The Dimensions of Selection

Peter Godfrey-Smith
Department of Philosophy
Stanford University†*

and

Richard Lewontin
Museum of Comparative Zoology
Harvard University

Appears in *Philosophy of Science* 60 (1993): 373-395.

Abstract

Proponents of genic selectionism have claimed that evolutionary processes normally viewed as selection on individuals can be "represented" as selection on alleles. This paper discusses the relationship between mathematical questions about the formal requirements upon state spaces necessary for the representation of different types of evolutionary processes, and causal questions about the units of selection in such processes.

1. Introduction

Among the many questions that become tangled together in units of selection debates, are questions about the possibility of "representing" various evolutionary processes within different theoretical frameworks. In particular, it has been claimed by Williams (1966), Dawkins (1982), Maynard Smith (1987), Sterelny and Kitcher (1988) and Waters (1991) that evolutionary processes usually described and modelled at the level of the individual organism as selection on genotypes, can always be redescribed in terms of competition between alleles. Critics such as Wimsatt (1980), Sober and Lewontin (1982), Sober (1984), and Lloyd (1988) reply that the only sense in which selection on individuals can be "represented" at the genic level is trivial and uninformative; that such a representation discards essential information, conceals the genotypic facts doing the real work, or both.[1]

---

† We would like to thank Marcus Feldman, Deborah Gordon, Philip Kitcher, Elisabeth Lloyd, Elliott Sober, Ken Waters and two anonymous referees for comments on earlier drafts. Thanks to Devin Muldoon for the graphics, and special thanks from PGS to Sarah Otto for numerous corrections and suggestions.

* Send reprint requests to: Peter Godfrey-Smith, Department of Philosophy, Stanford University, Stanford, CA, 94305-2155.

One aim of this paper is to clarify some of these questions by giving a precise sense to the idea that an evolutionary process can be "represented" in a given framework. Representing a process of evolution might involve merely plotting or tracking changes the system goes through, or it might involve predicting the development of the system, before the fact. In our discussion the second, predictive sense of representation will be adopted (as in Wimsatt 1980). A representation of an evolutionary process must be a dynamically sufficient model of that process.

We will then argue that, even for 1-locus cases, it is not possible to represent all evolutionary processes normally understood as selection upon genotypes in the standard language of genic selection -- with the (unconditional) frequencies of alleles. However, this is not because such models compute allele frequencies; it is because they do not have the required number of dimensions to represent the process. If an allelic model is enriched to the right dimensionality, then it can represent all processes of selection on genotypes. This enrichment can be done with conditional allele frequencies. Our conclusions about rival modes of "representing" evolutionary processes are conventionalist. However, this does not solve the units of selection problem in general, for there is more to an understanding of evolution than having a dynamically suficient model. At least some of the remaining questions are questions about the reality of various causal forces, and properties of models are not always good guides to properties of the processes the models describe (Sober 1981, 1984). In particular, counting the dimensions of a model is distinct from counting the causal forces underneath.

2. Representation and Dynamic Sufficiency

In this paper "representing" evolution will mean representing it in a mathematical model. Our aim at first is to make as much progress as possible on the units of selection problem with a stripped-down, minimal framework, so the only relevant components of a model are: (i) a space of N dimensions, where each dimension is used to represent the value of some variable in the model. It is not necessary that these variables correspond in any intuitive way to the natural properties of real entities; (ii) a set of rules transforming some (combinations of) variables into others; that is, a set of rules describing how the system moves from one point in the space to another; and (iii) a set of parameter values, contained in the rules of the model, that are given "from without." They do not change as the state of the system changes. As with the variables, these parameters need not be natural-looking properties of anything. They are just numbers.

For a (deterministic) model to be dynamically sufficient, given the system's location in state space at some time, it must be possible to compute the system's location at all later times, using only the rules and the system's present location.[2] To represent something, for us, is to represent it in a dynamically sufficient model. So if it is claimed that some process of evolution can be represented with a model whose dimensions are all frequencies of alleles, to satisfy us it is not sufficient that the progress of the system

---

1 It is unclear who exactly should be in these lists. Lloyd (1988) has argued that Williams' 1966 position is different to that of Dawkins and the others, but see Sober 1984 chapter 7. It may be that Maynard Smith would not accept as strong a version of this claim as the others.

2 In the case of a stochastic model, the present state of the system along with its rules determine a probability distribution over possible subsequent states.

simply be tracked by noting a succession of points in allele frequency space that the system occupies. It must be possible to predict the path of the system, using only the present location and a set of rules containing allele-frequency variables and fixed parameters.

In this discussion, we do not assume that theories are to be analyzed as models, or families of models, as the semantic view of theories holds (Lloyd 1988). We do assume that achieving a theoretical understanding of evolution involves at least the construction of dynamically sufficient models (Lewontin 1974 Chapter 1).

For any dynamical process, there is a required dimensionality N for models of that process. There will be a certain minimum number of independent dimensions (variables) without which it is not possible to build a dynamically sufficient model of that process. Moreover, the identity of those variables is restricted. Only some sets of N variables will form a dynamically sufficient space.

If a process has a required dimensionality of N, there will be rules mapping all points in the N-dimensional space to points in various (N-c)-dimensional spaces, in a many-to-one fashion. Models of these lower dimensionalities will not be dynamically sufficient. They can be used to track properties of the higher-dimensional process, but they will not tell us how it moves from point to point. There will not be rules mapping all points in the (N-c)-dimensional space one-to-one to points in N-dimensional space. If there were such rules, it would be possible to move from a point in (N-c)-dimensional space to a unique point in N-dimensional space, apply the rules of that richer space, and then drop down to (N-c) dimensions again. Then (N-c) dimensions would be sufficient after all.

There will also be mappings from N-dimensional space to various (N+c)-dimensional spaces, and there must be models in (N+c) dimensions that are dynamically sufficient. Not all spaces of (N+c) dimensions will accommodate dynamically sufficient models; some ascensions to higher dimensionality throw away essential information. But if a process can be represented in N dimensions, there is certainly some way of representing it in (N+c).

