P. KYLE STANFORD

THE UNITS OF SELECTION AND THE CAUSAL STRUCTURE OF THE WORLD

ABSTRACT. Genic selectionism holds that all selection can be understood as operating on particular genes. Critics (and conventional biological wisdom) insist that this misrepresents the actual causal structure of selective phenomena at higher levels of biological organization, but cannot convincingly defend this intuition. I argue that the real failing of genic selectionism is pragmatic – it prevents us from adopting the most efficient corpus of causal laws for predicting and intervening in the course of affairs – and I offer a Pragmatic account of causation itself which ultimately bears out the claim that genic selectionism misrepresents the causal structure of selective contexts.

1. INTRODUCTION

The units of selection debate concerns what entities are appropriately regarded as subject to the force of natural selection in evolutionary contexts. A number of thinkers, including most famously G. C. Williams (1966) and Richard Dawkins (1976, 1982), have argued for genic selectionism: the claim that selection acts exclusively upon particular genes.[1] By contrast, Elliot Sober (1984; 1994, with Wilson; 1998, with Wilson), Richard Lewontin (1982, with Sober), David Sloan Wilson (1994a, with Sober; 1994b, with Sober), Elisabeth Lloyd (1988, 1989), Stephen J. Gould (1980), Ernst Mayr (1963), Robert Brandon (1984), and William Wimsatt (1980, 1981) have all variously argued that genic selectionism is misguided and/or that selection can or does act on entities at a variety of different levels of biological organization: chromosomes, individual organisms, and groups of organisms, in addition to genes.[2] These critics of genic selectionism seem to share the intuition that to represent genes as the only entities subject to natural selection is to somehow *mis*represent the causal facts of the matter about selective processes.

Even these critics acknowledge that genic selectionism can accomplish what they call the 'bookkeeping' of natural selection, however: that is, given a set of gene frequencies, the genic selectionist can assign fitness coefficients to each allele and use the equations of population genetics to correctly predict the distributions of alleles that will be produced by natural selection in subsequent generations. But they insist that such formal

Erkenntnis **54:** 215–233, 2001.

manipulation simply obscures the fact that evolutionarily significant properties arise at a variety of levels of biological organization. Consider the classic case of industrial melanism in peppered moths: when industrial pollution darkens the trees in an English forest, melanic moths acquire a selective advantage over those who are not. But it is the moths and not the genes which are melanic, and we are inclined to say that it is being melanic which confers a particular selective advantage or disadvantage upon the bearer of that property in a particular environment. This is not to deny, of course, that properties at higher levels of biological organization are *themselves* often caused by particular genetic constitutions (in the right environments); nonetheless, the likelihood of survival of any particular moth seems to be impacted most directly by its coloration, and thus only indirectly by its genetic constitution. Thus, the critics insist that in construing all selection as the impact of particular genes in particular environments, the genic selectionist story is leaving out a crucial piece of causal information: the level of organization at which the evolutionarily significant properties actually arise.[3]

But genic selectionists like Kim Sterelny and Philip Kitcher (1988) and Kenneth Waters (1991) have offered the powerful rejoinder that these intuitions about the genuine causal efficacy of various properties at various levels of biological organization *simply reflect our prejudiced unfamiliarity with the genic point of view*. In cases like melanism, they suggest, the intuition that phenotypic properties are the evolutionarily relevant ones results from our deployment of an orthodox conception of the environment of causal interaction, rather than the one which is appropriate for the genic perspective. If we take the gene's-eye point of view seriously, we see that a crucial feature of a gene's environment is the other genes (including the allelic partner at its own locus) with which it is combined. Thus, any story that we are inclined to tell in terms of selective advantages caused by (for example) phenotypic properties can be straightforwardly translated into a corresponding genic selectionist story that treats an allele *in the presence of the relevant sort of genic environment* as the cause of that same selective advantage. Thus, the claim that it is melanism which confers a selective advantage upon moths, given a particular organismic environment, can be reformulated in the genic idiom as the claim that it is the presence of a particular allele which confers this benefit, given an environment including both genetic factors (like its allelic partner) and non-genetic ones (like pollution-darkened trees).

The availability of such genic schemes of translation is the crucial point behind Sterelny and Kitcher's (1988, 347, 354) insistence that if we are to evaluate genic selectionism fairly, we must apply it *consistently*, and

accept new conceptions of what count as environments and populations. Thus they reject Gould's (1980) and Mayr's (1963) claim that selection 'sees' organisms rather than genes as well as Brandon's (1984, 1990) more sophisticated argument that phenotypic properties screen off genotypic ones: such arguments depend upon conceiving of environments in the traditional (that is, biased against genic selectionism) way. The claim that the genic perspective must be applied consistently if we are to evaluate it fairly also grounds Waters's (1991, 563n) criticism of Wimsatt and Lloyd's approach:[4] applying the analysis of variance presupposes (as Lloyd recognizes) a particular construal of the relevant environment; thus, Waters argues, genic selectionists can always individuate genic environments so as to convert nonadditive into additive variance in the corresponding gene population.

Sober and Wilson's (1994; see also Wilson and Sober 1994a) recent proposal to identify units of selection in terms of 'shared fate' (or 'common fate') and differential fitness (cf. Sober and Wilson 1998, Ch. 3) begs the question against the consistent application of the genic perspective in just the same way. They suggest that we identify units of selection on a trait-by-trait basis, proceeding up the hierarchy of biological organization and asking at each level (genes, individuals, groups, populations, etc.) whether the individual entities at that level within a single group differ in fitness with respect to the given trait.[5] If not, the entities at that level have a 'shared fate' and are part of the same selective vehicle; the level at which the individuals within a single group differ in fitness is the level at which selection is really operating with respect to the trait in question. But this proposal picks out the traditional units of selection in a given context only by construing environments in the traditional, question-begging way. To see this, notice that Dawkins, in his (1994) commentary on the proposal, simply reformulates Wilson and Sober's claims about particular cases of selection acting above the level of the gene in the genic selectionist's idiom, as claims about the relative fitnesses of genes in genic environments. Sober and Wilson elsewhere insist that this sort of genic causal redescription "is not relevant to the empirical problem of deciding what types of adaptation are found in nature" (1994, 552), but this is simply to repeat and not to defend the assertion that genic selectionism misrepresents the causal structure of selective phenomena.

