf{e}}$. Now since there is only one locus set, $\sum_{j\in\mathcal{L}} p_j\beta_j = \sum_j p_j\beta_j = \mathbf{p}\,\boldsymbol{\beta}^\top = 0$, and so by lemma 3, we have $\boldsymbol{\beta}T = \boldsymbol{\beta}$. Using this and lemma 4 gives us $(\boldsymbol{\beta} + \frac{\bar{u}}{N}\check{\mathbf{e}})T = \boldsymbol{\beta}T + \frac{\bar{u}}{N}\check{\mathbf{e}}T = \boldsymbol{\beta} + \bar{u}\,\check{\mathbf{e}} = \mathbf{u}$. □

Now let $\mathbf{x}$ be a vector of proposed excess allele values and $\mathbf{y} = \mathbf{x}\,C^\top$. Then we have $y_i = \sum_j c_{ij}x_j$. The vector $\mathbf{y}$ will be the vector of genic values.

We want to minimise $Q = \sum_i \varphi_i(y_i - a_i)^2$. Now we have $\frac{\partial Q}{\partial x_j} = \sum_i \varphi_i 2(y_i - a_i)c_{ij}$ for all j. We want these partial derivatives to be zero. That is, for all j, we want $\sum_i \varphi_i(y_i - a_i)c_{ij} = 0$. That is, $(\mathbf{y}-\mathbf{a})\,\tilde{\boldsymbol{\varphi}}\,C = 0$. This means $\mathbf{y}\,\tilde{\boldsymbol{\varphi}}\,C = \mathbf{a}\,\tilde{\boldsymbol{\varphi}}\,C$. By the definition of $\mathbf{y}$ and equation (42), this becomes $\mathbf{x}\,C^\top\tilde{\boldsymbol{\varphi}}\,C = N\boldsymbol{\alpha}\,\tilde{\mathbf{p}}$. In other words, $\mathbf{x}\,C^\top\tilde{\boldsymbol{\varphi}}\,C\tilde{\mathbf{p}}^{-1} = N\boldsymbol{\alpha}$. So $\mathbf{x}\,T = \boldsymbol{\alpha}$. But by equation (40), we have $\boldsymbol{\alpha}T = \boldsymbol{\alpha}$, so $(\mathbf{x} - \boldsymbol{\alpha})T = 0$. By lemma 5, T is nonsingular, so $\mathbf{x} = \boldsymbol{\alpha}$. Consequently, by equation (26) we have $\mathbf{y} = \boldsymbol{\alpha}\,C^\top = \hat{\mathbf{a}}$.

11 Fisher's Theorem for Complete Recombination

We now let $\boldsymbol{\varphi}$ vary, using $'$ to represent time derivatives. From equation (18) we have

$$p_j{}' = \frac{1}{N}\sum_i \varphi_i{}' c_{ij} \qquad\qquad \mathbf{p}' = \tfrac{1}{N}\boldsymbol{\varphi}' C \quad . \tag{47}$$

For example, suppose we have a natural reward scheme, so equation (13) holds. From equations (47), (13), and (42), we have $\mathbf{p}' = \frac{1}{N}\boldsymbol{\varphi}'C = \frac{1}{N}K\mathbf{a}\tilde{\boldsymbol{\varphi}}C = \frac{1}{N}KN\boldsymbol{\alpha}\tilde{\mathbf{p}}$. In other words,

$$\mathbf{p}' = K\boldsymbol{\alpha}\tilde{\mathbf{p}} \qquad\qquad p_j{}' = Kp_j\alpha_j \tag{48}$$

if we have a natural reward scheme.

Now whatever our reward scheme we have the following.

From equation (22) we have $log(\bar{\varphi}_i) = \sum_j c_{ij} log(p_j)$, so

$$\frac{\bar{\varphi}_i'}{\bar{\varphi}_i} = \sum_j c_{ij}\frac{p_j{}'}{p_j} \quad ,$$

which gives

$$\bar{\varphi}_i' = \bar{\varphi}_i \sum_j c_{ij}\frac{p_j{}'}{p_j} \qquad\qquad \bar{\boldsymbol{\varphi}}' = \mathbf{p}'\tilde{\mathbf{p}}^{-1}C^\top\tilde{\bar{\boldsymbol{\varphi}}} \quad . \tag{49}$$

From equation (31), we have

$$\sum_i \bar{\varphi}_i' = 0 \quad . \tag{50}$$

We define the quantity

$$\dot{m} = \mu(\bar{\boldsymbol{\varphi}}) \quad .$$

By analogy with equation (8), the chain rule gives us

$$\dot{m}' = (\mu(\bar{\boldsymbol{\varphi}}))' = \sum_i \frac{\partial\mu(\bar{\boldsymbol{\varphi}})}{\partial\bar{\varphi}_i}\bar{\varphi}_i' \quad . \tag{51}$$

Now $\frac{\partial\mu(\bar{\boldsymbol{\varphi}})}{\partial\bar{\varphi}_i}$ is the partial derivative of μ in the i direction at the point $\bar{\boldsymbol{\varphi}}$. By contrast, $\frac{\partial\mu(\boldsymbol{\varphi})}{\partial\varphi_i}$ is the partial derivative of μ in the i direction at the point $\boldsymbol{\varphi}$. If the points $\bar{\boldsymbol{\varphi}}$ and $\boldsymbol{\varphi}$ happen to be the same, then $\dot{m} = \bar{m}$, and $\frac{\partial\mu(\bar{\boldsymbol{\varphi}})}{\partial\bar{\varphi}_i} = \frac{\partial\mu(\boldsymbol{\varphi})}{\partial\varphi_i} = m_i$, for all i.

Now let us assume linkage equilibrium.

That is, let us assume that at the current moment we have $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$. The change in $\boldsymbol{\varphi}$ causes a change in $\mathbf{p}$ and hence in $\bar{\boldsymbol{\varphi}}$, but of course $\bar{\boldsymbol{\varphi}}'$ may be quite different from $\boldsymbol{\varphi}'$. From what we said in the last paragraph, we see that our $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$ assumption means that equation (51) can be written

$$\dot{m}' = \sum_i \bar{\varphi}_i' m_i \quad . \tag{52}$$

By equations (5) and (50), $\sum_i \bar{\varphi}'_i m_i = \sum_i \bar{\varphi}'_i(\bar{m} + a_i) = \bar{m}\sum_i \bar{\varphi}'_i + \sum_i \bar{\varphi}'_i a_i = \sum_i \bar{\varphi}'_i a_i$. By this and equation (52) we have

$$\dot{m}' = \sum_i \bar{\varphi}'_i a_i = \mathbf{a}\,\bar{\boldsymbol{\varphi}}'^{\top} \quad . \tag{53}$$

If $\boldsymbol{\varphi}'$ is the change specified by the reward scheme, then $\bar{\boldsymbol{\varphi}}'$ and $\dot{m}'$ are the changes that actually occur if we have complete recombination.

From equations (49) and (47) we have $\bar{\boldsymbol{\varphi}}' = \frac{1}{N}\boldsymbol{\varphi}' C \tilde{\mathbf{p}}^{-1} C^{\top} \tilde{\boldsymbol{\varphi}}$, so $\bar{\boldsymbol{\varphi}}'$ is a linear function of $\boldsymbol{\varphi}'$. Similarly, from equations (26) and (23) we have $\hat{\mathbf{a}} = \frac{1}{N}\mathbf{a}\tilde{\boldsymbol{\varphi}} C \tilde{\mathbf{p}}^{-1} C^{\top}$, so $\hat{\mathbf{a}}$ is a linear function of $\mathbf{a}$. So we have $\mathbf{a}\,\bar{\boldsymbol{\varphi}}'^{\top} = \frac{1}{N}\mathbf{a}\tilde{\boldsymbol{\varphi}} C \tilde{\mathbf{p}}^{-1} C^{\top} \boldsymbol{\varphi}'^{\top}$, and we have $\hat{\mathbf{a}}\,\boldsymbol{\varphi}'^{\top} = \frac{1}{N}\mathbf{a}\tilde{\boldsymbol{\varphi}} C \tilde{\mathbf{p}}^{-1} C^{\top} \boldsymbol{\varphi}'^{\top}$, and from equations (23) and (47) we have $N\boldsymbol{\alpha}\mathbf{p}'^{\top} = \frac{1}{N}\mathbf{a}\tilde{\boldsymbol{\varphi}} C \tilde{\mathbf{p}}^{-1} C^{\top} \boldsymbol{\varphi}'^{\top}$. So from these last three equations we see that if $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$ then $\mathbf{a}\,\bar{\boldsymbol{\varphi}}'^{\top} = \hat{\mathbf{a}}\,\boldsymbol{\varphi}'^{\top} = N\boldsymbol{\alpha}\mathbf{p}'^{\top}$.

From this and equation (53) we have

$$\dot{m}' = \mathbf{a}\,\bar{\boldsymbol{\varphi}}'^{\top} = \hat{\mathbf{a}}\,\boldsymbol{\varphi}'^{\top} = N\boldsymbol{\alpha}\,\mathbf{p}'^{\top} \tag{54}$$

if we have linkage equilibrium.

So suppose we have a natural reward scheme, as given by equation (13). Then from that and equations (54) and (44) , we have $\dot{m}' = \hat{\mathbf{a}}\,\boldsymbol{\varphi}'^{\top} = K\hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}\,\mathbf{a}^{\top} = KV_g$. That is,

$$\dot{m}' = KV_g \tag{55}$$

if we have a natural reward scheme and linkage equilibrium.
This is Fisher's Theorem for the case of complete recombination. Compare with equation (15).

Before leaving these matters, it is worth noting that if we have a natural reward scheme and linkage equilibrium, then from equations (49), (48), (26), and $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$, we have
$\bar{\boldsymbol{\varphi}}' = \mathbf{p}'\,\tilde{\mathbf{p}}^{-1}\,C^{\top}\,\tilde{\boldsymbol{\varphi}} = K\boldsymbol{\alpha}\,\tilde{\mathbf{p}}\,\tilde{\mathbf{p}}^{-1}\,C^{\top}\,\tilde{\boldsymbol{\varphi}} = K\boldsymbol{\alpha}\,C^{\top}\,\tilde{\boldsymbol{\varphi}} = K\hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}} = K\hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}}$. In other words,

$$\bar{\boldsymbol{\varphi}}' = K\hat{\mathbf{a}}\,\tilde{\boldsymbol{\varphi}} \qquad\qquad \bar{\varphi}'_i = K\varphi_i\hat{a}_i \tag{56}$$

if we have a natural reward scheme and linkage equilibrium.
Compare this with equation (13).