Before it can be determined whether a given dimensionality is sufficient, it is necessary to have a clear conception of what is to be explained. In the present context of evolutionary genetics, it must be decided whether changes in population composition from generation to generation are all that is to be predicted, or whether a more fine-grained description, including changes within generations, is demanded. This decision is necessary because of a peculiarity of sexual reproduction; reproduction by a genotype is not the same as reproduction of a genotype. If we consider a diploid, sexually reproducing population segregating for two alternative alleles, A and a, at some locus, there will be three genotypes AA, Aa and aa. The frequencies of A and a are represented by p and q respectively. The relative frequency of any one of these genotypes, say AA, in two successive generations does not depend only on the survival and fertility of AA itself, because AA genotypes do not reproduce by themselves. Depending upon which genotype an AA individual mates with, it will produce some mixture of offspring genotypes which may or may not include its genotype. For example, if AA mates with aa no AA offspring are produced at all, but only heterozygotes Aa. Thus there is a distinction between the differential rate of production by a genotype (its relative survivorship and fertility) and the differential reproduction of a genotype, which depends also on mating frequencies and the operation of Mendelian segregation.

As a consequence one needs to look at one generation of evolutionary change as made up of two stages: (i) a survivorship stage that biases the frequencies of genotypes to increase the representation of the more fit, and (ii) a reproduction phase that depends on differential fertility and the properties of meiosis and sex. This phase may return to the population genotypes removed by differential survival. For example consider an allele a that is in the process of being lost from a population because in homozygous (aa) form it is completely lethal and its carriers die before reproduction. In any generation there will be no aa homozygotes among adults, but in the next generation some aa will reappear, in immature individuals produced by the mating of Aa x Aa. Of course, as the frequency of the allele a is reduced in successive generations, the frequency of aa among juveniles will also decrease. But whether or not the frequency of aa is seen to decrease monotonically with time will depend on whether the population is censused once a generation, perhaps at birth, or whether the intra-generational trajectory is also followed. The actual trajectory of change in aa frequency, if we include changes within generations, is a saw-toothed oscillation, but the frequency of the allele a does not go through this oscillation because the sexual reassortment of genes that restores aa genotypes has no effect on the allele frequency. Figure 1 illustrates the process of change for aa genotypes and a genes.

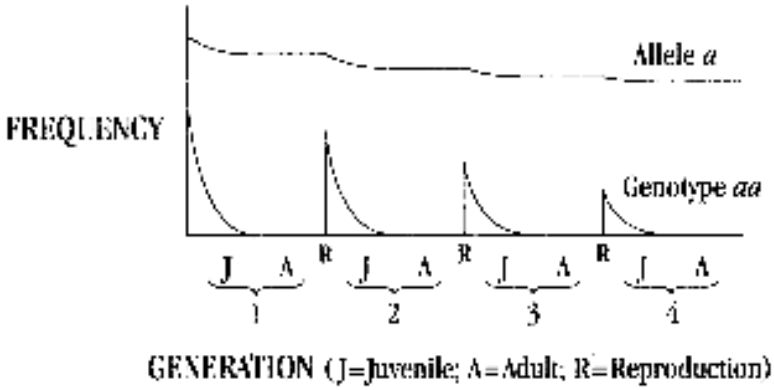

Figure 1: The loss and restoration of a deleterious recessive.

These details become important when we consider cases of stable equilibria of intermediate gene frequencies because of balanced polymorphism. For example, if the fitnesses (reproductive rates by genotypes) are such that heterozygotes Aa leave more offspring than either AA or aa, a stable intermediate pair of allele frequencies will be established such that there is no change in allele frequencies or genotype frequencies between generations. However, within generations the sawtooth oscillation of genotype frequencies continues. Heterozygotes are increased in frequency and homozygotes are decreased in frequency by the selection process, but sex restores the frequencies to their previous state each generation. The equilibrium is, in a sense, an equilibrium between selection and sex, as illustrated in Figure 2.

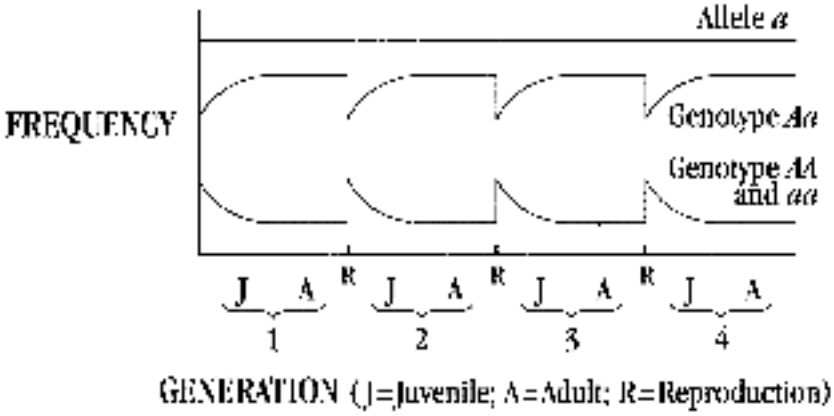

Figure 2: Balanced polymorphism

Figure 2 shows a qualitative difference in the behavior of genes and genotypes. Genes remain constant in frequency both within and between generations. No observations of simple allele frequencies can reveal whether natural selection (differential reproduction by genotypes) is operating or not. Observations of genotype frequencies between generations are similarly uninformative. A more fine-grained perspective is needed; observations at the level of genotypes within generations will reveal the operation of this selective process.

If we have certain fine-grained explanatory interests the inadequacy of the simple view which tracks evolution with p and q is immediately apparent. The example above appears in compressed form in Sober and Lewontin 1982, where it is used to argue that the genotype is the appropriate level of description. But that argument was grounded in a realist assumption that actual causal forces were to be distinguished. Through much of the present paper we will avoid assumptions of this kind, and see what follows from a leaner set of axioms. Though we are chiefly interested in the dimensionality required for the prediction of evolutionary trajectories, it is still necessary to make a decision about the temporal grain of these trajectories. If the entire trajectory within and between generations is demanded, then it is clear from the argument of Sober and Lewontin that something more than simple allelic information is needed in some cases. If only the trajectory between generations is of interest, the question of the required dimensionality is different. In the remainder of this paper, we will generally assume the latter. If a model predicts gene and/or genotype frequencies from generation to generation, it is a dynamically sufficient representation of that evolutionary process. Later we will find that even intra-generational dynamics are comprehensible for some allelic models, but we will not demand this level of resolution in general.