Sterelny and Kitcher's conclusion is that the units of selection debate is a tempest in a teapot. They hold that our choice of description or conceptual representation for selective processes is a matter of convention, although the genic language has the systematic advantage of always being available:

> Just as conventionalists have insisted that there are alternative accounts of the phenomena which meet all our methodological desiderata, so too we maintain that selection processes can usually be treated, equally adequately, from more than one point of view. The virtue of the genic point of view, on the pluralist account, is not that it alone gets the causal structure right but that it is always available. (1988, 359)

Such conventionalism also seems to be the upshot of Waters's claim that we must 'temper' our realism about the forces of selection:

> We can no longer maintain that a true description of a selection process provides a uniquely correct identification of the operative selective forces and the levels at which each impinges. Instead, we must accept the idea that the causes of one and the same selection process can be correctly described by accounts which model selection at different levels. (1991, 572)

Sterelny and Kitcher go on to argue that the genic representation of selective phenomena is universally available because it applies to cases in which the selective advantages of some alleles over others have nothing to do with the phenotypic properties of organisms. In the case of meiotic drive genes or segregation distorters, for example, one allele undermines its allelic partner in meiosis, gaining an advantage in the Mendelian lottery, and is therefore disproportionately represented in the surviving gametes. Likewise, so-called 'selfish DNA' or 'jumping genes' (genes which have no organismic function, but which increase their representation because they are good at inserting additional copies of themselves into other parts of the genome) may multiply and spread in virtue of a selection process which takes place at the molecular level. There is no story about selection for advantageous traits of organisms which can accommodate phenomena like segregation distorters[6] and jumping genes, but these cases can easily be described in terms of the selective advantages of particular genes, a language uniformly available, insist genic selectionists, for correctly describing all other cases of selective phenomena.

Thus, the case against genic selectionism will require more than vague intuitions about the genuine causal agents and forces at work in selective contexts. The critics must explain why we should not join the genic selectionists in identifying *those very intuitions* as a product of unfamiliarity with the genic perspective, eschewing the idea that there is any one real or privileged causal structure in selective contexts, and embracing the genic perspective as one whose universal availability promises a more systematic treatment of all selective phenomena.

2. CAUSATION AND UNITS OF SELECTION

To accomplish this task, critics of genic selectionism will have to provide some convincing rationale for identifying and individuating 'real' or 'genuine' causal agents and forces that shows just why the genic selectionist's procedure is illegitimate. Sober's classic and influential (1984) effort to develop such an account begins with the Pareto formulation of his contextual unanimity criterion: the claim that a true cause must raise the probability of its effect in at least one set of the relevant[7] background conditions and must not lower it in any such set. He then proceeds to argue that the phenomenon of heterozygote superiority illustrates why genes will not always satisfy this condition. The classic case of heterozygote superiority involves sickle-cell anemia in human beings: homozygotes for the sickling allele (*aa*) have anemia, while homozygotes for the normal allele (*AA*) are not anemic but are instead susceptible to malaria. Heterozygotes (*Aa*), however, avoid serious anemia while possessing malarial resistance as well (clearly the fittest combination in an environment of pandemic malaria). Sober points out that in this case of heterozygote superiority, neither of the allelic forms of the gene satisfies the Pareto contextual unanimity criterion: the normal allele lowers fitness when paired with another normal allele and the sickling allele lowers fitness when paired with another sickling allele, but each raises fitness when paired with the other. Thus, no allele qualifies as a genuine cause of any selective advantage or disadvantage, for none raises the probability of such an effect in at least one set of background conditions without lowering it in any. According to the critics (see Sober 1984, Chs. 7–9; Wilson and Sober 1994b; Sober and Wilson 1994, 1998), such cases illustrate that genic selectionism is able to balance the books – correctly predicting the distributions of alleles produced by natural selection from one generation to the next – only by *averaging together* the outcomes of the various *distinct* selective processes in which each gene is involved: if a particular allele produces sickle-cell anemia in combination with one allelic partner and malarial resistance in combination with another, the genic selectionist is simply averaging the fitnesses of the two outcomes together (discounted by the frequency of their occurrence) to get the overall fitness of the allele in question.

One problem with this analysis (see Sterelny and Kitcher 1988, 344–345; Waters, 1991, 567–568) is that virtually *no* genuine cause in evolutionary theory satisfies this Pareto contextual unanimity condition. Returning to the famous melanic moths, for example, we find that melanism itself fails Sober's criterion in this case: after all, there are *some* relevant background contexts, like the few remaining light trees or light patches

on trees, in which it confers a selective *disadvantage* to be melanic (i.e., in which melanism hurts a moth's chances of surviving and reproducing). Melanism is not an advantage for the moths in every background context, it is revealed to be an advantage only on the whole and for the most part, that is, only when we *average over* the existing environmental backgrounds. Genic selectionism cannot be condemned for assigning fitnesses by averaging over environments in reckless disregard of contextual unanimity, then, for *the* classic case of evolutionary explanation, industrial melanism, shows why any reasonable account of causation in selective contexts must do so.