Note that if we have a natural reward scheme, linkage equilibrium, and no epistasis (zero epistatic values), then $\hat{\mathbf{a}} = \mathbf{a}$ and by equations (13) and (56) we have $\bar{\boldsymbol{\varphi}}' = \boldsymbol{\varphi}'$.

12 Fisher's Wording

Fisher's statement of his theorem (see section 1) is ambiguous as to what sort of recombination he assumes and as to whether he assumes linkage equilibrium. If we look at his proof, we see that in places he assumes linkage equilibrium and in places he doesn't. For example, he is very careful to make a distinction between what he calls average excess and what he calls average effect.[6]

Suppose we look at the example in section 7. The *average excess* of allele A (over allele B) is $\alpha_1 - \alpha_2$, which is 6. The *average effect* of allele A is the average of the change in value if you start with an individual that has a B at locus 1 and replace the B with an A. So in the example, since every individual that has a B is of genotype BD, the average effect of A is 100.

Now it is clear that at linkage equilibrium, an allele's average effect is the same as its average excess, so clearly Fisher is not assuming linkage equilibrium, and his wording of the theorem does not imply equilibrium.

But then, Fisher's proof seems to contain errors. On [6, page 32], he defines random variables x and X. The variable x seems to be what we have called m_i and the variable X seems to be what we have called $\bar{m} + \hat{a}_i$. He then says, "It follows that the sum for the whole population of the product $X(x - X)$ derived for each individual must be zero, ..." He seems to be claiming that the quantity $(\bar{m}\mathbf{e} + \hat{\mathbf{a}})\tilde{\boldsymbol{\varphi}}(\mathbf{m} - (\bar{m}\mathbf{e} + \hat{\mathbf{a}}))^{\top}$ is zero. If we use $\mathbf{e}\tilde{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$ and $\mathbf{m} - \bar{m}\mathbf{e} = \mathbf{a}$, the quantity becomes $\bar{m}\boldsymbol{\varphi}(\mathbf{a} - \hat{\mathbf{a}})^{\top} + \hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}(\mathbf{a} - \hat{\mathbf{a}})^{\top}$. By equations (6) and (27), the quantity is simply $\hat{\mathbf{a}}\tilde{\boldsymbol{\varphi}}(\mathbf{a} - \hat{\mathbf{a}})^{\top}$. He is claiming this must be zero. But in our example in section 7 it is -18.

Of course the quantity is zero at linkage equilibrium, as equation (44) shows.

[6] Crow and Kimura [2, page 210] say "Fisher's [4, 6] treatment of the subject is recondite. A clearer discussion of the circumstances under which the principle [Fisher's theorem] works was given in Fisher [5]. For a straightforward general derivation, see Kimura [15]." In Kimura [15], he evidently determines average *effect* by least mean square approximation. I haven't read [15].

On [6, page 34], Fisher says "The genetic variance as here defined is only a portion of the variance determined genotypically, and this will differ from, and usually be somewhat less than, the total variance to be observed." The wording is typically confusing, but if it means that genic variance is never greater than total variance ($V_g \leq V$), as it seems to, then it is wrong. It is true at linkage equilibrium, as equation (45) shows, but it is not true in the example in section 7.

Of course by "genetic variance" he might have meant something other than the variance of the excess genic values $\hat{a}_i$. These values were sums of allele excess values (see equation (26)). Perhaps he meant us to use sums of allele average effects. We can see that the average effects of alleles A, B, C, and D in the example are 100, -3, 3, 94, respectively. The analogue of $\hat{a}_1$, $\hat{a}_2$, $\hat{a}_3$, $\hat{a}_4$, using average effects, is 103, 194, 0, 91, respectively. The variance of these, remembering that half of the population is genotype 1 and half is genotype 4, is 9445. No, surely he didn't mean that variance.

Later, on page 34, Fisher re-defines x and X, and proceeds as if he is not assuming linkage equilibrium. But then on page 36 he says, "It follows that the variance of the genetic value ξ is equal to the covariance of genetic and genotypic values ξ and x." Again, it looks like an assertion of equation (44), which we have seen fails unless we have linkage equilibrium. Unless of course he wants us to use the average excesses, in which case his variance is 9445, as we have seen, and his covariance is $\frac{1}{2}(3)(103) + \frac{1}{2}(-3)(91)$, which is 18. No, he couldn't mean that.

Either Fisher's reasoning or his exposition is faulty.

13 Problems with Fisher's Approach

Having said this, we need to admit some of the problems with Fisher's approach. The most glaring problem is that Fisher's theorem requires a natural reward scheme, and implementing such a scheme requires a way of estimating the m_i's. This may not be as easy as it seems. A cannibal genotype may have very low m_i, but its reproduction rate will be proportional to the resources it acquires, and this may be very high. The correct m_i is something like the marginal utility of genotype i. Correct estimating procedures may require either group selection [24] or something analogous to capitalist pricing theory. Holland has employed the capitalist approach in developing his bucket brigade [13] for classifier systems. Whatever the approach, it relies ultimately on something associated with genotype i that is not affected too much by the other genotypes. Some restriction on genotype interaction is required.

There is a second problem with Fisher's approach. This stems from the finiteness of our populations. Even if we do have a way of estimating the m_i's, Fisher's theorem requires that reproductive rate be exactly proportional to the m_i's (or more generally, be exactly $Ka_i+\chi$). In practice, this requires a huge population, for otherwise the reproductive rate will be subject to sampling error noise. To see how large a population might be required, we need only reflect that the analysis assumes each φ_i is positive. Thus the population should have at minimum as many individuals as there are possible genotypes. This is ridiculous. No feasible population is anywhere near that large.

But if we look at Fisher's theorem for complete recombination, and restrict ourselves to populations at linkage equilibrium, we can view all quantities as functions of the p_j's, and we see that what we need is to have every p_j positive. Now a population that contains every possible allele has every p_j positive, and this is achieved in most modest size populations. The problem is that even with a high recombination rate, such populations do not achieve exact linkage equilibrium. If they did, every φ_i would be positive, since every $\bar{\varphi}_i$ is assuredly positive if the population contains every allele.

So we see that in any real population, sampling error noise makes both the reproductive rates and the population structure different from what the analysis requires.

So the population does not evolve deterministically as predicted by Fisher's theorem. The noise introduces a randomness, and the population $\boldsymbol{\varphi}$ travels through the positive orthant not along a deterministic smooth track, but in tiny probabilistic jumps that occur when individuals are born or die. In modeling populations, we assume (as do evolutionary biologists and ecologists) that the actual track of $\boldsymbol{\varphi}$ is near the deterministic track, in this case the track that Fisher predicts, and that the larger the population (the larger the sample size) the less the noise and the closer the population will follow Fisher's track. For many years there was no proof in the literature of the convergence to the deterministic track as the population size increased, (neither for this case nor for the analogous ecological models). Biologists simply assumed the convergence. Trevor Fenner and I have sketched a convergence proof that we are convinced will work when it is tidied up. Probably there are now convergence proofs in the literature.

In any case, Fisher gives us an expected direction for the movement of $\boldsymbol{\varphi}$, and deviations from that direction are due to unbiased noise.

Of course this argument is assuming complete recombination, or at least very high recombination rate, and that is not the situation we are interested in. Nevertheless, if we look at short schemata rather than alleles, and view all quantities as functions of the probabilities of the short schemata, then we can have a recombination rate just high enough that the short schemata are linked only loosely to each other, yet low enough that the short schemata are comparatively undisrupted. Of course this is the situation we are interested in. The short schemata overlap, and the situation is messy, but the hope is that an approximate analogue of Fisher's theorem for complete recombination will hold, but with short schemata in place of alleles.

The third problem with Fisher's approach has to do with the amount of epistasis. Just as we need a bound on genotype interaction, so too do we need a bound on interaction between alleles in the same individual. Suppose we have a conceptual catalog that lists all the possible genotypes and suppose to each genotype we assign a random integer number between 0 and 100000, each integer being equally likely to be chosen. Suppose we enter these numbers in the catalog. For each genotype i, the number ν_i will be the number entered for i in the catalog. We let $\mu(\boldsymbol{\varphi}) = \sum_i \varphi_i \nu_i$. Then whatever the population, for each i we have $m_i = \nu_i$, the number in the catalog. Fisher's theorem says that with complete recombination, $\bar{m}$ increases at a rate proportional to V_g. So presumably a genetic algorithm will work on a population of these genotypes, in the sense that $\bar{m}$ will increase.

But of course the genetic algorithm will actually fail. The increase in $\bar{m}$ will be microscopic. The difference in genotype values is almost entirely due to epistasis. In a linkage equilibrium population, V_g is tiny and V_e is almost equal to V. The alleles' values are almost identical and selection for better alleles accomplishes virtually nothing. Similarly selection for better short schemata by a genetic algorithm accomplishes virtually nothing, so genetic algorithm linkage doesn't help.

So we see that some restriction on the amount of epistasis is required for evolution to work. Strong interaction between the effects of genes closely linked is okay, but if there is too much interaction between weakly linked genes, a genetic algorithm cannot work. It may be that a Walsh Transform analysis can separate the nearby interactions from the more distant interactions, but I have not investigated this. The hope would be that such an analysis could extend Fisher's theorem to systems with a linkage pattern like that produced by crossover. But even without such an extension, the conjecture that $\bar{m}$ increases at a rate proportional to some number between V_g and V seems reasonable. The problem here is that with high epistasis and low V_g, that number could be tiny.

But provided we assume some restrictions on interaction between genotypes in a population and on interaction between alleles in an individual, Fisher's analysis makes sense, and evolution and genetic algorithms presumably work.