3. Representing Selection

We turn now to a series of cases, displaying the required dimensionalities associated with different evolutionary processes, and showing what freedom of movement there is between theoretical frameworks. We assume two alleles A and a at a locus, with genotypes AA, Aa and aa. These genotype frequencies will be labelled D, H and R respectively (for dominants, heterozygotes and recessives). These frequencies are

evaluated each generation, just after reproduction. Though there are three genotypes, their frequencies must sum to one, so two dimensions suffice to represent the location of a population in a genotypic frequency state space. There are two "degrees of freedom." Similarly, though there are two alleles A and a, their frequencies p and q must also sum to one. So one dimension suffices for an allelic frequency space; there is one degree of freedom. The frequencies of the alleles are computable from the genotype frequencies by enumerating the alleles comprising each genotype.

(1) $p = D + H/2$
(2) $q = R + H/2$

These rules map a 2-dimensional onto a 1-dimensional space.

Many discussions of the units of selection focus on heterozygote superiority, where the fittest genotype of the three is Aa. As we noted above, this has become something of a test case for claims about the power of genic selectionism (Sober and Lewontin 1982, Sterelny and Kitcher 1988). For a change, we will begin with a different phenomenon, heterozygote inferiority. We begin by outlining a standard genotypic model of this kind of selection, and our conclusions about it are familiar. Nonetheless, our presentation will go through some detailed book-keeping, as the details of this book-keeping are essential to our later arguments.

We will suppose AA and aa are equally fit; $W_{AA} = W_{aa} = 1$. The fitness of Aa, W_{Aa}, is (1-s). Here s is a parameter measuring the penalty associated with the genotype Aa. These fitness differences should be taken to imply that some types are systematically favored over others, that there is a disposition or propensity of some sort underlying differences in reproductive success. Thus we are concerned with selection of some sort, rather than drift, though it is not yet established what the unit of this selection is. We assume, as is standard, that this penalty involves reduced viability rather than reduced fertility. Suppose that before mating the genotype frequencies are D, H and R. Individuals mate at random, and the six mating types have their frequencies displayed in Table 1.

TYPE OF MATING	FREQUENCY OF MATING	OFFSPRING		
		AA	Aa	aa
AA x AA	D^2	D^2		
AA x Aa	2DH	DH	DH (1-s)	
AA x aa	2DR		2DR (1-s)	
Aa x Aa	H^2	$H^2/4$	$[H^2/2]$ (1-s)	$H^2/4$
Aa x aa	2HR		HR (1-s)	HR
aa x aa	R^2			R^2

Table 1: Uniform heterozygote inferiority

With random mating, the chance of a mating of two dominants is just the chance of picking two individuals at random and finding them to be dominants, which is D^2. The frequencies in the second column sum to $(D + H + R)^2$ which is one.

For a moment, ignore the (1-s) terms in the Aa column, and consider the frequencies of the offspring types before selection. That is, consider the genotypic frequencies of the newly formed zygotes. These frequencies are determined by summing the columns for each offspring type. Without selection, these zygote frequencies would also be the adult genotype frequencies. The new frequency for dominants, D', would be:

$$(3) \qquad D' = D^2 + DH + H^2/4 = p^2$$

Similarly, if the Aa column had not been penalized by (1-s) each time, the new frequency H' of heterozygotes would be $2(D + H/2)(R + H/2) = 2pq$, and R' would be q^2. So whatever D, H and R were originally, if there is random mating and no selection, one generation of mating produces the Hardy-Weinberg equilibrium of genotype frequencies.

If a population is in Hardy-Weinberg equilibrium, there are rules transforming allele frequencies directly into genotype frequencies. The position of a population in 2-dimensional genotype frequency space is determined by its position in 1-dimensional allele frequency space; it is possible to move freely between the two spaces without loss of information. The applicability of the Hardy-Weinberg rule eliminates a degree of freedom.

Now consider the effect of selection against heterozygotes, represented by the (1-s) terms in Table 1. This knocks the adult population out of Hardy-Weinberg equilibrium. Also, the sum of the column for each offspring type is not the new frequency for that genotype, as these three totals no longer sum to one. However, this combination of mating and selection still has only one degree of freedom. To see this, note first that the (1-s) terms can be factored out of the sum in the column for Aa. This sum is $2(D + H/2)(R + H/2)(1-s)$. Further, the sums of the genotype columns can be converted into the new genotype frequencies by dividing each by the sum of all terms in the three offspring columns of the table, or $\bar{W}$:

$$(4) \qquad \bar{W} = (D + H/2)^2 + 2(D + H/2)(R + H/2)(1-s) + (R + H/2)^2 = 1 - 2pqs$$

The new frequencies of genotypes AA and Aa are:

$$(5) \qquad D' = p^2/[1 - 2pqs]$$

$$(6) \qquad H' = 2pq(1-s)/[1 - 2pqs]$$

These determine R'.

We now have formulas for the new genotype frequencies in terms of the old allele frequencies and parameter s. This space has one dimension, so this evolutionary process

has a required dimensionality of one. It easy to derive a formula for the new frequency of allele A, p'.

(7) $$p' = D' + H'/2 = p(1-qs)/(1 - 2pqs)$$

We also have a formula for the amount of change in the frequency of allele *A* in a generation, Δp:

(8) $$\Delta p = p' - p = pqs(2p - 1)/(1 - 2pqs)$$

Finally, we can say in allelic terms what this system will do. There is an unstable equilibrium at p=0.5. Though the system will sit still on this exact value of p, a slight perturbation in either direction will produce more change in that same direction until one allele is lost. The speed of this removal increases with s.

We have seen that a 1-dimensional allelic space can represent this evolutionary process. This result is well known, and opponents of genic selectionism do not generally challenge the dynamic adequacy of the allelic perspective here; rather, they claim that the gene's eye view of this process distorts its real structure. For instance, allele A has no unequivocal role of its own. It has a certain value when part of the AA genotype, and another when associated with a. This can be seen readily when the average or "marginal" fitnesses of the two alleles are calculated. These fitnesses are weighted sums of the two genotype fitnesses associated with each allele.