Most damaging to the critics' case, however, is the fact that a genic scheme of translation *does* preserve the causal complexity of allegedly problematic cases like heterozygote superiority after all (cf. Sterelny and Kitcher 346–347). Genic selectionism's opponents accuse it of obscuring the fact that the fitnesses of the *A* (normal) and *a* (sickling) alleles are each an averaged product of distinct causal processes (degree of anemia and degree of malarial susceptibility) operating upon *pairs* of alleles (*AA*, *aa*, and *Aa*). But the genic perspective is perfectly capable of recognizing and describing the distinct contributions of these processes to the averaged fitnesses of *A* and *a*. The average fitness of *A*, for example, is generated by two distinct causal processes: one is the selectively advantageous effect of *A* in an environment that includes *a* as its allelic partner, and the other is the selectively disadvantageous effect of *A* in an environment that includes *A* as its allelic partner. Thus, precisely *because* any claim about a causal process at some higher level of biological organization can be straightforwardly transformed into a corresponding claim about the effect of a particular allele in a particular allelic environment, genic selectionism is indeed able distinguish the separate causal processes contributing to the average fitness of a particular allele at a particular locus,[8] and the critics' complaint misfires.

Despite serious and sustained efforts, then, the critics have failed to parlay the tenacious intuition that genic selectionists are distorting the causal structure of selective phenomena into a charge against genic selectionism that will stick.[9] Our general predicament is illustrated well by Godfrey-Smith and Lewontin's (1993), in which the authors demonstrate with elegance and precision precisely *how* to enrich genic models to represent all selective processes (viz., by using conditional allele frequencies), but proceed to complain that this result does not solve the units of selection problem, insisting that "there is more to an understanding of evolution than having a dynamically sufficient model" (p. 374). They maintain that genic selectionism ignores distinct questions about the 'reality' of causal forces,

but they have no substantive case to make (as they seem to realize, see p. 393) as to why the genic selectionist's account of the relevant causal forces in selective contexts is mistaken. They can only insist stubbornly that genic selectionism is inconsistent with our intuitions about the actual causal agents and forces at work in selective contexts. We remain in need, however, of some defense of the claim that this conflict represents some genuine failing of genic selectionism, rather than simple unfamiliarity with the genic perspective.[10] The rest of this paper will seek to provide such a defense.

3. THE PRAGMATIC THEORY OF CAUSATION

Contra Sober and Wilson (1994) and others, the genic selectionist's point does not *simply* concern the possibility of "representing" selective phenomena in genic language. A substantive question remains concerning the defensible general constraints on the individuation of causal agents and forces: we want to know whether there is anything wrong with accepting genes in genic environments (rather than entities at a variety of levels of biological organization) as the *genuine* causes in all biological selective contexts, and if so, precisely what it is.

Posing the problem in this way suggests that we cannot respond to the genic selectionist without considering how to identify genuine causal agents and forces in general. I have recently proposed a Pragmatic Theory of Causation (Stanford forthcoming) which allows us, I suggest, to see precisely what is illegitimate about the genic selectionist's proposal to construe all selection as operating on particular genes. Briefly, it turns out that genic selectionism prevents us from adopting as our corpus of causal laws the most *useful* set of rules we could employ in predicting and intervening with regard to selective phenomena, and, on the Pragmatic Theory of Causation, this *just is* to misrepresent the causal structure of those phenomena. Let us begin, then, with a sketch of the Pragmatic Theory itself.

Suppose we take for granted the notion of a property and of what it is for a property to be instantiated in a particular time and place. Let us then define an *event* as the instantiation of a collection of properties in the same place at the same time, and an *event-class* as a set of events whose members all share some particular set of properties, what we might call the *qualifying properties* of the class. With this machinery in place we may conceive of a *causal law* as a rule for predicting the occurrence of a member of one event-class from the occurrence of a member of another, or for intervening to produce a member of on event-class by producing a member of another.

Let us proceed to ask what characteristics we might find pragmatically desirable in the corpus of causal laws we choose to adopt – what characteristics will make a corpus of rules the *best* one we could have, pragmatically speaking? The preeminent desideratum, it would seem, is that our corpus of rules be *maximally effective* in predicting and intervening in the course of events, which is to say that it must allow us to predict and intervene successfully with respect to every single event which will ever occur throughout the course of our collective experience, or (should no such set of rules prove possible) with respect to as many such events as possible.[11]

Of course, there will be many different possible sets of causal laws that are maximally effective. For one thing, given any maximally effective set of causal laws, innumerable equally effective alternative sets can be generated from it by adding 'irrelevant' rules whose antecedent clauses will never or could never be fulfilled. More generally, for any given maximally effective corpus of causal laws there will be innumerable predictively equivalent alternatives whose rules are nonetheless more convoluted, complex or unsystematic than those of the original[12] (along with, perhaps, some predictively equivalent alternatives whose rules are more simple and systematic). These possibilities illustrate that effectiveness is not the only pragmatic consideration that concerns us: for any given level of effectiveness in prediction and intervention, a more *efficient* set of rules will clearly be preferable to a less efficient one. That is to say, what we *really* want to have is the single most simple and systematic member *of* the various maximally effective candidate sets of rules for prediction and intervention. And the suggestion of the Pragmatic Theory of Causation is that the set of laws picked out in this way – what we might call the *Pragmatic Corpus* of laws – just is the set of causal laws true in our world.