14 The Proportionality Question

In practice, genotype values m_i are expressions of payoff received from the environment. In Genetic Algorithms, reproductive rates are set on the basis of this payoff. But there are many ways of doing this. A natural reward scheme is only one way. Which ways are good and which bad? Let's state the question as follows.

For each i, let r_i be the excess reproductive rate of genotype i. That is, r_i is the reproductive rate of genotype i minus the weighted average of the reproductive rates of all the genotypes (weighted by proportions $\boldsymbol{\varphi}$). Of course we have $\sum_i \varphi_i r_i = 0$. If for the moment we let $\bar{R}$ be the average reproductive rate, then for asexual reproduction we have $s_i' = s_i(\bar{R} + r_i)$. Therefore $\bar{s}' = \bar{s}\bar{R}$, and since $\varphi_i = \frac{s_i}{\bar{s}}$, the quotient rule tells us that $\varphi_i' = \varphi_i r_i$. We see that excess reproductive rate r_i is $\frac{\varphi_i'}{\varphi_i}$, or if we prefer, $(\log(\varphi_i))'$. We shall take that as the definition of excess reproductive rate. Whatever the reproductive rates, equation (9) tells us that $\bar{m}' = \sum_i \varphi_i r_i a_i$. The question is, what should the reproductive rates r_i be?

Of course one possible answer is that the reproductive rates should be determined by a natural reward scheme. Remember that a reward scheme is natural if and only if equation (13) holds. Then by the definition of excess reproductive rate, a reward scheme is natural if and only if for all i we have $r_i = Ka_i$, for some non-negative K that is independent of i. Let $\mathbf{r}$ be the vector of reproductive rates, so for any i, the i'th entry in $\mathbf{r}$ is r_i. So the natural reward schemes are those in which $\mathbf{r}$ is a scalar multiple of the vector $\mathbf{a}$ (non-negative scalar). This is a useful alternative definition of "natural reward scheme".

The literature contains a long history of attempts to show that $\mathbf{r}$ should be determined by a natural reward scheme, that $\mathbf{r}$ should be a scalar multiple of $\mathbf{a}$, that excess reproductive rate should be proportional to excess value. This proportionality (which Holland calls "fitness proportional selection") is seen by Holland as having advantages, and Kimura tries to outline an advantage of it. [2] (Note that if reproductive rate is

proportional to value then excess reproductive rate must be proportional to excess value, but the implication doesn't hold in reverse.)

I am going to outline an advantage of the proportionality that I have not seen reported elsewhere. The argument is not compelling, but I find it persuasive.

Let us be clear what we do and do not mean by proportional reproductive rates. Proportionality means $\mathbf{r} = K\mathbf{a}$ for some scalar K. But K need not be a constant. It could depend on time, on $\boldsymbol{\varphi}$, or on anything else. Note that what Genetic Algorithms people call "fitness scaling" destroys proportionality.

15 Noise

So now let's return to the question. What should the vector $\mathbf{r}$ be? Well, $\boldsymbol{\varphi}$ is traveling over the population domain, and we want it to move uphill on the function μ. With no recombination the rate of climb is $\bar{m}'$. With complete recombination, the rate of climb is $\dot{m}'$, as given by equation (53), or equation (54).

Let's go back and ask what we can get out of Fisher's Theorem. It tells us that if we use proportional reproductive rates then the rate of change of average value (the climb rate) will be some variance or other. And a variance is positive (unless the ground is level) so we are going in a direction that takes us uphill.

But this is not the only uphill direction. If you are standing on a hillside you have a full 180 degrees of uphill directions. Remember that $\bar{m}' = \sum_i \varphi_i r_i a_i$, so if we look at excess values a_i and excess reproductive rates r_i then as long as for each genotype the two are the same sign, we head uphill in the asexual case. Scaling of excess reproductive rate seems to work fine. And we can even head down in some coordinates, provided we head up in others.

But what an engineer might ask is this. Okay, all these directions are uphill, but which takes us uphill fastest? The engineer would be attracted to the steepest ascent direction. Proportional reproductive rates do not give us steepest ascent, but we have seen in section 5 that if we scale the axes by taking the square root then proportional reproductive rates do give us steepest ascent. What does this tell us?

Very little. The problem is obviously this. Once we stop talking about trying to ensure we are going uphill and start trying to increase the rate at which we climb, we run into the obvious formal problem. If we are going uphill, and we want to double the rate at which we climb, it's trivial. Just double all the excess reproductive rates. If we double reproductive rates, we double each r_i, double each φ_i', double the length of $\boldsymbol{\varphi}'$, and double $\bar{m}'$. By equations (47), (49), and (53), we also double the length of $\bar{\boldsymbol{\varphi}}'$ and double $\dot{m}'$. Thus without changing the direction we are moving, we double climb rate by doubling speed. In effect, we have simply doubled the learning rate. This is not biology, it's engineering, and we, like god, can simply double the rate at which things change.

For biologists looking at rate of increase in average value, the fact that in biological populations it's exactly the size of the variance of the values is important. For Genetic Algorithms researchers, who can change at will the rate of learning, it's only the fact that the variance is positive that is important. Its size is irrelevant because we can always increase the rate of evolution by multiplying all the excess reproductive rates by the same constant of proportionality.

So why don't we do that anyway? Because it increases the sampling error noise that we discussed in section 13. If we double the rate of evolution (double all the excess reproductive rates) then this doubles the standard deviation of the noise, multiplies the noise variance by four. But notice that the noise is unbiased. Our analysis has been using average values, and the noise is the difference between the true values and these average values. So the noise is unbiased. So the noise will all average out in the end and cause no damage.

But of course that's not true. One reason is that in a finite population, such noise can cause extinctions of good schemata, what is usually called premature convergence. I'm going to assume that the standard deviation of the noise in the reproductive rate is a good measure of this danger.

As an illustration, let's assume we are using a natural reward scheme. This is of course much discussed in the biological literature, for example in Kimura [2].

As Kimura points out, the variance in payoff received by genotypes can often be thought of as composed of three variances, the genic variance, the epistatic variance, and the environmental variance. Environmental variance is the variance of values that vary because the environment is moving from state to state. In our analysis we pretended that environmental variance was zero and that at any point $\boldsymbol{\varphi}$, each genotype value was fixed and equal to m_i. In practice, the environment state[7] changes in a pattern influenced by $\boldsymbol{\varphi}$, the genotype values fluctuate, and m_i is merely the average value of genotype i at $\boldsymbol{\varphi}$. Environmental variance is a measure of the fluctuations. In our analysis above, V was all the variance except environmental variance. Loosely speaking, it included genic variance and epistatic variance.

[7]The location of the organism in the environment is part of the environment state.

Kimura talks as if you can simply add genic variance, epistatic variance, and environmental variance. (as in equation (45)) This is not true in general, but let's for the moment talk loosely as if it were. The point is this. Some of that variance is useful and some just causes trouble. Variance is useful if the reward scheme is built to take advantage of it. Let me illustrate, still assuming we are using a natural reward scheme.

If reproduction is asexual, then the environmental variance is useless to us, but the rest of the variance is useful, and the rate of change of average value can use that useful variance, as in equation (15) for natural reward schemes. On the other hand, suppose we have complete recombination and consequently the population has Hardy Weinberg ratios. Then that useful variance is in two bits, the genic variance and the epistatic variance. But now that's not all useful. Equation (15) becomes equation (55). Only the genic variance is useful, and the epistatic variance as well as the environmental variance is useless.

Suppose we select a genotype and test it against the environment to get payoff. The payoff we get is due to three things, the alleles present (their values), the epistasis, and the environment (the state the environment happens to be in when we test). The three variances are respectively the variances due to the variation of these three things. With complete recombination, only the first is useful, and the other two are noise, whereas with no recombination, the first two are useful and only the third is noise.

I'm going to simplify the discussion by assuming that environmental variance is zero as we have done in the previous sections, but I think the discussion can be extended to the nonzero environmental variance case. In this discussion, in a given population a genotype gives the same payoff every time it is tested, and that payoff is the value of the genotype, the partial derivative. (There are no cannibals or freeloaders.) So with no recombination there is no noise. With complete recombination, the epistatic variance is the noise. If there is complete recombination, the system merely learns which alleles are good, and doesn't learn which combinations of alleles are good. If a genotype is good merely because of epistasis, the system will never learn that and the epistasis merely provides noise to a system trying to learn which alleles are good.

That discussion of noise was assuming a natural reward scheme, so it was really only illustrative. How much of what we said holds in general? Let's start over, now not making any assumption as to what reward scheme we are using.

We want to discover what reward scheme is good. We want to discover what the excess reproductive rate vector $\mathbf{r}$ should be.

Suppose we have the following simple situation. We have complete recombination and zero environmental variance. The population is at Hardy Weinberg Equilibrium.

Let's define for each locus set $\mathcal{L}$ the variance of the allele values at that locus. $V_{\mathcal{L}} = N \sum_{j \in \mathcal{L}} p_j \alpha_j^2$ We will call that the locus variance. Then by equation (46) we have $V_g = \sum_{\mathcal{L}} V_{\mathcal{L}}$.

The genic variance is the sum of n different locus variances, one for each locus. The variances all add nicely. The total variance is the sum of these n locus variances and the epistatic variance.

All these quantities are defined from the genotype values m_i. For our point $\boldsymbol{\varphi}$, these genotype values are given.

Now the game is as follows. We want for each genotype i to specify an excess reproductive rate r_i. That is, we want to specify an excess reproductive rate vector $\mathbf{r}$. The idea is to select $\mathbf{r}$ to get high rate of increase in average value (high climb rate) and low noise. Suppose we have a candidate vector $\mathbf{r}$.

Now the r_i's are numbers just as the a_i's are, and we can use them to calculate analogues of total variance, genic variance, and epistatic variance, just as if the r_i's were the excess genotype values. These new quantities will of course be different from the old quantities, which we calculated using the excess values rather than the excess reproductive rates. The new variances too are additive, since we have complete recombination and Hardy Weinberg ratios.

We need some terminology to distinguish the quantities based on excess values a_i from the analogous quantities based on the excess reproductive rates r_i. We will distinguish them by the prefixes value- and rate- respectively. Thus we have a value-genic-variance and a rate-genic-variance, for example.