(9) $$W_A = pW_{AA} + qW_{Aa}$$
$$= 1 - qs$$

(10) $$W_a = pW_{Aa} + qW_{aa}$$
$$= 1 - ps$$

These quantities are dynamically useful; they can be used to predict the movement of the population in allelic frequency space. For Δp can be expressed in terms of p and the marginal fitnesses:

(11) $$\Delta p = pq(W_A - W_a)/[pW_A + qW_a]$$

However, the allelic fitness of A is frequency dependent; it is high when p is high and low when p is low. For Sterelny and Kitcher (1988) and Maynard Smith (1987) this is just a case of frequency dependent selection, akin to well-understood cases where genotypic fitnesses vary according to genotype frequencies, such as systems where rare types are at an advantage. Waters (1991) views it as a combination of frequency-dependent selection and selection in a spatially heterogeneous environment. For Sober and Lewontin (1982) on the other hand, these frequency-dependent allele fitnesses suggest the real process of selection here is a genotypic one, as the genotypic fitnesses are not context-sensitive in the way the allelic fitnesses are. Wimsatt (1980) and Lloyd (1988) make claims in a similar spirit.[3] We will not address these issues yet. For our immediate purposes, the genotypic and allelic models of this situation are as good as each

[3] Wimsatt's and Lloyd's views are discussed in Godfrey-Smith (forthcoming).

other, as they are both dynamically sufficient representations. But now we will present another example of selection against heterozygotes, in which the simple allelic perspective does not have such power.

Although our first example was schematic, the second is closer to reality. All human populations are polymorphic with respect to the Rh blood groups, which are controlled by a single locus with two alleles. Genotypes AA and Aa are Rh-positive (Rh^+) and aa is Rh negative (Rh^-). A is dominant over a, and it codes for an antigen on the surface of red blood cells. Antigens are markers recognized by the immune system. In most situations the genotypes are thought to be equal in fitness. The exception is the case of a heterozygote (Aa) fetus in a homozygous recessive (aa) mother. If an Rh^+ fetus is inside an Rh^- mother, the mother's immune system is often activated by the antigen on the fetus' blood cells, and the mother produces antibodies against her own fetus. In the absence of treatment this results in the death of the infant from anemia. The only way for an Rh^+ fetus to be conceived in an Rh^- mother is for the fetus to be a heterozygote, as it must receive one a allele from its aa mother. The father could be AA or Aa.[4] So there is selection against heterozygotes, as represented in Table 2. In the first column, females are on the left.

TYPE OF MATING	FREQUENCY OF MATING	OFFSPRING		
		AA	Aa	aa
AA x AA	D^2	D^2		
AA x Aa	DH	DH/2	DH /2	
Aa x AA	DH	DH/2	DH/2	
AA x aa	DR		DR	
aa x AA	DR		DR(1-s)	
Aa x Aa	H^2	$H^2/4$	$H^2/2$	$H^2/4$
Aa x aa	HR		HR/2	HR/2
aa x Aa	HR		[HR/2](1-s)	HR/2
aa x aa	R^2			R^2

4 We have simplified the situation. Often, the first heterozygote child of one of the dangerous mating types is not affected, but the second is. As the chance of having more than one heterozygote child is higher for an AA x aa mating than for an Aa x aa mating, the parameter s should not be the same for these two types of mating in Table 2. For a detailed examination of this system, see Feldman, Nabholz and Bodmer 1968.

Table 2: The Rh case

In this case more book-keeping must be done. We must distinguish between two types of AA x aa and Aa x aa matings. Only those in which the mother is aa result in selection against heterozygote offspring. The critical difference between this case and the previous one is the fact that it is not possible here to factor out the term (1-s) from the sum of the entries in the column for heterozygotes, because only some heterozygotes are affected. This makes a reduction to a 1-dimensional space impossible.

We will go through this result in detail. Consider the zygotes first. Whatever the old D, H and R might be, if mating is random the zygotes are in Hardy-Weinberg equilibrium. Then genotype frequencies are functions of allele frequencies. If there is selection, but selection is uniform within a genotypic class (as in Table 1) then whatever D, H and R were, the term (1-s) can be factored out and applied to the zygotic genotype frequency. Because this zygotic genotype frequency is a function of the allelic frequencies, the genotype frequency after selection is a function of the old allelic frequency (and s) as well. The dimensionality is 1.

In the Rh case selection is not uniform within the class of heterozygotes. Consequently, though the zygotes are in Hardy-Weinberg equilibrium before selection, the number of heterozygotes subject to selection depends on the frequencies of the mating types, the components of the sum which produces the zygotic frequency for heterozygotes. The frequencies of mating types depend on D, H and R.

In the Rh case the entries in the heterozygote column sum to $[H + R(2 - s)](D + H/2)$, which is $2pq - Rsp$. Note that if there is no differential mortality ($s=0$) this reduces to the heterozygote frequency under Hardy-Weinberg equilibrium. The sums of the columns for dominants and recessives are as they were in the previous case: p^2 and q^2. To attain the new genotypic frequencies D', H' and R', the sums of the genotype columns are divided by $\bar{W}$. $\bar{W}$ is $(1 - Rsp)$, so the new genotype frequencies are:

(12) $D' = p^2 / (1 - Rsp)$
(13) $H' = [2pq - Rsp]/ (1 - Rsp)$

These formulas cannot be reduced to expressions in terms of p and s. It is not possible to predict the new genotypic frequencies or allele frequencies in terms of the old allele frequencies. The allele frequency space has a dimensionality of 1, and the Rh selection process has a required dimensionality of 2.

It is still possible to track the population's movement through 1-dimensional allele space, if alleles are what we are interested in following. When we have computed the population's next step in genotypic frequency space, we can drop a dimension, and derive the new allele frequencies resulting from that step:

(14) $p' = D' + H'/2 = p(1 - Rs/2)/(1 - Rsp)$

But once we have dropped into allele space, we are dynamically stuck there. It is not possible to move back to a unique position in genotypic frequency space, and it is not possible to determine the population's next position in either space, from a given position in allele space.

Earlier we resisted the temptation to indulge in philosophical commentary on the status of frequency dependent allele fitnesses. The allele fitnesses were functions of p and s, and could be used to describe a dynamically sufficient model of the population's movement in both genotype and allele frequency spaces. In the Rh case the allelic fitnesses are also frequency dependent. They are not dependent on the allelic frequencies however, but on the genotype frequencies.

This is because the fitness of A is not constant within a genotypic combination. The marginal fitnesses of A and a can be derived with the "method of weights" of Lewontin 1958. With this method, the outcome of selection is described as if a uniform fitness-like weighting was applied to genotypic classes in Hardy-Weinberg equilibrium. That is, the sum of each offspring column is divided by its expected value under Hardy-Weinberg equilibrium with no selection. So the numerator of D' (D' un-normalized by $\bar{W}$) is divided by p^2, the numerator of H' by 2pq, and that of R' by q^2. The resulting expressions can be treated as genotypic fitnesses.