This approach to causal theorizing shares the metaphysically deflationary ambitions of Hume's classic constant conjunction account of causation and resists the current fashion for brute metaphysical postulation (whether of primitive natural necessitation between universals (Armstrong 1983; Tooley 1987), nature's causal capacities (Cartwright 1989), literally existent possible worlds (Lewis 1973a, b), or what have you), but it avoids the implausible implications that undermine Hume's account and its deflationary successors[13] (for details, see [Stanford, forthcoming]). Its Pragmatic dimension consists in holding our concept of causation to be directed towards picking out just that set of rules or patterns in the course of events that would be of the greatest *utility* to creatures *like us*, i.e., with just *our* set of practical needs, cognitive limitations, location in the ongoing temporal stream of experience, etc. The Pragmatic Theory's insistence on

the importance of both effectiveness and efficiency recalls David Lewis's proposal (1973b) that strength and simplicity trade off in picking out the causal laws true at a world and Philip Kitcher's suggestion (1989) that the causal structure of the world is established by tradeoffs among the fecundity and stringency of the argument patterns needed to explain it. These proposals are ultimately undermined, however, by the fact that there is simply no 'common currency' in which gains and losses in their respective competing desiderata can be compared, so neither offers any concrete account of how the causal laws or relations (at our world or any other) are actually fixed. Thus, it is important to recognize that the Pragmatic Theory does *not* trade off considerations of effectiveness and efficiency: instead, effectiveness is the dominant pragmatic consideration and the Pragmatic Corpus is simply the single most efficient *of the* various maximally effective sets of rules we could be using to predict and intervene in the course of events.

We would do well to wonder what the Pragmatic Theory's demand for efficiency is supposed to amount to, especially in light of the troubled philosophical history of related notions like simplicity. As a general matter, a corpus of rules will be more efficient as it is *easier for us to learn and to apply*. The Pragmatic Theory thus embraces the implication that the set of causal laws thus generated will be in some sense relative to us and our cognitive constitutions (quite broadly construed[14]): after all, the Pragmatic Theory insists that we find ourselves with a need and a use for our concept of causation only because we are the kinds of creatures we are, with an interest in the best possible set of rules that *we* could be using (in light of our actual cognitive constitutions and limitations) to predict and intervene in the course of affairs. We want the smallest number of rules, of course, but this cannot be gained by simple syntactic tricks of making the rules hopelessly convoluted or conjunctive. Thus, a formal statement of efficiency will require a formal notion of a semantic or cognitive unit, and the most desirable set of rules for prediction and intervention will be the maximally effective set with the lowest degree of semantic or cognitive cardinality.

Fortunately, we need not have any such formal account in hand in order to see how these pragmatic desiderata will favor the traditional construal of selective phenomena over the genic selectionist's alternative. As it turns out, we will be able to show that *endorsing genic selectionism will force us to adopt a larger corpus of equally simple and systematic rules for prediction and intervention than would be open to us under the standard construal of selection as operating at many different levels of biological organization*. This will surely constitute grounds of efficiency favoring

the standard construal on *any* reasonable account of what efficiency might amount to, ensuring in turn that the genic selectionist's causal claims are not reflected in the Pragmatic Corpus of laws and therefore *mis*describe the causal structure of selective phenomena.

4. PRAGMATIC DESIDERATA AND THE UNITS OF SELECTION

We have already seen why genic selectionism will not suffer any disadvantage in effectiveness when compared with the alternatives: any predictive claim about selection operating at a higher level of biological organization can be straightforwardly translated into a claim about selection acting upon a gene in the appropriate (allelic) environment. If genic selectionism runs afoul of the Pragmatic Theory's desiderata, then, it must be because there are grounds of efficiency favoring a corpus of rules that recognizes selective forces at multiple levels of biological organization over a corpus that adopts their genic redescriptions.

Of course, no single case of natural selection will alone reveal such grounds for preferring to construe selective forces as operating at the organismic (or higher) level of organization instead of the genic level or vice versa. In each particular case, the gene's-eye point of view generates exactly one causal rule for every rule endorsed by the more traditional construal: if we simply need to choose between 'anemia confers selective disadvantage X in (organismic) environment Eo' and 'allele a confers selective disadvantage X in (allelic) environment Ea' (which includes allelic partner a along with other factors, genetic and otherwise), then it seems that we have no grounds of efficiency for choosing one way of construing the causal situation over the other.[15] The case of human sickle-cell heterozygote superiority demands three rules for prediction and intervention whether we characterize differential selective advantages as the products of anemia, malarial susceptibility, and malarial resistance or appeal instead to the corresponding genic redescriptions of these properties.

The question ignored by this conventionalist point, however, is how these alternative accounts of the causal structure of selective processes will fit into our framework of rules for systematically predicting and intervening with respect to *other* biologically significant phenomena. The key issue is that we will surely want our corpus of rules for predicting and intervening to take advantage of the rules available to us for (1) simply predicting the *occurrence* of properties which arise at higher levels of biological organization, independent of their selective consequences (e.g., for predicting that homozygotes for the sickling allele will in fact, in the relevant environments, be anemic) and (2) predicting systematic selective

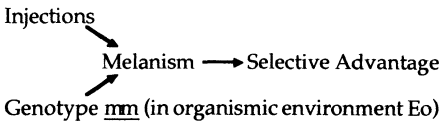

Figure 1. The traditional selectionist's causal representation of selective phenomena in the case of industrial melanism.

advantages arising from causes which are not genetic. Given our need for rules of these two kinds (and the need to integrate the rules we use into the most efficient systematic corpus possible) we can show that the rules for predicting and intervening favored by genic selectionists are indeed at a pragmatic disadvantage when compared with their traditional counterparts: any corpus of causal laws containing the genic selectionist's rules will be systematically less efficient in predicting and intervening in the course of events than the otherwise identical corpus containing the more traditional selective rules instead; thus the genic selectionist's rules cannot qualify as members of the Pragmatic Corpus. To see why this is so, let us once again revisit the case of industrial melanism and compare the number of rules for prediction and intervention required by the traditional construal of selection to the number required by the genic selectionist's alternative.