If we are worried about extinctions and premature convergence, it is the rate-variances that produce the dangerous noise. Look for example at locus 1. Here the useful variance is the rate-locus-1-variance. This variance drives the adaptation at locus 1. The rate-epistatic-variance just provides dangerous unbiased noise at locus 1. And so does the rest of the rate-genic-variance, the rate-locus-variances for the other loci. Most of the rate-variance is just causing trouble.

So we want to choose $\mathbf{r}$ to get the rate-total-variance as low as we can, since the rate-total-variance is mostly just noise. Define $W^2 = \sum_i \boldsymbol{\varphi}_i r_i^2$, so W^2 is the rate-total-variance. I said I would assume that the standard deviation of the noise is a good measure of the danger noise can do. So W is a good measure of the noise danger. I shall call W the *noise measure*. We want to get W as low as we can.

But of course if W is zero then evolution will stop and the climb rate $\dot{m}'$ will also be zero, so W must not be too low. But it is clear that if two different $\mathbf{r}$'s give the same W but different climb rates $\dot{m}'$, then

we prefer the $\mathbf{r}$ with the higher climb rate. For any given value of W, we prefer the $\mathbf{r}$ with the highest climb rate.

Define the *effectiveness ratio* to be the ratio of the climb rate to the square root of the rate-total-variance. Here, the effectiveness ratio is $\frac{\dot{m}'}{W}$. It is something like a signal to noise ratio. We said that if we leave $\mathbf{r}$ pointed in the same direction but double its length, then this doubles both $\bar{m}'$ and $\dot{m}'$. It also doubles W, so the doubling has no effect on the effectiveness ratio. The effectiveness ratio depends on the direction of $\mathbf{r}$, but is independent of its length.

What we want to do is find the $\mathbf{r}$ direction that gives the largest effectiveness ratio. Then for any given value of W, an $\mathbf{r}$ aimed in that direction and of a length to give the rate-total-variance of W^2 will give the largest climb rate possible with that value of W. (And for any given climb rate value $\dot{m}'$, an $\mathbf{r}$ aimed in that direction and of a length to give a climb rate of $\dot{m}'$ will give the *smallest* rate-total-variance W^2 possible with that value of the climb rate.) We shall show that this desirable direction is the same as the direction in which $\hat{\mathbf{a}}$ points.

16 The Asexual Case

Let's begin by looking at the simpler asexual case. Here the climb rate is $\bar{m}'$, and the effectiveness ratio is $\frac{\bar{m}'}{W}$. Of course in the asexual case, what the rate total variance W^2 measures is all useful variation, not noise, so W is not a measure of noise and the effectiveness ratio is not a signal to noise ratio. It is merely a measure of how well the variation produces climb. Like the complete recombination effectiveness ratio, the asexual effectiveness ratio too is independent of the length of $\mathbf{r}$. We now ask which $\mathbf{r}$ direction gives the largest asexual effectiveness ratio.

As in section 5, we look at the square root space. We have $\boldsymbol{\varphi}$ traveling over the population domain in our original space, and its image $\boldsymbol{\psi}$ traveling over a portion of the unit sphere in the square root space. We have $\varphi_i = \psi_i^2$ for all i. We have the value function μ over the original space and the corresponding value function f over the square root space. The value functions are defined over the whole positive orthant. If $\boldsymbol{\psi}$ is the image of $\boldsymbol{\varphi}$ then $f(\boldsymbol{\psi}) = \mu(\boldsymbol{\varphi})$.

At any time, the velocity of $\boldsymbol{\psi}$ as it travels over the square root space is the velocity vector $\boldsymbol{\psi}'$. Its speed of travel is the length of this vector. I will write σ to mean the speed of travel. So $\sigma^2 = \sum_i (\psi_i')^2$.

Now $\varphi_i = \psi_i^2$. So we see that $\varphi_i r_i = \varphi_i' = 2\psi_i\psi_i'$. If we divide this equation by ψ_i, we get $\psi_i r_i = 2\psi_i'$. If we now square both sides, we get $\varphi_i r_i^2 = 4(\psi_i')^2$. If we now sum over i, we obtain $W^2 = 4\sigma^2$. So we see that $W = 2\sigma$.

So the effectiveness ratio is half of the climb-speed ratio $\frac{\bar{m}'}{\sigma}$. The numerator of this climb-speed ratio is the rate of climb of $\boldsymbol{\varphi}$ in the original space, which is of course the rate of climb of $\boldsymbol{\psi}$ in the square root space. The denominator is the speed of $\boldsymbol{\psi}$ over the ground in the square root space. Like the effectiveness ratio, the climb-speed ratio is independent of the length of $\mathbf{r}$ and depends only on the direction of $\mathbf{r}$. We ask which direction of $\mathbf{r}$ gives the largest climb-speed ratio.

The direction of $\mathbf{r}$ determines the direction of $\boldsymbol{\psi}'$, the direction in which $\boldsymbol{\psi}$ is moving. Clearly we get the largest climb-speed ratio if and only if $\boldsymbol{\psi}$ is taking the path of steepest ascent in the square root space. In section 5 we saw that if the ground is not level then there is one unique steepest ascent direction in the square root space, and that $\boldsymbol{\psi}'$ points in that direction if and only if we have a natural reward scheme, that is, if and only if $\mathbf{r}$ is a scalar multiple of $\mathbf{a}$. So $\mathbf{r}$ must point in the same direction as $\mathbf{a}$ for maximum effectiveness ratio. If $\mathbf{r}$ points in any other direction, the effectiveness ratio will be less.

So in the asexual case, the effectiveness ratio is at its maximum when $\mathbf{r}$ is a positive scalar multiple of $\mathbf{a}$. In that case $\mathbf{r} = K\mathbf{a}$ for some positive K. Then we have $\varphi_i' = \varphi_i r_i = K\varphi_i a_i$, and equation (10) becomes $\bar{m}' = \sum_i \varphi_i' a_i = K\sum_i \varphi_i a_i^2 = KV$. (equation (15))
Furthermore, $W^2 = \sum_i \varphi_i r_i^2 = K^2 \sum_i \varphi_i a_i^2 = K^2 V$. So $\frac{\bar{m}'}{W} = \sqrt{V}$. So to sum up:

Lemma 6 *In the asexual case, anywhere the ground is not level we have the following: the effectiveness ratio equals the standard deviation of the genotype values if the excess reproductive rates are proportional to the excess genotype values, and it is less than that otherwise.*

17 The Complete Recombination Case

Suppose we have complete recombination. Then the rate of change of $\bar{m}$ is $\dot{m}'$. We must have linkage equilibrium, so by equation (54) we have $\dot{m}' = \hat{\mathbf{a}}\,\boldsymbol{\varphi}'^{\top}$.

Note that here the effectiveness ratio is $\frac{\dot{m}'}{W}$, where W is the square root of $\sum_i \varphi_i r_i^2$.

Now $\boldsymbol{\varphi}$ is our current population. We want it to climb the value function μ.

We could define a different value function $\dot{\mu}$ to be $\dot{\mu}(\boldsymbol{\vartheta}) = \bar{m} + (\boldsymbol{\vartheta} - \boldsymbol{\varphi})\mathbf{m}^\top$, where $\mathbf{m}$ is the vector whose i'th entry is m_i, the partial derivative at $\boldsymbol{\varphi}$ of $\mu(\boldsymbol{\varphi})$ with respect to φ_i . (see equation (2)) So $\dot{\mu}$ is simply the linear approximation to μ near $\boldsymbol{\varphi}$. Climbing this function is exactly the same problem near $\boldsymbol{\varphi}$ as climbing μ.

But more interesting is the approximation $\ddot{\mu}$ defined by $\ddot{\mu}(\boldsymbol{\vartheta}) = \bar{m} + (\boldsymbol{\vartheta} - \boldsymbol{\varphi})(\bar{m}\mathbf{e}^\top + \hat{\mathbf{a}}^\top)$. We see that $\ddot{\mu}(\boldsymbol{\varphi}) = \mu(\boldsymbol{\varphi})$. More interestingly, $\frac{\partial}{\partial \vartheta_i}\ddot{\mu}(\boldsymbol{\vartheta}) = \bar{m} + \hat{a}_i$.

We look now at what certain quantities would be if our value function were $\ddot{\mu}$ rather than μ and our current population were still $\boldsymbol{\varphi}$.

$\bar{m}$ would still be $\bar{m}$.
m_i would be $\bar{m} + \hat{a}_i$.
a_i would be $\hat{a}_i$.

(In other words, the μ excess value vector is $\mathbf{a}$, but the $\ddot{\mu}$ excess value vector is $\hat{\mathbf{a}}$. In effect, the switch to $\ddot{\mu}$ removes the epistasis.)

V would be V_g.

So the $\ddot{\mu}$ ground is level if and only if $V_g = 0$.
The asexual climb rate $\bar{m}'$ would be $\sum_i \varphi_i' \hat{a}_i$ instead of $\sum_i \varphi_i' a_i$. That is, $\bar{m}'$ would be $\hat{\mathbf{a}}\boldsymbol{\varphi}'^\top$. Or to put it another way, (by equation(54))

$\bar{m}'$ would be $\dot{m}'$.

Whatever reproduction rates $\mathbf{r}$ we choose, the asexual climb rate on $\ddot{\mu}$ at $\boldsymbol{\varphi}$ is identical to the complete recombination climb rate on μ at $\boldsymbol{\varphi}$.

The definition of W is independent of the value function, so the asexual effectiveness ratio using $\ddot{\mu}$ at φ is the complete recombination effectiveness ratio using μ at φ.

In the previous section, we showed that if the $\ddot{\mu}$ ground is not level ($V_g \neq 0$) then this effectiveness ratio is highest when the excess reproductive rates $\mathbf{r}$ are proportional to the $\ddot{\mu}$ excess values, which are $\hat{\mathbf{a}}$. So we see that the effectiveness ratio under complete recombination is highest when $\mathbf{r}$ points in the same direction as $\hat{\mathbf{a}}$ and is lower than that if $\mathbf{r}$ points in any other direction. If $\mathbf{r}$ points in the same direction as $\hat{\mathbf{a}}$, then we have $\mathbf{r} = K\hat{\mathbf{a}}$ for some scalar K. Then, $\dot{m}' = \hat{\mathbf{a}}\boldsymbol{\varphi}'^\top = \sum_i \varphi_i' \hat{a}_i = \sum_i \varphi_i r_i \hat{a}_i = K \sum_i \varphi_i \hat{a}_i^2 = KV_g$, and $W^2 = \sum_i \varphi_i r_i^2 = K^2 \sum_i \varphi_i \hat{a}_i^2 = K^2 V_g$. So the effectiveness ratio is $\sqrt{V_g}$.