The homozygotes have a fitness of 1 and the heterozygotes have a fitness of (1 - Rs/2q). The marginal fitnesses can be derived in the usual way, with (9) and (10) above.

(15) $W_A = 1 - Rs/2$

(16) $W_a = 1 - Rps/2q$

As we claimed earlier, Δp can be expressed in terms of p and the two marginal fitnesses: (11) above. So while knowledge of the marginal fitnesses of alleles, and p, is sufficient to construct a dynamically sufficient model of the Rh case, the dimensionality of the representation is still 2. Consequently, though we have no argument with the claim that genic fitnesses can be used to represent selection (in our current sense), it is not true that this move eliminates the need for knowledge of genotype frequencies.

In the Rh case, it is selection which blocks a 1-dimensional representation of the evolutionary process. There are other ways of achieving this effect. The equivalence of genotypic and allelic representations in the first example discussed depends also on the randomness of mating, represented in the second column of Table 1. Much is made of the Hardy-Weinberg equilibrium as a kind of "basic law" of population genetics. But assortative mating for a trait is extremely common and may even be more common than random mating. Human populations, for example, mate assortatively by height, color, shape and so on. When mating is not random a model requires at least one extra parameter describing the mating pattern. Then even with no selection, genotype frequencies may not be expressible as functions of allele frequencies. Strictly, it is not Hardy-Weinberg equilibrium which is required to reduce the dimensionality, but some rule determining a unique position in genotype frequency space from a position in allele frequency space. Though the existence of such a rule is necessary, it is not sufficient, as the Rh case shows. Selection must also act in a certain uniform way. If either assumption is violated, a 1-dimensional model cannot represent the process.

4. Richer Allelic Spaces

As the allele frequency space we considered in the previous section was 1-dimensional, there is no way it can represent a process with a required dimensionality of

2. However, this is not because the space is an allele space. An allele frequency space with 2 dimensions can represent the Rh example just as well as the genotypic frequency space can. Such a space can be constructed with the aid of conditional allele frequencies.

A conditional allele frequency is the frequency of an allele given a certain genetic context. For example, f{A|A} is the frequency of A alleles within the class of alleles which are in combination with other A alleles; f{A|a} is the frequency of A within alleles in combination with a. These frequencies must be interpreted carefully: f{A|a} is not the frequency of Aa combinations. The frequency of Aa combinations is H. Rather:

(17) $H = 2\, f\{A|a\}q = 2\, f\{a|A\}p$
(18) $D = f\{A|A\}p$
(19) $f\{A|A\} + f\{a|A\} = 1$
(20) $f\{A|A\}p + f\{A|a\}q = p$

These equations are based on a standard rule for conditional probability; the overall probability of an event is the product of the conditional probability of the event given a certain condition and the probability of that condition, summed over all conditions. That is, $Pr(A) = \Sigma Pr(B_i)Pr(A|B_i)$.

We will show that both (i) a space of any one conditional allele frequency and any one unconditional allele frequency, and (ii) some spaces of two conditional allele frequencies, can represent anything a genotypic frequency space can represent. To show that these allele spaces can represent anything genotype spaces can represent, it is sufficient to show that a population's position in these allele spaces uniquely determines its position in genotype frequency space. For once the transition to genotypic frequency space has been made, the marginal allelic fitnesses can be derived, and used to find Δp. The allelic fitnesses will be functions of the conditional allelic frequency combinations described below.

The dimensions of dynamically sufficient allelic spaces can be:

(i) Any conditional allele frequency and any unconditional allele frequency:
Suppose you know p and f{A|a}. Then:

(21) $R = (1-p)(1- f\{A|a\})$
(22) $H = 2\, f\{A|a\}(1-p)$
(23) $D = 1 - H - R$

(ii) The two conditional frequencies of an allele:
Suppose you know f{A|a} and f{A|A}. Then if you can derive one unconditional allele frequency, you can use the results in (i) above to derive D, H and R.

(24) $p = f\{A|A\}p + f\{A|a\}(1-p)$
$= f\{A|a\} / [1- f\{A|A\} + f\{A|a\}]$

Not all spaces using conditional allele frequencies have two degrees of freedom. Suppose you know two conditional allele frequencies with the same conditionalization: f{A|A} and f{a|A}. These sum to one, and D, H and R cannot be derived from them.

Conditional allele frequency spaces of 2 dimensions can be used to represent anything a genotype frequency space can represent. Writers like Sterelny, Kitcher and Waters claim the gene's eye view can successfully mimic the usual framework of genotypic selection. We have shown that if the gene's eye view is of the same dimensionality as the genotypic, then this claim is true.

5. Dimensions and Causes

On the criteria used so far in this paper, there is no way to choose between allelic and genotypic representations of the selective processes that are standardly viewed as selection on genotypes. Either point of view can produce a dynamically sufficient model as long as the required dimensionality of the process is respected. In this section, we will discuss relations between these formal questions and other aspects of the units of selection debate.

The question remaining about the right way to view these selective processes is in part metaphysical. The dispute arises even in cases where all the empirical facts are in. Genic selectionism and genotypically-oriented orthodoxy agree that the world contains both genes and genotypes. They agree also that the fate of a gene is entwined with that of other genes in the same genome, and they agree that genes, rather than genotypes, are passed from generation to generation, in sexually reproducing organisms. They also agree about the dynamics of the cases we have discussed. They describe these dynamics differently, but there is no problem of incommensurability. It is easy to translate between one framework and another. There is still room for disagreement, however, about the extent to which the two models are faithful to the causal forces operating in evolution, and about which models provide the best causal explanations of various evolutionary processes. A Humean genic selectionist and a Humean proponent of genotypic orthodoxy might find nothing to argue about at this point. But we are not Humeans.

An initial line of objection to the genic perspective focuses on the status of the conditional allelic frequencies used in the genic reconstruction of 2-dimensional cases. Genotypes are, intuitively, natural properties of organisms, and plausible bearers of causal influence. In the Rh case we began with a story about an incompatibility between a maternal genotype and a fetal genotype, and we ended with a dynamic in the space of alleles. This might look like an elaborate dodge, a mathematical trick making use of allele frequencies which are "contaminated" with genotypic information. It might be claimed that using conditional allele frequencies is simply an evasive way of building a genotypic model.