On the traditional selectionist story, a particular genotype (say *mm*) is the cause of melanism in moths, and this organismic-level property of melanism is in turn the cause (given a particular range of environmental backgrounds) of a selective advantage for those moths who possess it. Like most phenotypic properties, however, melanism can arise even in organisms which lack the corresponding genetic basis: there is a series of injections, for instance, that will render moths melanic no matter what genotype they possess, and such moths will have all the selective advantages of true (genetic) melanics. On the traditional selectionist account, these injections are simply another cause of the organismic-level property of melanism, which is in turn the cause of its usual selective advantage (again, given the relevant environmental background). Thus, the traditional selectionist account gives us, for the system we have just described, the causal representation depicted in Figure 1 (where each arrow corresponds to one rule needed for prediction and intervention): the series of Injections and the Genotype *mm* (in organismic environment Eo) are each a cause of the organismic property Melanism, which is in turn the cause (in a range of relevant environments) of a given Selective Advantage.

What structure of rules must the genic selectionist invoke to predict and intervene with respect to these same phenomena? She will, of course, have to represent the presence of a particular allele *m* (in a particular environmental background, Ea, including another copy of *m* as an allelic partner)

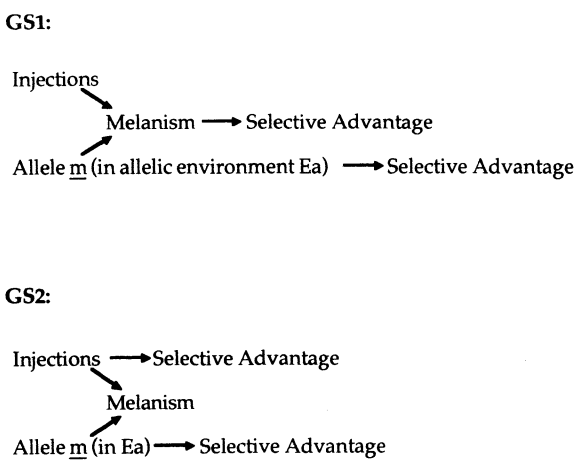

Figure 2. The two alternative schemes of causal representation available to the genic selectionist (GS).

as the immediate cause of the selective advantages enjoyed by melanic moths. But to equal the effectiveness of the traditional selectionist account, the genic selectionist's rules must also enable us to predict the occurrence of the same organismic properties (here, melanism) that can be predicted on the traditional account. To do so, the genic selectionist will need two additional rules: one predicting the organismic property Melanism from the occurrence of *m* (in Ea) and another predicting Melanism from the series of injections. This structure does not yet enable her to predict the selective advantages enjoyed by non-genetic melanics, because it predicts selective advantages by appeal to conditions obtaining at the genetic level, rather than by appealing to melanism itself. Thus, to predict the selective advantages enjoyed by non-genetic melanics, the genic selectionist will have to add *either* a rule predicting selective advantages from the organismic property of melanism itself (as in the traditional story) or a rule predicting selective advantages directly from the melanism-inducing injections, generating the two alternative genic selectionist schemes of causal representation depicted in Figure 2 as GS1 and GS2 respectively.

The crucial thing to notice, of course, is that both of these schemes for causal representation are demonstrably less efficient than the traditional causal story. That is, in order to predict and intervene with respect to all of the biologically significant events we are considering, either of these genic selectionist construals requires a *larger number* of (equally complex) rules for prediction and intervention than does the traditional construal.

The genic selectionist's natural reply is to claim that I have rigged the game by taking the genic perspective only halfway. We have already seen that a fair evaluation of the genic selectionist causal perspective must ap-

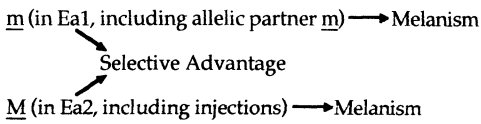

Figure 3. The causal representation generated by a *thoroughgoing* application of genic selectionism.

ply it in a thoroughgoing fashion: I have cheated by treating injections or melanism as the cause of selective advantages in non-genetic melanics, for the whole point of the genic selectionist perspective is that *all* instances of selection must be described as the products of particular alleles in particular environments. How, then, can we translate the causal facts about injections and non-genetic melanics into genespeak?

Notice that from the genic point of view, a series of injections is simply something which, like an allelic partner or a pollution-darkened tree, occurs in the extended environment of a particular allele.[16] Thus, the genic selectionist will represent the selective advantages of what I have called non-genetic melanics as an effect of their genetic composition *in the relevant* (injection-including) allelic environments. Of course, we must still insist that the occurrence of melanism at the organismic level be predicted for both genetic and non-genetic melanics.[17] Thus, the most systematic and thorough application of the genic selectionist perspective will generate the causal representation depicted in Figure 3.[18] Even when it is applied consistently, then, the genic selectionist's scheme once again requires us to adopt four rules for prediction and intervention, and is thus demonstrably less efficient than the traditional causal representation which recognizes selection operating at the level of biological organization at which the evolutionarily relevant property (in this case, melanism) arises.

We can be sure that this pragmatic superiority of the traditional, hierarchical account of the causal structure of selective phenomena will survive changes of example because it is generated by our need to predict both the occurrence of properties at higher levels of biological organization and selective advantages quite generally. Notice that the traditional selectionist representation takes predictive advantage of a certain general phenomenon: in a given environment, organisms having a given property at the organismic level (e.g., melanism) have a particular selective advantage, *no matter how they acquired* this phenotypic property. By representing this property as itself the cause of selective advantages, the traditional account reduces the number of causal connections which must be drawn: the selective advantages of any of the various causes of melanism are given automatically by the rule which specifies the selective advantages of melanism itself. Indeed, it seems quite natural to think that this is just

what our persistent intuitions about the genuine evolutionary significance of various properties at a variety of levels of biological organization really amount to: that the presence of a property at a given level confers precisely the same fitness benefits, no matter how the gene, chromosome, organism, group, deme, or species came to possess that property.[19]