To summarize:

Lemma 7 *In the complete recombination case, anywhere $V_g \neq 0$ we have the following: The effectiveness ratio equals the standard deviation of the genic values if the excess reproductive rates are proportional to the excess genic values, and is less than that otherwise. In other words, the effectiveness ratio is $\sqrt{V_g}$ if the vector $\mathbf{r}$ is a positive scalar multiple of $\hat{\mathbf{a}}$, and is less than that otherwise.*

18 What Does This Tell Us?

The first thing we see is that all this tells us virtually nothing about the learning rate K and how big it ought to be. The learning rate K is in effect a measure of sample size, how many samples we take for a given amount of change we make. The smaller K is, the larger the sample size. And we know that a larger sample size reduces noise but slows the changes. It does so in exactly the way we see here. As far as K is concerned, our analysis tells us precisely what we knew already.

The second thing we see is that Lemma 7 doesn't discuss the $V_g = 0$ case. But from equation (46) we see that this case consists only of situations where every allele has zero excess value, only of situations where one allele is just as good as another. If that's the case and there is complete recombination, then it doesn't matter what scheme we use.

But all this does seem to tell us that we should arrange that our excess reproductive rates are proportional to excess genic values. This is pleasing, but also a little bit worrying because that is not what we usually do. Usually we make our excess reproductive rates proportional to excess values. (We make $\mathbf{r}$ a scalar multiple of $\mathbf{a}$, not of $\hat{\mathbf{a}}$.) We use a_i rather than $\hat{a}_i$ for the obvious reason that calculating $\hat{a}_i$ is a big mess, and a_i falls out easily and implicitly in the natural reproductive scheme. Let's talk about this discrepancy between what we do and what we ought to do.

Let's continue with the complete recombination case. Suppose there were no epistasis. Then $\hat{\mathbf{a}} = \mathbf{a}$, and we see that what we do is exactly what we ought to do. In this case, $\mathbf{r} = K\mathbf{a} = K\hat{\mathbf{a}}$, and since the $\hat{a}_i$'s

are linear combinations of the allele values, the r_i's are linear combinations of the allele reproductive rates. Let's look at this no-epistasis case from the point of view of an allele.

Suppose that there are 50 loci and that the value-variance at each locus is about the same. Suppose there are three alleles that can be at locus 1. Let's look at things from the point of view of one of those alleles. I will call that allele "our allele". Suppose we are doing what we are supposed to do and setting excess reproductive rates proportional to excess values. To simplify, suppose we set K equal to 1, so we have $\mathbf{r} = \mathbf{a} = \hat{\mathbf{a}}$. Then the excess reproductive rate of our allele equals its excess value, which is what we want. But what about the noise? That allele would like the noise to be zero. That is, it would prefer that every time an individual containing our allele is tested, the number of excess children it gets should be exactly the excess value of our allele. All the individuals containing our allele should have the same reproductive rate, and it should be the excess value of our allele.

But it isn't. Some of these guys get more children than that and some get less, depending on what alleles they have at other loci. That noise is unfortunate and threatens our allele with unjustified extinction, but if we eliminate that noise to please our allele, then we get no evolution at the other loci and we lose the opportunity the other loci give us for further improvement. There is a trade off. We want to use the information at the other loci, but that gives us noise at the first locus. We want to use that information some, but not too much. Our analysis here tells us that the optimum is achieved when we use it just enough to make the excess reproductive rates proportional to excess values.

The point is that the value difference information at the other loci is both useful and destructive. We want to use enough of it to benefit from its usefulness, but not so much that we suffer a lot from its destructiveness. It turns out that there is an optimal proportion of it that we should use, an optimal K for those loci, and the optimal K is whatever K we are using for the first locus.

Now suppose we have epistasis. Now there is even more information at the other loci and we would expect it too to be both useful and destructive. But actually, because we have complete recombination, this epistatic enformation is entirely useless because the recombination will destroy any use to which we put it. So clearly we want to use none of this information if possible. We want $\mathbf{r}$ to be a scalar multiple of $\hat{\mathbf{a}}$, not $\mathbf{a}$.

The analysis here doesn't give us any reason to make $\mathbf{r}$ a scalar multiple of $\mathbf{a}$ rather than $\hat{\mathbf{a}}$, but here are a couple of qualitative comments. I have already mentioned that one reason might be that $\hat{a}_i$ is difficult to calculate. Another reason might be that if there is any linkage, that is if the recombination is not complete, then a lot of the epistatic information is actually useful and we wouldn't want to ignore it.

I will leave this matter now. We have an argument for proportionality, but whether we want the r_i's proportional to the $\hat{a}_i$'s or the a_i's is a more complicated issue.

So we seem to have produced a persuasive, if not compelling, argument for proportionality. On the face of it, it provides an argument against what GA people call "fitness scaling". Let's take a look and see if we really should avoid that scaling.

I had the pleasure of being involved in discussions with Yuval Davidor while he was working on his PhD. He was using a genetic algorithm to evolve a good program to move a robot arm along a target path. He used some rather violent "fitness scaling", that is, his reproductive rates were not proportional to what I would call value. I thought then and still think that the work was excellent and very interesting and that his "fitness scaling" was perhaps appropriate. But let's pretend for a moment that I had tried to talk him out of his "fitness scaling". The conversation might have gone something like this.

> **Yuval:** You say, Tom, that you are worried about sampling error noise because it might cause random loss of good alleles. But there are many copies of each allele in the population. It would take an awful lot of noise to cause such a loss.
>
> **Me:** But my comment about allele loss during complete recombination was only illustrative. In practice we have linkage, and it is random loss of good short schemata that is the problem. The number of copies of a short schema in the population might not be large.
>
> **Yuval:** But if a short schema is lost it can be regenerated by crossover.
>
> **Me:** Yes and if an allele is lost it can be regenerated by mutation, but such regeneration introduces a sort of random element. It replaces the power of evolution with unguided random walk. This is clearest in the case of mutation. If a scheme needs mutation to work, there is something wrong. That's why I assume no mutation. Similarly, consider a good short schema preserved by linkage, a schema whose good quality is due to epistasis within the schema, not to the good quality of its alleles. Regenerating such a schema through recombination is just as random and unguided as regenerating a good allele through mutation. Furthermore, with your scaling there is in general no tendency for average value in the population to increase.
>
> **Yuval:** So what? I don't care about average value. I care about the value of the best individual.

Me: An adaptive system must both explore and exploit, properly dividing its effort between exploration and exploitation. Fisher's Theorem tells us that as the genetic algorithm evolves, exploring new promising genotypes, it continually increases the average value in the population at a rate proportional to some variance of values. Thus the system's performance tends to improve. But look at the mess your scaling causes. Your reproductive rates are the cubes of the values.[8] Now all Fisher's Theorem can tell us about your system is that the average of the cubes of the values is increasing at a rate proportional to the variance of the cubes of the values. With all your recombination, the average of the values themselves could actually be falling. The values measure how well individuals are doing. A system or population with falling values is like a business running into debt. The fact that its research is superb won't prevent bankruptcy if it can't exploit its knowledge while doing research. A population with falling values will die out from starvation or be eliminated by other more successful populations. An individual that intelligently learns but doesn't exploit what it knows will die.

Yuval: What universe are you living in? My system isn't fighting for survival in a jungle, capitalist or otherwise. It's learning to control a robot arm. As far as I am concerned, the evolving population is doing exploration. When the exploration period is finally over, I select the best individual and only then does exploitation begin. I'm no more interested in exploitation during the exploration period than Chris Watkins was when he proved his Q-learning convergence theorem.[9] And anyway, who are you to say that I'm using the cubes of values instead of values. What is the value of a program that guides this robot arm? Is it the reciprocal of the mean distance from the target path, as you seem to think it is, or is it the cube of the reciprocal of the mean distance? It comes down to this. You think I should be looking for the program with the highest value, and I'm looking for the program with the highest cubed value. Isn't that the same program?

Me: A genetic algorithm doesn't look for good genotypes. It looks for good schemata and produces a population of them. Getting good genotypes is a side effect. A genetic algorithm works when there is not too much epistasis, when the values of the building blocks (schemata) tell us quite a bit about the values of the genotypes. Then new genotypes generated by recombination will be promising ones. Then the populations of schemata that exploit our current knowledge are the very populations from which further exploration should depart. When you cube the values, you destroy the relationship between the schema values and genotype values. If epistasis was low before you cubed the values, it could become very high afterward. Then your genetic algorithm is little better than random search.

Yuval: Okay, maybe cubing increases epistasis. But maybe it decreases it. You have no way of knowing which. What matters is what works, and my system works. What you have been stating is not a bunch of proven principles, but merely generalized waffle illustrated by some formalism that works only when there is no recombination or when there is complete recombination and so is irrelevant to the situations we care about and is particularly irrelevant to my work. You fear large epistasis. We avoid large epistasis if the building blocks and their values are in some sense natural, meaning relevant to the particular problem. These matters cannot be properly discussed without bearing in mind the actual details of the problem the genetic algorithm is trying to solve, without actually running the experiments and finding out what works. If you ever ran any experiments you would know that your conclusions are too categorical and your criticisms misguided.

The above conversation is entirely imaginary and in it I have in places purposely misrepresented both Yuval's views and mine. But I thought such an imaginary conversation was the quickest way to sketch the context in which the results given here belong. As the conversation indicates, the issues are complicated and the situation is unclear. The reader presumably has other arguments that were not mentioned above.

19 Is This Proportionality Argument New?

I have not seen in the literature the argument that we gave above for proportionality. Kimura ([2] section 5.9) discusses what he calls a Maximum Principle for natural selection. In [21] I gave the standard warning against sampling error noise,[10] and this warning also runs through the population genetics literature. In a

[8] Actually, Yuval's reproductive rates were exponential functions of what I would call values. That's even more violent scaling.