We doubt if this argument can be made to work, in a general form. Genotypes are plausible natural properties of organisms. But genotype frequencies are properties of populations, not of individuals. Similarly, conditional allelic frequencies are properties of populations. The real entities being counted in a conditional allelic model are individual genes, which are as natural as genotypes, in a sexual population. For the genic selectionist, the dimensions of both the allelic and the genotypic models count genes. Both count genes in complicated ways; ways which preserve information about how the genes are grouped together in organisms. To claim that counting genotypes is something over and above counting genes as they appear in different contexts is to beg the question against the genic selectionist.

In a heterogeneous ensemble in which entities appear grouped in different subpopulations, the unconditional frequency of an entity in the ensemble is a weighted average of its conditional frequency in each subpopulation. From the realist standpoint, it is as plausible to claim that the weighted average frequency of an entity is an abstract summary of its "real" conditional frequencies in each of the various contexts it appears in, as it is to claim a privileged status for the unconditional frequency. A more appropriate question to ask is whether the frequency in the total ensemble is the causally salient property, and the division into subclasses accidental, or vice versa. The question of which kind of frequency is fundamental or basic depends on the forces operating in the population.

There is a symmetry here between fitnesses and frequencies. Unconditional gene frequencies like p are weighted averages, where the terms averaged are the frequencies of the allele in various contexts (see (20) above). A marginal fitness like W_A is a weighted average in the same sense. A genotypic fitness like W_{Aa} is an allele's fitness in a context, and these contexts are averaged over to obtain a marginal fitness. It is our view that in the case of both fitnesses and frequencies, there is no a priori ontological privilege attaching to either the conditional or the unconditional mode of accounting, but one or the other can be more causally appropriate in particular cases. We will illustrate this later.

The final reason to reject the view that conditional allelic frequencies are merely evasive ways to represent genotypic selective processes is the fact that there are some selective processes which are 2-dimensional, and which can be represented both with genotypic and with conditional allelic frequencies, but which are best understood causally as competition at the allelic level. The clearest way to present this point is with a hypothetical case, though it is related to a real example.

The case we will discuss next involves gametic selection, selection on eggs and sperm, which each carry half the normal number of chromosomes for the species. Suppose the only difference between alleles A and a is that sperm carrying a alleles swim faster than sperm carrying A alleles. Then in matings involving heterozygote males, it there is a fertilization, it is more likely to be achieved by an a-bearing gamete. Homozygotes for A are not significantly less fertile than males with a-bearing sperm however, and once fertilization has occurred, the alleles behave identically. In Table 3, the advantage possessed by a-bearing sperm is expressed with parameter m. Normally, A and a would have even chances of achieving a fertilization. Here a-bearing sperm achieve m of the fertilizations and A-bearing sperm achieve $(1-m)$, where $m>0.5$. Ignore for the moment the asterisks against some entries.

TYPE OF MATING	FREQUENCY OF MATING	OFFSPRING		
		AA	Aa	aa
AA x AA	D^2	D^2		
AA x Aa	DH	DH(1-m)	DHm	
Aa x AA	DH	DH/2	DH/2	
AA x aa	DR *		DR **	
aa x AA	DR *		DR	
Aa x Aa	H^2	$(1-m)H^2/2$	$H^2/2$	$mH^2/2$ *
Aa x aa	HR *		HR/2 **	HR/2 **
aa x Aa	HR *		HR(1-m)	HRm
aa x aa	R^2*			R^2 **

Table 3: Gametic selection

We will go though some book-keeping as before. To find the new frequency of A homozygotes, we sum the entries in the first column and divide by $\bar{W}$. Consider first the sum of entries in the first column:

(25) $D' = p(D + H(1-m))$

If m=1/2, as in the absence of gametic selection, this expression reduces to p^2. As m increases D' decreases. Expression (25) gives the new frequency of AA homozygotes, as$\bar{W} = 1$.

The method of weights can be used to derive the marginal fitnesses of the alleles.

(26) $W_A = [p + D + H(1-m)]/2p$

(27) $W_a = [q + R + Hm]/2q$

(28) $\Delta p = H(1/2 - m)/2$

As in the Rh case, these expressions cannot be reduced to functions of p and m. The process is 2-dimensional. In the last section we showed that genotypic models can always be replaced by models using conditional allelic frequencies. Such a replacement can be

performed in both the Rh and the gametic selection cases.[5] In the present case the switch is trivial:

(29) $\Delta p = (1-p)\, f\{A|a\}(1/2 - m)$

The translations into a conditional allelic framework have a different status, however, in the Rh and the gametic selection cases. In the Rh case, the force driving the population through genotypic and allelic spaces was an elevated death rate in individuals of a certain genotype born of certain matings. In the gametic selection case, the process is a competition between alleles and their associated gametes. The a allele produces certain properties in gametes bearing it -- a high swimming speed -- and this produces a reproductive advantage for a against A. The entities engaging in this causal process do not even have diploid genotypes; sperm are either a or A, not AA, Aa or aa. We assume that being of genotype AA, Aa or aa has no effect on the fitness of the individual. Now, the advantage of the a allele is only manifested in a certain kind of context -- matings involving Aa males. The number of situations in which the advantage of a is manifested, and hence the rate of change in the frequency of a, depends on the frequency with which heterozygotes mate. As a consequence, the process has a dimensionality of 2. In the Rh case the process was 2-dimensional because of the complex way in which some individuals, characterized by genotypes, are penalized. In the gametic selection case the process is 2-dimensional because of the complex nature of the contexts in which gametes, characterized by their alleles, compete. So while it may be that the switch to the allelic point of view distorts the causal structure of the Rh case, in the gametic selection case the genotypic perspective is the ontologically dubious one, if any is, and the conditional allelic model seems to reflect more faithfully the underlying mechanics of the process.