Notice, however, that by insisting that selective advantages are produced exclusively by the presence of genes in particular environments, the genic selectionist precludes herself from taking predictive advantage of these patterns of invariable co-occurrence at higher levels of organization. Thus, when there are a variety of combinations of genes and environment (e.g., *m* in Ea1 (including another copy of *m*) and *M* in Ea2 (including the series of injections)) which produce precisely the same phenotypic (or even higher level) properties *and* precisely the same selective advantages, the genic selectionist will need independent rules connecting each combination to the higher-level property *and* each to the relevant selective advantage.[20] The more traditional hierarchical account of selection does not restrict itself to representing all selection as acting at the genic level, and therefore does not multiply causal connections in this way: the traditional selectionist account is free to adopt only a single rule for predicting the selective advantages in question from the relevant higher-level properties themselves. Thus, genic selectionism forces us to adopt a larger set of rules for successfully predicting and intervening with respect to biological phenomena precisely because it refuses to accept rules for predicting selective advantages from properties which arise at higher levels of biological organization.

5. CONCLUSION: CAUSATION AND THE UNITS OF SELECTION

The Pragmatic Theory of Causation thus offers us, as previous proposals have not, concrete and compelling grounds for holding genic selectionism to misrepresent the causal structure of selective phenomena and for defending our stubborn intuitions regarding the causal efficacy of properties at various levels of the traditional selective hierarchy: on the Pragmatic theory, considerations of efficiency play a *constitutive* role in picking out the causal laws of our world, so the pragmatic shortcomings of genic selectionism ensure that it misrepresents the actual causal structure of selective phenomena. But one need not embrace the Pragmatic Theory of Causation (nor any particular theory of causation) in order to understand the failing of genic selectionism to be a pragmatic one or to embrace the critique of genic selectionism offered here. Genic selectionism forces us to adopt a less efficient but no more effective set of causal laws than does the traditional

construal of selective processes, and this counts as a pragmatic failing of genic selectionism no matter what theory of causation we embrace. We are therefore free to reject the Pragmatic Theory of Causation, but nonetheless reject genic selectionism as well, simply because it forces us to adopt an unnecessarily unwieldy corpus of causal laws. The appeal of the Pragmatic Theory of Causation in this connection is not that it alone entitles us to reject genic selectionism, but rather that it alone explains why the failing of genic selectionism is more than *merely* pragmatic: that is, why genic selectionism does indeed, as we always suspected, misconstrue the very causal structure of selective phenomena.

ACKNOWLEDGEMENTS

I would like to thank an anonymous referee for this journal and a number of my colleagues, especially Philip Kitcher, for helpful comments and suggestions regarding the argument of this paper.

NOTES

[1] Dawkins's later work (1982) defends the somewhat weaker claim that selection can always be *represented* as operating on particular genes. As we will see, however, this claim supports an argument for thinking of all selection in genic terms, because only the genic representation of selection is always available.

[2] Of course, there is considerable intergroup variation among these thinkers: Gould (1980), for instance, argues simply that genic selectionism is misguided because it is organisms and not genes which are visible to selection, while Sober (1984, Chs. 7–9) offers an extended defense of the claim that selection processes can indeed operate at a variety of levels of biological organization.

[3] It is this intuition which is at stake in Sober's famous insistence (1984, passim) that we distinguish the selection *of* a particular set of entities (genes, organisms, groups, whatever) from the selection *for* the properties which are of genuine evolutionary significance (that is to say, the properties *in virtue of which* the entities enjoy their selective advantages), at whatever level of biological organization those properties arise.

[4] Wimsatt (1980, 1981) attempts to resolve the units of selection issue by appeal to context-independent variance in fitness among entities at one level which does not appear as context-independent variance in fitness at any lower level of organization, and Lloyd (1988, 1989) refines this account into one based on additive components of variance and decomposition between levels of organization. Sober (1984, Section 7.8) offers an extended attack on the attempt to understand units of selection in terms of the analysis of variance. As far as I can see, Lloyd's (1989) reply to Sober's critique establishes simply that a suitably idealized application of ANOVA (the analysis of variance) is a better guide to causal structure than Sober allows, but it seems to me that Sober is right to insist that it is the causal structure of selection processes themselves which is at issue. Thus, I do not take Lloyd's insistence on the importance of the ANOVA criterion (perhaps even its *constitutive* role when its application is suitably idealized, depending on one's view of causation) to be

inconsistent with the proposal I will make below (or indeed with any attempt to resolve the controversy by appeal to the causal structure of selection processes), but I do not see how that insistence is itself supposed to help resolve the controversy.

[5] It is not entirely clear to me how this can work, as the choice of a trait to which we are relativizing in this way would seem to determine the choice of unit of selection for that trait, because most kinds of traits can only be possessed by entities at a given level of biological organization: genes cannot be anemic, organisms do not have proportions of altruistic and selfish members, and populations are not segregation-distorting. Sober and Wilson's (1998) formulation of the procedure for identifying units of selection *may* avoid this problem, but nothing I say here depends upon this worry about their view in any case.

[6] Indeed, in easily identifiable cases of meiotic drive (see, for example, Lewontin and Dunn's famous (1960) study of the *t* allele in *Mus musculus*) the distorting allele is typically harmful to the organism (at least in homozygous form): otherwise (barring complications), the driving gene becomes fixed in any population in which it occurs.

[7] Sober (1984, Ch. 8) offers a particular account of how we can pick out 'relevant' background contexts, but nothing I say turns on the adequacy of this account.

[8] Thus, Sober and Wilson's (1994, 552) argument that different hypotheses about the units of selection typically make different empirical predictions is misguided: while there are indeed differences in the empirical predictions made by hypotheses of selection acting at different levels in the traditional selective hierarchy (and differences between the corresponding genic reconstructions, appealing to distinct genic environments, of each of these hypotheses), there are *no* differences in empirical predictions made by any *given* hypothesis of selection operating at a particular level of the traditional hierarchy and a genic reconstruction of *that very hypothesis* (that is, of selection operating on genes in genic environments that include those features that the given traditional selective hypothesis holds to be relevant).