[9] Okay, Yuval wouldn't have said that. His PhD predates Watkins's proof.

[10] but in a classifier systems context

sense, the proportionality argument above combines Kimura's Maximum Principle with that warning, but with a strange twist.

If we translate Kimura's statement of the Maximum Principle ([2] p. 234 bottom) into our notation for asexual reproduction, it reads roughly as follows.[11]

> Natural selection acts to change gene frequencies in such a way that $\sum_i \varphi'_i m_i$ is maximum for a given value of $\sum_i \varphi_i a_i^2 = \sum_i \frac{1}{\varphi_i}(\varphi'_i)^2$.

($\sum_i \varphi'_i m_i$ is of course the climb rate $\bar{m}'$.)

Now we see that Kimura's statement is confusing and in fact seems to be circular. The problem stems from the usual problem we encounter in the biological literature, the confusion of the two meanings of "fitness": value and reproductive rate. Maximizing $\sum_i \varphi'_i m_i$ for a given value of $\sum_i \varphi_i a_i^2$ makes no sense. The sum $\sum_i \varphi_i a_i^2$ is V, and it is whatever it is, whatever choice of φ'_i's we make, so we can make $\sum_i \varphi'_i m_i$ as large as we like.

But Kimura didn't mean $\sum_i \varphi_i a_i^2$, he meant $\sum_i \varphi_i r_i^2$, which is the right hand side of his little equation (since $\varphi'_i = \varphi_i r_i$). His proof uses that little equation[12], which he proved earlier using the very proportionality we need to justify. The circularity in Kimura's statement permeates his illustrative proof justifying it, and I found his argument terribly confusing.

The variance V is a measure of the variation on which selection can operate. Because Darwin didn't understand genetics as we know it, he feared this variance would drop too low for evolution to work. Since then, it has been clear that it is best for evolution if the variance V stays high. The fear is that evolution might itself cause V to drop, and to drop all the faster the bigger the climb rate. Kimura's $\sum_i \varphi'_i m_i$ is the climb rate $\bar{m}'$.

In his illustrative proof of the Maximum Principle, Kimura uses Lagrange multipliers to maximize climb rate $\bar{m}'$ subject to two conditions. The first is $\sum_i \varphi'_i = 0$. "The second is that [variance V][13] is (momentarily) constant." [p. 231] It sounds as if, like Darwin, he is worried about a drop in V and is trying to get the biggest $\bar{m}'$ he can without (momentarily) dropping V, though it's hard to be sure what his motive was.

To see things correctly, we must realize that while we would indeed like V to be high, V is not the relevant variance. The relevant variance is $\sum_i \varphi_i r_i^2$, and we want that variance to be *low*. If like most biologists you use the word "fitness" to mean reproductive rate, then

> We want the variance in fitness to be low.

This statement sounds so outrageously counter to all the biological language since Darwin that I was blinded to its truth.

One year in the 1970's, Holland was on sabbatical from the University of Michigan and Bernard P. Zeigler taught his Adaptive Systems course. In his lecture notes for that course, Zeigler removed the circularity in Kimura's argument, producing a coherent proof. I later modified Zeigler's proof in a small way by a square root scaling of the axes, thus turning it into a steepest ascent proof, the Lagrange multipliers proof given in section 5. I discovered later that taking square roots of probabilities is fairly standard now, and that Fisher is connected with this technique, though evidently in a context different from what we have here.

But of course in removing the circularity in the argument, Zeigler had to remove the motivation. The steepest ascent proof in section 5 is on the face of it unmotivated. Why should one want to follow the steepest ascent path in the square root space?

The key to the argument here is connecting the admittedly damaging sampling error noise to the variance $\sum_i \varphi_i r_i^2$, as we did above in section 15. Once we note that the learning rate (the length of $\mathbf{r}$) affects the climb rate and the square root of that variance in the same way, we can quantitatively connect the two and see that we can get the variance low if we follow the steepest ascent path in the square root space. That we should try to get a variance low, not high, is the surprising twist.

If we look again at Kimura's statement of the maximum principle and remove the circularity by removing the left hand side of the little equation, Kimura's statement becomes this: "Natural selection acts to change gene frequencies in such a way that the climb rate $\bar{m}'$ is maximum for a given value of the variance W^2 ." This is true, and follows from lemma 6. It also acts to change gene frequencies in such a way that variance W^2 is *minimum* for a given value of the climb rate $\bar{m}'$, but Kimura doesn't say this.

The details of the Lagrange multipliers steepest ascent proof in section 5 are almost identical to the details in Kimura's argument. This is not surprising. The steepest ascent proof is the guts of the proof

[11] Kimura is actually there discussing the one locus case with recombination.

[12] See his equation 5.9.7

[13] He actually says "genic variance", and calls it V_g , because he is discussing the one locus case with recombination.

of lemma 6, and Kimura's corrected statement is a weaker version of lemma 6. But *why* do we want to maximize the climb rate for a given value of the variance W^2 ? Because we want the variance W^2 to be *low*. Kimura never says this. Presumably because he doesn't believe it. He confuses W^2 with V, and he knows V should be high.

It is possible that the Maximum Principle argument has been presented elsewhere as an argument similar to the one I've given here. And it is possible that the argument here is a special case of arguments given elsewhere, perhaps by Fisher. Investigation of these possibilities might provide useful insight.

20 Suppose there is no Epistasis

We said in section 13 that if there were too much epistasis there was a problem with Fisher's approach. But if epistasis is low, and Fisher's approach works, does Holland's story work too? Indeed, if there is no epistasis do schemata have static values?

The vector $\mathbf{m}$ is a vector of partial derivatives of μ. Suppose over a subset $\mathcal{B}$ of the population domain the vector $\mathbf{m}$ is constant. ($\mathcal{B}$ could be the whole population domain, or it could be just the $\boldsymbol{\varphi}$'s for which $\bar{\boldsymbol{\varphi}} = \boldsymbol{\varphi}$, or any subset.) Then by equation (3) we see that $\bar{m}$ is linear over $\mathcal{B}$. (The m_i's in equation (3) are constant.)

Now suppose in addition that $V_e = 0$ everywhere in $\mathcal{B}$. Then we always have $a_i = \hat{a}_i$.

> **Trivializing Assumptions:** Suppose that over $\mathcal{B}$: the m_i's are constants independent of $\boldsymbol{\varphi}$, and $V_e = 0$.

Then if our distribution $\boldsymbol{\varphi}$ is in $\mathcal{B}$, we have $\hat{\mathbf{a}} = \mathbf{a}$. (Since $V_e = 0$)

Lemma 8 *Suppose the trivializing assumptions hold.*
Then if k and γ are alleles in the same locus set, $\alpha_k - \alpha_\gamma$ is the same for every $\boldsymbol{\varphi}$ in $\mathcal{B}$.

Proof: Let i be a genotype for which $c_{ik} > 0$. (There must be one.) Then let ℓ be a genotype identical to i except that one of the k alleles is replaced by γ. So $c_{ik} = c_{\ell k} + 1$, and $c_{i\gamma} = c_{\ell\gamma} - 1$, and for any allele j different from k or γ we have $c_{ij} = c_{\ell j}$. Then by equation (26) we have $\hat{a}_i - \hat{a}_\ell = \alpha_k - \alpha_\gamma$. Since $V_e = 0$, we have $a_i = \hat{a}_i$ and $a_\ell = \hat{a}_\ell$. By equation (5), $m_i - m_\ell = a_i - a_\ell$. Thus $m_i - m_\ell = \alpha_k - \alpha_\gamma$. And both m_i and m_ℓ are independent of $\boldsymbol{\varphi}$, so long as $\boldsymbol{\varphi} \in \mathcal{B}$. □

We now look at two distributions in $\mathcal{B}$, which we can call $\boldsymbol{\varphi}$ and $\ddot{\boldsymbol{\varphi}}$. We will use the ¨ to distinguish quantities given by $\ddot{\boldsymbol{\varphi}}$ from the corresponding quantities ($\mathbf{p}$, $\boldsymbol{\alpha}$, etc.) given by $\boldsymbol{\varphi}$. For example, the conclusion of lemma 8 tells us that $\alpha_k - \alpha_\gamma = \ddot{\alpha}_k - \ddot{\alpha}_\gamma$.

Lemma 9 *Suppose the trivializing assumptions hold.*
Then $\bar{m} = \ddot{\bar{m}} + N \sum_k p_k \ddot{\alpha}_k$; and if j is in locus set $\mathcal{L}$, we have $\alpha_j = \ddot{\alpha}_j - \sum_{k \in \mathcal{L}} p_k \ddot{\alpha}_k$.

Proof: By lemma 8, if k and j are alleles in the same locus set $\mathcal{L}$, then $\alpha_k - \alpha_j = \ddot{\alpha}_k - \ddot{\alpha}_j$. So $\ddot{\alpha}_k - \alpha_k = \ddot{\alpha}_j - \alpha_j$. Multiplying by p_k gives $p_k\ddot{\alpha}_k - p_k\alpha_k = p_k\ddot{\alpha}_j - p_k\alpha_j$. Summing over all k in $\mathcal{L}$ and using equations (20) and (24) gives $\sum_{k\in\mathcal{L}} p_k\ddot{\alpha}_k = \ddot{\alpha}_j - \alpha_j$, provided $j \in \mathcal{L}$. (This yields the second half of the lemma.) We now multiply by c_{ij} and sum over all j in $\mathcal{L}$. By equation (19), this gives $N\sum_{k\in\mathcal{L}} p_k\ddot{\alpha}_k = (\sum_{j\in\mathcal{L}} c_{ij}\ddot{\alpha}_j) - (\sum_{j\in\mathcal{L}} c_{ij}\alpha_j)$. We now sum over all $\mathcal{L}$, to obtain $N\sum_k p_k\ddot{\alpha}_k = (\sum_j c_{ij}\ddot{\alpha}_j) - (\sum_j c_{ij}\alpha_j)$. By equation (26), the right hand side is $\ddot{\hat{a}}_i - \hat{a}_i$. But by the trivializing assumption $V_e = 0$, we have $\hat{a}_i = a_i$ and $\ddot{\hat{a}}_i = \ddot{a}_i$, so the right hand side is $\ddot{a}_i - a_i$, or $(\ddot{m}_i - \ddot{\bar{m}}) - (m_i - \bar{m})$. But by the trivializing assumptions, $\ddot{m}_i = m_i$, so the right hand side is $\bar{m} - \ddot{\bar{m}}$, and the whole equation is $N\sum_k p_k\ddot{\alpha}_k = \bar{m} - \ddot{\bar{m}}$. □

We see from lemma 9 that if the trivializing assumptions hold, then $\bar{m}$ is a linear function of the p_j's. This is true even if the population is not at linkage equilibrium, since a_i is always equal to $\hat{a}_i$. We can let $\ddot{\boldsymbol{\varphi}}$ be the uniform distribution and then the $\ddot{\alpha}_j$'s are what some genetic algorithms people call the static values of the alleles. We see that in this case, static values do indeed have some significance, but they have this significance only when the trivializing assumptions hold. It may be useful to re-write the first equation in lemma 9 as

$$\bar{m} = N \sum_k p_k \left(\frac{\ddot{\bar{m}}}{NL} + \ddot{\alpha}_k \right) \quad , \tag{57}$$

so the constant $\frac{\ddot{\bar{m}}}{NL} + \ddot{\alpha}_k$ could be thought of as the static value of allele k.