A phenomenon closely related to gametic selection is segregation distortion. Segregation distortion occurs when the proportions of germ cells produced by a heterozygote individual are unequal. The basic dynamics of segregation distortion are the same as above. The cause for the bias in favor of a is different -- different sperm phenotypes, in gametic selection, and different sperm proportions, in segregation distortion -- but the effects are the same. Segregation distorters should then drive their rival alleles out of the population. A large number of segregation distorters have been observed in nature, however. The general reason is that the bias in favor of a during meiosis is counteracted by selection against aa homozygotes. The segregation distorters of the t allele family in the mouse, for instance, are either lethal or cause male sterility when homozygous (Lewontin and Dunn 1960, Lewontin 1962). As a consequence, distorters are kept in intermediate frequency in almost all mouse populations.[6]

The effects of opposing selective forces are represented by the asterisks in Table 3. The most philosophically significant case is that in which the distorter allele a is lethal in homozygous form. In the second column of Table 3, some of the mating frequencies are

---

5 There is a third way to construct a 2-dimensional model of this kind of system. Several models of segregation distortion and similar phenomena use as variables the frequency of the distorter allele in the two types of gamete (see Lewontin 1968, Feldman and Otto 1991, and references in the latter paper).

6 The equilibrium values in house mouse cases may also be due to inter-demic selection against populations in which all the males are tt (Lewontin 1962).

marked with a single asterisk. These are matings in which either the male or female, or both, is aa. If the aa combination is lethal, these mating frequencies are set at zero, and all the offspring entries in those rows are zero. Another offspring entry is also zero: aa offspring of Aa x Aa matings. Note that when the table is compressed in this way, the entries in the second column again sum to 1: $(D + H)^2 = 1$.

The essential point is the effect this selective regime has on the relation between allele frequencies and genotype frequencies. Because there are no aa individuals ($R=0$), all the a alleles are found in heterozygotes. Consequently, $q = H/2$, while $p = D + H/2$ as before. It is possible to move between the spaces of allele and genotype frequencies without loss of information. The dimensionality is 1.

(30) $W_A = [1 + p - 2qm]/2p$

(31) $W_a = 1/2 + m(p - q)$

Since it is 1-dimensional, this evolutionary process can be represented in an unconditional allelic frequency space. Though it can be modelled with a simpler framework than gametic selection or segregation distortion alone can, it is causally more complex. Evolution in the previous gametic selection case is driven by a single competitive process: a-bearing sperm are more likely to achieve a fertilization than A-bearing sperm, in a heterozygous individual. In the case of a balance between segregation distortion and aa lethality, we have the original competition between gametes plus differential mortality at the level of genotypes. In moving from segregation distortion alone to a balanced mix of segregation distortion and selection on genotypes, we have added a causal force but dropped a dimension. The reduction in dimensionality is achieved by the fact that the new causal force simplifies the way genes are collected into genotypes.

Further, the new causal force is different in kind to the old one. In the case of pure gametic selection and (perhaps not so obviously) segregation distortion, the selective process sorts gametes, entities which do not have diploid genotypes. If aa is lethal, however, this is a fact about the prospects of diploid individuals. Though the unconditional allelic space is sufficient to represent the dynamics of this process, it is not as causally faithful to it as a genotypic perspective is. On the other hand, simple competition between gametes requires, dynamically, a genotypic or conditional allelic space, though here the individual gene's eye view is more causally reasonable.

We pause to express a general conclusion: The required dimensionality of a system is one thing, and its causal structure another (see also Sober 1981). The "units of selection" question about which processes can be represented in which frameworks is quite distinct from the "units of selection" question about which kinds of entities are doing causal work.

This conclusion is supported by additional facts about these scenarios. The combination of segregation distortion and selection on genotypes is only 1-dimensional in the extreme case where all aa individuals die before reproduction. That is, parameter s is equal to 1. If parameter *s* is less than 1, then Table 3 cannot be simplified by discarding all offspring of matings involving aa individuals and discarding aa offspring of heterozygote matings. If $s<1$, all the entries must be counted as usual, though the entries

in the final column are penalized by (1-s). Here the simple mapping between gene and genotype spaces is not available, and the dimensionality is 2.

Here is a summary of the cases discussed.[7]

MODE OF GENIC SELECTION	MODE OF GENOTYPIC SELECTION	PARAMETER VALUES	DIMENSIONALITY
None	Uniform heterozygote inferiority	All	1
	Rh case	All	2
Segregation distortion or gametic selection	None	All	2
	Balancing viability selection, against aa	All values of m, s = 1	1
		All values of m, s < 1	2

Table 4: Summary

Earlier we discussed the objection that the only legitimate genic models use unconditional allelic frequencies, and that conditional allelic frequencies are a dodge. From this viewpoint, when s<1 combinations of segregation distortion and selection on genotypes should be modelled genotypically (or in the way outlined in note 5). But when s reaches 1, an allelic framework is available. There are two points to note here. Firstly, it would be unwise to infer from this change in state space that there is a change in the level of selection when s reaches 1. Secondly, it appears to us that conditional allelic frequencies are useful in situations like this, as a sort of intermediate between the genotypic view and the simple, 1-dimensional genic perspective. If alleles are to be counted, then when s<1 alleles must be counted in more complicated ways than when s=1. As s approaches 1, the correction made for the very small number of aa individuals gets less and less significant, and the conditional frequency f{ala} approaches 0. If conditional allelic frequencies are admissible, there is no need to switch between entities as s moves between 1 and lesser values.

In this section we have focused on a negative point: mathematical facts about the required dimensionality of a process are distinct from causal facts about the units upon which selection operates. But if questions of dimensionality will not answer all units of

---

7 It is important that all the genotypic selective forces discussed concern viability, rather than fertility. Generally selection on fertility is more complex to model than selection on viability. We will not work through an example of fertility selection, but if the offspring entries with double asterisks in Table 3 are set at zero, we have total sterility of aa males. This case has causal similarities to that of aa lethality -- an advantage at the genic level is balanced by a disadvantage at the genotypic level -- but it must be modelled in a 2-dimensional space.

selection questions, what will answer them? How does one determine what the "real" causal forces operating in a selective situation are? As our chief aim is to distinguish these questions from those about dimensionality, we will not attempt to solve this problem here (see Sober 1984 for one attempt). But we will make some brief suggestions.

First, one can examine the role played by parameters in a model, and the relation between parameters and the world. What is being measured by m, or by s? In both the uniform heterozygote inferiority case and the Rh case, parameter s measures the systematic loss of individuals of a certain genotype. Looking across any row in which an s appears, we see that the mating in question has produced fewer surviving individuals than would have resulted if penalty s had not been applied. Parameter m behaves differently. It is true that m only has its influence in matings which involve Aa males (hence the 2-dimensionality), but the total number of offspring left by these matings (the number of offspring in a row of the table) is not reduced or increased by m. The parameter alters the distribution of genotypes produced by the mating, with no net loss of individuals. This is reflected by the fact that in the pure gametic selection case evolution occurs though $\bar{W} = 1$. In the terminology of section 2, if gametic selection is the only force then being an Aa male has no influence on the reproduction by an individual. It does affect what this is reproduction of. On the other hand, in the last case discussed, being aa has dire effects on the reproduction by an individual; effects so dire that the question of reproduction of is idle.