[9] Indeed, Sober and Wilson themselves (1994) provide an insightful summary of the previous attempts to make this intuition concrete, along with powerful criticisms of each; as we've seen above, however, their own alternative does no better in undermining or replying to genic selectionism.

[10] Sober and Wilson's most recent reply to genic selectionism (1998, Ch. 2) softens their earlier charges, suggesting that "selfish gene theory" is simply an alternative 'perspective' on the same selective processes that can also be described in other ways. But this concession seems to give the *conventionalist* genic selectionist (like Sterelny, Kitcher or Waters) all that she wants or needs to argue that representing selection exclusively at the genic level is the most convenient and systematic option: if genic and "multilevel" selection are merely different ways of looking at the same process, Sober and Wilson cannot argue that the genic selectionist is misdescribing the causal structure of selective phenomena. Their further suggestion (1998, 91–92) that the 'selfish gene perspective' requires the "tag-along concept of vehicles [actually, genic environments], which is not nearly as well developed as multilevel selection theory" is quite a weak reply to the conventionalist's case for appealing exclusively to the genic perspective. I hope to show that the underlying conventionalism Sober and Wilson are accepting here is misguided in any case: that is, I hope to show that to characterize all selection as operating at the genic level is not merely inconvenient, but is instead to get the causal story *wrong*.

[11] It might be thought that the notion of intervention introduces an inherently modal or counterfactual element into the account. If so, a maximally effective corpus of rules should be understood to be one whose rules correctly describe all (or as many as possible) of

the successions of events in the actual course of events and in all of the counterfactually possible courses of events that either are or were actually open to us at some time. That is, the Pragmatic Theory can accommodate the modal or counterfactual aspect of intervention without becoming a disguised counterfactual analysis of the causal relation itself.

[12] At the extreme, we could always achieve maximal effectiveness by adopting an enormous corpus of event-specific rules, each of which predicted exactly one actual event by appeal to some unique characteristic(s) (e.g., specific spatiotemporal position) of the event immediately preceding it, or (incorporating counterfactual intervention) exactly one *possible* event by appeal to some likewise unique characteristic(s) of the event which immediately precedes it in a given counterfactual course of events.

[13] Accounts of causation, causal law or natural law which respect Hume's deflationary ambitions have been offered by Braithwaite (1927), Ramsey (1928, 1929), Strawson (1952), Goodman (1954), Ayer (1956), Popper (1959, 1968), Nagel (1961) and Skyrms (1980). These and related accounts are criticized in Kneale (1950, 1961), Molnar (1969), Dretske (1977), Armstrong (1983), Tooley (1987), and elsewhere.

[14] That is, while the efficiency demand surely renders causal relations relative to the cognitive constitutions of *human beings*, I am skeptical that criteria of efficiency in selecting rules for prediction and intervention will have to be so variable as to demand a different Pragmatic Corpus for members of different cultures, or for adherents of competing 'conceptual schemes'. The account could, of course, accommodate this outcome as well, but its attractions would be diminished.

[15] One might think that it is the *complexity* of the rules which favors the efficiency of the traditional selectionist story. This may be, but without a convincing formal account of semantic or cognitive complexity, I am not free to assume so, and these two sorts of rules (along with all the others we will consider here) are at least *roughly* equal in the magnitude of their complexity. Accordingly, I will try to show that genic selectionism has much bigger problems in the efficiency department, for reasons having nothing to do with the complexity of the rules it requires.

[16] Failure to appreciate this point is what undermines Brandon's (1984) argument against genic selectionism by appeal to screening off relations. See Sterelny and Kitcher, pp. 351–354.

[17] As this demand makes clear, the genic selectionist could avoid the trouble by refusing to talk about phenotypic properties or their causes in any context, selective or otherwise. But predicting phenotypes from genotypes in environments is surely one of the most fundamental and important things that we want biological science to be able to do. Perhaps giving up all talk of phenotypes and their causes is what Waters means when he grants (p. 562) that Williams's perspective may be "strange", but if so, the genic selectionist perspective is much stranger (and much less useful) than Williams, Waters, Sterelny, Kitcher or even Dawkins is letting on or is willing to accept.

[18] Here M is the alternative allelic form to m, and the one that nonmelanics share; that is, suppose that genotypes MM and Mm generate ordinary, nonmelanic peppered moths in typical (non-injection-including) allelic environments. Incidentally, nothing turns on the assumption that M is dominant over m: if we reverse the dominance and adjust the environments to include allelic partner M in Ea2 and drop the requirement that allelic partner m be in Ea1, an identical problem results for the genic selectionist.

[19] It is worth noting that this analysis bears out Sober's (1984) insistence that it is the properties which are selected *for* (e.g., selection for being tall vs. selection for being a member of a group with many tall members) which determines the level of biological organiza-

tion at which selection is genuinely acting. Genic selectionist redescription is misguided because tall organisms experience whatever selective advantages they do, no matter how they came to be tall and organisms in groups with many tall members experience whatever selective advantages they do, no matter how the group came to possess many tall members, *and this makes a difference to the efficiency of the schemes of causal representation which are available to us.*

[20] Of course, the genic selectionist could reduce the number of rules she requires by adopting rules predicting a single selective advantage from either genetic condition and a further rule predicting *melanism* in turn from this selective advantage, but this would simply be an empirical mistake: it is not the case that any moth enjoying some particular degree of selective advantage over its competitors becomes melanic.

REFERENCES

Armstrong, D. M.: 1983, *What Is a Law of Nature?*, Cambridge University Press, Cambridge.