Let's abbreviate that static value as $\ddot{\beta}_k$. Suppose allele k is in locus set $\mathcal{L}$. Then at $\boldsymbol{\varphi}$, the average of all the static values of the alleles in locus set $\mathcal{L}$ is $\sum_{j\in\mathcal{L}} p_j\ddot{\beta}_j = \frac{\ddot{\bar{m}}}{NL} + \sum_{j\in\mathcal{L}} p_j\ddot{\alpha}_j$. We can subtract this from $\ddot{\beta}_k$ to obtain the excess static value of k at $\boldsymbol{\varphi}$. It is $\ddot{\alpha}_k - \sum_{j\in\mathcal{L}} p_j\ddot{\alpha}_j$.

Now suppose we have a natural reward scheme. We can derive a formula for $p_j{}'$ in terms of the static allele values. From equation (48) and the second half of lemma 9, we have $p_j{}' = Kp_j(\ddot{\alpha}_j - \sum_{k\in\mathcal{L}} p_k\ddot{\alpha}_k)$. So p'_j is Kp_j times the excess static value, a static version of equation (48).

But of course the trivializing assumptions do not usually hold. Allele values are population dependent quantities. Fisher's analysis, as discussed in this paper, points the way toward the proper approach. The trivializing assumptions are unnecessary, and they throw away very important features of the real world.

21 An Attempted Defense of Holland's Story

But wait a minute. If we are already assuming restrictions on interaction between genotypes, then if there are enough restrictions, each m_i will be nearly independent of the genotype frequencies. And if we are assuming restrictions on distant genes, then if we look at short schemata we see that epistasis between the various short schemata is low. So we see that if we think of short schemata where we previously thought of alleles, then the trivializing assumptions nearly hold. Then by lemma 9 and equation (57), each such schema has a value that is nearly independent of the population and we see that $\bar{m}$ is a nearly linear function of the schema probabilities, where the coefficients are the (nearly constant) schema values. Of course the fact that schemata overlap complicates the story, but not fundamentally, as far as the point here is concerned. We can think of the schema values as static values, or true values, independent of the population. We admitted above that a population was only a sample, and now Holland's story that successive populations are successive samples used to estimate these schema values makes sense.

So schemata have static values (or nearly static). The purpose of the evolutionary mechanism is to estimate these values. The value of an individual to a population (it's marginal utility) is given by adding up the values of its schemata (with provisions for overlapping schemata and with some small epistatic unbiased noise). Any value in what an individual does is dictated (in the long run) by the genes of which the individual is made. These genes say how well the individual responds to the problems set by the environment. Holland's story leads us to this Fisher-like view, a view popularized in its extreme form by Dawkins. [3] The evolutionary process will work better if it quickly replaces bad schemata (low value schemata) in the population by good schemata. Evolution has become an efficient process and so this replacement is very fast. Evolutionary pressure on schemata is strong, and evolution is fast, as in the case of the Birmingham moths. Thus concerns of eugenicists like Fisher are plausible concerns. Social safety nets can hinder the elimination of bad schemata, or even cause them to spread through the population. Equality of opportunity is both impossible (since different individuals hold different genes) and dangerous (since such equality could cause bad schemata to spread). An allele for low salary causes economic difficulties in the population, and economic distress would be alleviated if we could reduce the prevalence of this allele. (I'm thinking here of the brown-eye allele, which has a much lower salary than the blue-eye allele. Remember that an allele's value (salary) is the average of the values (salaries) of the individuals containing it. Diploidy doesn't affect this argument.)

22 The Defense Fails

But that's ridiculous. It's true that both Fisher's approach and Holland's story use linear approximations and assume some restriction on the interaction between genotypes, but the way they do it is very different. Both start with the same value function μ, and both use linear approximations to the value, but they do it very differently.

Holland's story uses a linear approximation based on partial derivatives at a particular special point $\ddot{\boldsymbol{\varphi}}$. The resulting linear approximation to the value function μ, to the value $\bar{m}$, is exactly correct at point $\ddot{\boldsymbol{\varphi}}$ and good at points near $\ddot{\boldsymbol{\varphi}}$, but increasingly bad as we get farther and farther from $\ddot{\boldsymbol{\varphi}}$. Holland's story would have us climb this linear approximation. As we go up the linear slope, getting farther and farther from $\ddot{\boldsymbol{\varphi}}$, the function we are climbing is increasingly losing touch with reality. Eventually we find our current population $\boldsymbol{\varphi}$ at a point where the linear function and the true function μ have little to do with each other. The population $\boldsymbol{\varphi}$ is still climbing the linear function, but on the true function μ it could be heading straight downhill. This is not a noise problem. It is the deterministic track itself that can head downhill. The linear approximation in Holland's story is a ridiculous approximation.

None of this can happen in Fisher's approach. Fisher's partial derivatives are partial derivatives at the point we are at, at the point $\boldsymbol{\varphi}$. Fisher's linear approximation is not ridiculous, it's simply differential calculus, and Fisher's approach operates at points close to $\boldsymbol{\varphi}$, where the approximation is good. In Fisher's approach, the deterministic track of $\boldsymbol{\varphi}$ cannot head downhill if the m_i's are what they claim to be.

Let me say that more carefully. We said that in Fisher's approach there are problems estimating the correct m_i's, both because of biases introduced by cannibals and other genotype interactions and because of sampling error noise. Both of these problems are also present in Holland's story. In both stories, these problems can cause $\boldsymbol{\varphi}$ to take a step downhill.

But in Holland's story there is the much bigger problem that as $\boldsymbol{\varphi}$ gets farther and farther from the special point $\ddot{\boldsymbol{\varphi}}$, its motion is based on the increasingly irrelevant partial derivatives at $\ddot{\boldsymbol{\varphi}}$.

If the trivializing assumptions don't hold, equation (57) and the equations in Lemma 9 are seriously in error unless we are close to $\ddot{\boldsymbol{\varphi}}$. By contrast, Fisher's equations are all correct. Problems with Fisher's approach are matters like "Although equations (45) and (55) hold, it could be that V_e is close to V." Both Fisher's approach and Holland's story have such problems. Fisher's approach highlights those problems so we can address them. Holland's story buries those problems by introducing a much bigger unnecessary problem.

Fisher's approach tries to get $\boldsymbol{\varphi}$ to climb up the local hill in the function μ. Fisher's approach is imperfect and there are the problems we have mentioned. But Holland's story has $\boldsymbol{\varphi}$ heading off happily into the distance when the hill it was supposed to be climbing has long been left behind.

Holland's story makes no sense.

23 A More Sensible Approach

Luckily, Holland's Genetic Algorithms have not heard Holland's story. They do not head off into the distance. They work fine. Their schemata do not have true values. Their schemata do not have static values. Their schemata have population dependant values.

Holland's story depends on the belief that a gene or schema describes some characteristic of an organism's response to problems set by the environment, indeed that an organism's value is given by how well it responds to such problems. This is essentially a perception oriented view of adaptation, evolution, learning, and intelligence. Much of what Hebb said [10] is dominated by such perception orientation, though he was dealing with learning in Neural systems rather than evolution in populations. Fu's work (based on Hebb's ideas) explicitly restricts itself to perception. Holland hopes her experiments will eventually include a complete loop from the system through the environment and back to the system.[14] The loop would have a forward leg in which the system action causes a change in the environment (this change is the environment's response), and a return leg in which the environment's action causes a change in the system (this change is the system's response). Perception is merely the return leg. This return leg is the exclusive focus of Fu's PhD work. [7] By contrast, the forward leg, the action of the system on the environment, has been the focus of Paul Scott's work [18], particularly when he was studying the cat motor cortex. Indeed, Holland's work traditionally focused on the action of the system on the environment, rather than on the return (perception) leg of the loop from the environment back to the system [11]. My own work begins with Holland's traditional bias and modifies the system to include perception as a secondary mechanism [23]. It is the interaction of system and environment that is the starting point of a sensible analysis. The system's response to the environmental problems is only half of the interaction. The effect of system action on the environment is the other half. Organisms do not merely try to reproduce within a context of environmental problems set by the "laws of nature" in their environment.

> "It is a long way from the 'laws of nature' to the horse's hoof. Rabbits, kangaroos, snakes, and grasshoppers, all of whom traverse the same ground as the horse, do not have hooves. Hooves come not from the nature of the ground but from an animal of certain size, with four legs, running, not hopping, over the ground at a certain speed and for certain periods of time. The small gracile ancestors of the horse had toes and toenails, not hooves, and they got along very well indeed. So, too, our central nervous systems are not fitted to some absolute laws of nature, but to laws of nature operating within a framework created by our own sensuous activity. Our nervous system does not allow us to see the ultraviolet reflections from flowers, but a bee's central nervous system does. And bats 'see' what nighthawks do not. We do not further our understanding of evolution by general appeals to 'laws of nature' to which all life must bend. Rather we must ask how, within

[14] personal communication, to myself and also to Fu

the general constraints of the laws of nature, organisms have constructed environments that are the conditions for their further evolution and reconstruction of nature into new environments.