Secondly, there is a general principle which explains why considerations of dynamic sufficiency will not solve all units of selection problems, and which also goes some way to explaining why it can be hard to determine the real locus of causal power even when there appear to be no empirical disputes left. In general, if a dynamically sufficient description of a system can be fashioned with entities that are combinations or collections of lower-level constituents, and in which the state descriptors are frequencies or numbers of entities, it is always possible to find a representation with the same number of dimensions in which the entities are the lower level constituents and the state variables are their conditional frequencies. That is, when modelling a multi-level system there is always a choice between counting the higher-level units (such as genotypes) and counting the lower level components (genes) in a way sensitive to facts about their appearance in combinations. One way to view this is as the difference between thinking in terms of conjunctive probabilities -- Pr(A&B) -- and conditional probabilities -- Pr(A|B). Conjunctive probabilities are naturally taken as basic when whole higher-level units are causally salient, and there is no point in thinking in terms of a lower level constituent acting differently in various contexts. A biological illustration is found in the case of evolution in a diploid species which reproduces asexually. If there is no sex, then the three genotypes AA, Aa, and aa have entirely distinct evolutionary roles and, as far as population dynamics are concerned, there is no need to recognise common components such as A appearing in different contexts in the population. At no point is there any reshuffling of genic "components" to produce genotypes; reproduction by is the same as reproduction of.

The opposite is found in the pure gametic selection case. Here there is a competition between the components A and a, alleles which produce different sperm phenotypes. The evolutionary effects of this difference in sperm are determined by the frequency of the contexts in which A and a appear, so conditional rather than unconditional frequencies

must be used, but the causal process driving evolution in this case is not a competition between higher level combinations such as AA and Aa, but a competition between lower level units which is influenced by the distribution of these lower level units across different contexts. Here the conditional mode of accounting is preferable to the conjunctive.[8]

Between these clear cases there are many uncertain ones. The question of how often genotypes and other higher-level combinations interact as units, and how often the work is done by genes and other lower-level components having their own roles conditioned by context, is in part a live biological issue. For instance there may be uncertainty about the Rh case, discussed in section 3. In our presentation, we stressed that here evolution is driven by the differential mortality of individuals, characterized by genotypes. But at least part of the causal story can be told in genic terms: The presence of an A allele causes both the formation of an A antigen and the suppression of anti-A antibodies. A mother lacking the A allele produces the antibodies (and not the antigen). The mother's antibodies react with the fetus' antigen, producing anemia. This much can be said without mentioning diploid genotypes. It is a story about the effects of alleles in different contexts. Now this is only part of the story; the Rh case constitutes an instance of natural selection because the result of this process is the death of individuals of a certain genotype. There is no penalty for being A or a simpliciter. However, we do not pretend the right causal story is obvious, in these complex cases, and perhaps in some situations there is no fact of the matter. But the fact that there is no privileged state space, when the aim is simply dynamic sufficiency, does not imply there is no privileged causal story. Conventionalism about one theoretical project need not imply conventionalism about another.

* * *

References

Dawkins, R. (1982), The Extended Phenotype. Oxford: Oxford University Press.
Feldman, M.W., M. Nabholz and W.F. Bodmer. (1968), "Evolution of the Rh polymorphism: A Model for the Interaction of Incompatibility, Reproductive Compensation, and Heterozygote Advantage", American Journal of Human Genetics 21: 171-193.
Feldman, M.W., and S.P. Otto. (1991), "A Comparative Approach to the Population Genetics Theory of Segregation Distortion", American Naturalist 137: 443-456.
Godfrey-Smith, P. (forthcoming), "Additivity and the Units of Selection" to appear in PSA 1992 Volume 1.

---

8 Waters claims that the choice between genic and higher level models of selection is a choice about "where to draw the conceptual divide between the environment and the selected domain" (1991, 571). Our distinction between conjunctive and conditional accounting resembles Waters' framework here. Our differences from him concern the possibility of objective decision about whether conditional or conjunctive accounting is more appropriate in particular cases.

Lewontin, R.C. (1958), "A General Method for Investigating the Equilibrium of Gene Frequency in a Population", Genetics 43: 419-434.
Lewontin, R.C. (1962), "Interdeme Selection Controlling a Polymorphism in the House Mouse", American Naturalist 96: 65-78.
Lewontin, R.C. (1968), "The Effect of Differential Viability on the Population Dynamics of t Alleles in the House Mouse", Evolution 22: 262-273.
Lewontin, R.C. (1974), The Genetic Basis of Evolutionary Change. New York: Columbia University Press.
Lewontin, R.C. and L.C. Dunn. (1960), "The Evolutionary Dynamics of a Polymorphism in the House Mouse", Genetics 45: 705-22.
Lloyd, E. (1988), The Structure and Confirmation of Evolutionary Theory. New York: Greenwood Press.
Maynard Smith, J. (1987), "How to Model Evolution", in J. Dupré (ed.) The Latest on the Best: Essays on Optimality and Evolution. Cambridge, MA: MIT Press, pp.119-131.
Sober, E. (1981), "Holism, Individualism, and the Units of Selection", PSA 1980, volume 2. East Lansing: Philosophy of Science Association, pp.93-121.
Sober, E. (1984), The Nature of Selection. Cambridge: MIT Press.
Sober, E. and R.C. Lewontin (1982), "Artifact, Cause and Genic Selection", Philosophy of Science 49: 157-80.
Sterelny, K. and P.S. Kitcher. (1988), "The Return of the Gene", Journal of Philosophy 85: 339-361.
Waters, K. (1991), "Tempered Realism about the Force of Selection", Philosophy of Science 58: 553-573.
Williams, G.C. (1966), Adaptation and Natural Selection. Princeton: Princeton University Press.
Wimsatt, W.C. (1980), "Reductionist Research Strategies and their Biases in the Units of Selection Controversy", in T. Nickles (ed.) Scientific Discovery: Case Studies. Dordrecht: Reidel, pp. 213-259.