"It is difficult to think of any physical force or universal physical law that represents a fixed problem to which all organisms must find a direct solution. We think of gravitation as universal, but because it is such a weak force, it does not apply in practice to very small organisms suspended in liquid media. Bacteria are largely outside the influence of gravity as a consequence of their size, that is, as a consequence of their genes. On the other hand, they are subject to another universal physical force, Brownian motion of molecules, which we are protected from by our large size, again a result of our evolution. The most remarkable property of living organisms is that they have avoided biologically the chemical laws of mass action and the high energy needed to initiate most chemical reactions; both have been accomplished by structure. The structure of the genes themselves, and the way they are held together in very large macromolecular structures, makes it possible for gene replication and protein synthesis to take place even though there is only a single molecule of each gene in each cell. The structure of enzymes, in turn, makes it possible to carry out at ambient temperatures chemical reactions that would otherwise require great heat.

"It is impossible to avoid the conclusion that organisms construct every aspect of their environment themselves. They are not the passive objects of external forces, but the creators and modulators of these forces. The metaphor of adaptation must therefore be replaced by one of construction, a metaphor that has implications for the form of evolutionary theory."
(Levins and Lewontin [16, pages 103 and 104])

The approach in [23] tries to follow the spirit of the previous quote, but it falls short. The approach in [23] thinks of the environment as a large finite automaton with a payoff attached to each state. Following Holland's traditional bias, it is the action of the system on the environment that is fundamental, and the action of the environment on the system (perception) secondary. But in [23], the action of the system merely changes the environment state, it doesn't change the state diagram. To change the diagram, the approach in [23] would need to be extended. The system, which in [23] is a classifier system, could be expanded and part of it only could be what we would call the system. The rest could be part of the environment. In this coevolutionary context, we would see the environment state diagram changing, and the system action would affect these changes. Whether this extension captures the essence of the real world situation is not clear.

Even unextended, the approach in [23] captures some of what we need. It is easy to see that there the system-environment complex forms a large Markov chain, and that in this chain, the average payoff per time unit is a rational function of parameters in the system. It is possible to view the system there as a population, with genotype probabilities given very simply by the parameters. Thus the average payoff per time unit becomes a rational function of the genotype probabilities. This is just what we need for Fisher's analysis. The function is the function μ of section 4. Thus we can use the general Fisher results to analyze the system-environment interaction through analyzing μ and interpreting the results. That we can so use Fisher is not a peculiarity of the scenario in [23]. Many scenarios, both sequential and parallel, in genetic algorithms and biology produce the required differentiable μ. (For example, consider Axelrod's work. [1]) Thus Fisher's results give us an insight into system-environment interaction, the proper subject of our study.

If we make the trivializing assumptions, we conceptually break the interaction loop and are left with only the effect of the environment on the system, an effect which alone determines the schema values. Our results then become simplified to a situation that doesn't obtain in the situations we care about. Holland's story makes this very mistake. It concentrates on the perception side of the interaction.

Holland traditionally concentrated on the action part of the interaction. [11] His traditional approach was amenable to Fisher's general analysis. He insisted that all his students read Fisher. Holland's story is a turning away from Holland's traditional approach. Holland's story concentrates on the perceptual side of the interaction at the cost of the action side. It prevents us from understanding the nature of system-environment interaction.

24 Streamlining

Holland's traditional approach uses Fisher's theorem to make the general argument that if average reproductive rate is proportional to value, then φ will tend to move to increase $\mu(\varphi)$. This argument is general and does not depend on the trivializing assumptions required by Holland's current[15] story. Neither does it

[15]By "current" I mean 1998, the year of the argument this paper is part of. (See the note at the start of this paper.)

expect quick replacement of some schemata by others, as does his current story. The argument works when value differences are small as well as when they are large.

Much of adaptation (be it evolution or learning) works by simplification and streamlining, as redundant features are reduced or eliminated. More advanced organisms have got where they are by such simplification, be it simplification in flower construction, aortic arches, arthropod appendage numbers, or organization of prokaryote genomes.[16] Modern eubacteria are particularly advanced in this respect, having eliminated many redundancies. Such streamlining is the basis of the evolution of complex organization.

This streamlining and reduction of redundancy relies on taking advantage of very small value differences. Holland's traditional approach uses Fisher's results to argue that such slow adaptation will occur if reproductive rate is proportional to value. Streamlining and elimination of redundancy will slowly take place. The resources freed by this become available for new organization, the building blocks of which include the now streamlined pieces of the organization already achieved. For Holland, one of the main attractions of Hebb's Cell Assembly Theory was that, in that theory, as cell assemblies were used, they became more compact and streamlined, freeing redundant neuronal units for use in other cell assemblies.[17] I do not understand why Holland has rejected his traditional approach and produced a story that stresses quick schema replacement in special cases of high value difference and rejects the rich promise of general organizational adaptation, with its interplay of streamlining and complexity and its recognition that they are two sides of the same coin.

25 The New Deal

Schemata do not have static values; they do not have true values. Any values they have are effects of system-environment interaction and so depend on the current population. Individuals' values are not the result of their schema values, they are the result of interaction with the environment and the current population. The low salary (brown-eye) allele is a low salary allele in Detroit, but not in other populations where there is no such history of racial discrimination. If a Detroit African-American kid does badly in school, it probably does have to do with his genes, for if he had had different genes, he and his family would have had a different skin color and hence his educational background would have been different. But notice that his value and the value of his genes is given by the interaction of individual and environment, environment that the population can change. If there is a problem, it is that very interaction that needs to be addressed, not the genetics.

If a student of mine has trouble understanding me, the natural conclusion is that someone somewhere has failed to live up to what that student needs. [14] It is not easy to address such failings, but we must constantly try to do so. We do not need some scientist telling us that maybe we have overlooked the possibility that somewhere in the student's DNA is a string that when decoded spells FAILURE.

Our School systems and indeed our whole society has so many failings that all of us are very far from realizing our potential. By addressing these failings we can realize a vast untapped source of inspiration and productivity. By empowering those people whom the system has failed and oppressed, we not only bring about a more equal and more just society, we bring about a far more prosperous one. It's only Common Sense.

> "Now I come to the links which will build us a more lasting prosperity. I have said that we cannot attain that in a nation half boom and half broke. If all of our people have work and fair wages and fair profits, they can buy the products of their neighbors and business is good. But if you take away the wages and the profits of half of them, business is only half as good. It doesn't help much if the fortunate half is very prosperous – the best way is for everybody to be reasonably prosperous." (Franklin Delano Roosevelt, Fireside Chat, Monday, 24 July 1933)

References

[1] Robert Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In Lawrence Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41, Los Altos, California, 1987. Morgan Kaufmann.

[2] J. F. Crow and M. Kimura. *An Introduction to Population Genetics Theory.* Harper and Row, New York, 1970.

[16] See [22] for further discussion.

[17] My evidence that such compactification was a main attraction for Holland is second hand. He evidently stated this in a conversation with Fu in early 2001.

[3] Richard Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, England, 1989.

[4] R. A. Fisher. *The Genetical Theory of Natural Selection.* Clarendon Press, Oxford, 1930.

[5] R. A. Fisher. The theoretical consequences of polyploid inheritance for the midstyle form of *Lythrium salicaria*. *Ann. Eugen.*, 11:31–38, 1941.

[6] R. A. Fisher. *The Genetical Theory of Natural Selection.* Dover, New York, 1958. second revised edition of the work first published in 1929.

[7] Leeann Liang Fu. An analysis of Hebb's Cell Assembly as a mechanism for perceptual generalzation, PhD dissertation, 2003. University of Michigan, Ann Arbor.

[8] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, Reading, Mass., 1989.

[9] John J. Grefenstette. Deception Considered Harmful. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms*, pages 75–91, San Mateo, CA, 1993. Morgan Kaufmann.

[10] D. O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York, 1949.

[11] J. H. Holland. Adaptive plans optimal for payoff-only environments. In *Proceedings of the Second Hawaii Conference on System Sciences*, pages 917–920. University of Hawaii, Honolulu, 1969.

[12] J. H. Holland. *Adaptation in Natural and Artificial Systems*. Univ. of Michigan Press, Ann Arbor, 1975.

[13] John H. Holland. *Hidden Order*. Addison-Wesley, Reading, MA, 1995.

[14] John Holt. *How Children Fail*. Penguin, London, England, 1982. first published, 1964.

[15] M. Kimura. On the change of population fitness by natural selection. *Heredity*, 12:145–167, 1958.

[16] Richard Levins and Richard Lewontin. *The Dialectical Biologist*. Harvard University Press, Cambridge, MA, 1985.

[17] Hilary Rose. Colonising the Social Sciences? In Hilary Rose and Steven Rose, editors, *Alas, Poor Darwin – Arguments Against Evolutionary Psychology*. Jonathan Cape, London, England, 2000.

[18] P. D. Scott. *The Cerebral Neocortex as a Programmable Machine, Ph.D. Dissertation.* University of Sussex, Falmer, Sussex, 1975.

[19] John Maynard Smith. *Evolutionary Genetics*. Oxford University Press, Oxford, England, 1989.

[20] T. H. Westerdale. A reward scheme for production systems with overlapping conflict sets. *IEEE Trans. Syst., Man, Cybern.*, SMC-16(3):369–383, 1986.

[21] T. H. Westerdale. A defense of the bucket brigade. In J. David Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 282–290, San Mateo, CA, 1989. Morgan Kaufmann.

[22] T. H. Westerdale. Redundant classifiers and prokaryote genomes. In Richard K. Belew and Lashon B Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 354–360, San Mateo, CA, 1991. Morgan Kaufmann.

[23] T. H. Westerdale. An approach to credit assignment in classifier systems. *Complexity*, 4(2):49–52, 1999. electronic version, 40 pages, http://journals.wiley.com/complexity.

[24] D. S. Wilson. The group selection controversy: history and current status. *Ann. Rev. Ecol. Syst.*, 14:159–187, 1983.