Ayer, A. J.: 1956, 'What Is a Law of Nature?', reprinted in A. Ayer (1963), *The Concept of a Person and Other Essays*, St. Martin's Press, New York, pp. 209–234.

Braithwaite, R. B.: 1927, 'The Idea of Necessary Connexion', *Mind* **36**, 467–477.

Brandon, R. N.: 1984, 'The Levels of Selection', in R. N. Brandon and R. M. Burian (eds), *Genes, Organisms, Populations: Controversies over the Units of Selection*, MIT Press, Cambridge, pp. 133–141.

Brandon, R. N.: 1990, *Adaptation and Environment*, Princeton University Press, Princeton.

Cartwright, N.: 1989, *Nature's Capacities and Their Measurement*, Clarendon Press, Oxford.

Dawkins, R.: 1976, *The Selfish Gene*, Oxford University Press, Oxford.

Dawkins, R.: 1982, *The Extended Phenotype: The Gene as the Unit of Selection*, Freeman, Oxford.

Dawkins, R.: 1994, 'Burying the Vehicle' (Open peer commentary on Wilson and Sober's (1994a) 'Reintroducing Group Selection to the Human Behavioral Sciences'), *Behavioral and Brain Sciences* **17**, 616–617.

Dretske, F. I.: 1977, 'Laws of Nature', *Philosophy of Science* **44**, 248–268.

Godfrey-Smith, P. and R. Lewontin: 1993, 'The Dimensions of Selection', *Philosophy of Science* **60** 373–395.

Goodman, N.: 1954, *Fact, Fiction and Forecast*, The Athlone Press, Atlantic Highlands, NJ.

Gould, S. J.: 1980, 'Caring Groups and Selfish Genes', in *The Panda's Thumb: More Reflections in Natural History*, Norton, New York, pp. 85–92.

Kitcher, P.: 1989, 'Explanatory Unification and the Causal Structure of the World', in P. Kitcher and W. Salmon (eds), *Scientific Explanation*, University of Minnesota Press, Minneapolis, MN, pp. 410–505.

Kneale, W.: 1950, 'Natural Laws and Contrary-to-Fact Conditionals', reprinted in T. L. Beauchamp, ed.: 1974, *Philosophical Problems of Causation*, Dickenson, Belmont, CA.

Kneale, W.: 1961, 'Universality and Necessity', reprinted in T. L. Beauchamp, ed.: 1974, *Philosophical Problems of Causation*, Dickenson, Belmont, CA.

Lewis, D.: 1973a, 'Causation', *Journal of Philosophy* **70**, 556–567.

Lewis, D.: 1973b, *Counterfactuals*, Harvard University Press, Cambridge, MA.

Lewontin, R. and L. Dunn: 1960, 'The Evolutionary Dynamics of a Polymorphism in the House Mouse', *Genetics* **45**, 705–722.

Lloyd, E. A.: 1988, *The Structure and Confirmation of Evolutionary Theory*, Greenwood Press, New York.
Lloyd, E. A.: 1989, 'A Structural Approach to Defining Units of Selection', *Philosophy of Science* **56**, 395–418.
Mayr, E.: 1963, *Animal Species and Evolution*, Harvard University Press, Cambridge, MA.
Molnar, G.: 1969, 'Kneale's Argument Revisited', *Philosophical Review* **78**, 79–89.
Nagel, E.: 1961, *The Structure of Science*, Harcourt, Brace and World, New York.
Popper, K. R.: 1959, *The Logic of Scientific Discovery*, Hutchinson, London.
Popper, K. R.: 1967, 'A Revised Definition of Natural Necessity', *British Journal for the Philosophy of Science* **18**, 316–324.
Ramsey, F. P.: 1928, 'Universals of Law and Fact', reprinted in Ramsey: 1978, in D. H. Mellor (ed.), *Foundations*, Routledge and Kegan Paul, London, pp. 128–132.
Ramsey, F. P.: 1929, 'General Propositions and Causality', reprinted in Ramsey: 1978, in D. H. Mellor (ed.), *Foundations*, Routledge and Kegan Paul, London, pp. 133–151.
Skyrms, B.: 1980, *Causal Necessity*, Yale University Press, New Haven, CT.
Sober, E.: 1984, *The Nature of Selection: Evolutionary Theory in Philosophical Focus*, University of Chicago Press, Chicago.
Sober, E. and R. C. Lewontin: 1982, 'Artifact, Cause, and Genic Selection', *Philosophy of Science* **49**, 157–180.
Sober, E. and D. S. Wilson: 1994, 'A Critical Review of Philosophical Work on the Units of Selection Problem', *Philosophy of Science* **61**, 534–555.
Sober, E. and D. S. Wilson: 1998, *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Harvard University Press, Cambridge.
Stanford, P. K.: forthcoming, 'The Pragmatic Theory of Causation'.
Sterelny, K. and P. Kitcher: 1988, 'The Return of the Gene', *Journal of Philosophy* **85**, 339–361.
Strawson, P. F.: 1952, *An Introduction to Logical Theory*, Methuen, London.
Tooley, M.: 1987, *Causation: A Realist Approach*, Oxford University Press, New York.
Waters, C. K.: 1991, 'Tempered Realism about the Force of Selection', *Philosophy of Science* **58**, 553–573.
Williams, G. C.: 1966, *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*, Princeton University Press, Princeton.
Wilson, D. S. and E. Sober: 1994a, 'Reintroducing Group Selection to the Human Behavioral Sciences', *Behavioral and Brain Sciences* **17**, 585–608.
Wilson, D. S. and E. Sober: 1994b, Authors' response to open peer commentary on 'Reintroducing Group Selection to the Human Behavioral Sciences', *Behavioral and Brain Sciences* **17**, 639–647.
Wimsatt, W.: 1980, 'Reductionistic Research Strategies and Their Biases in the Units of Selection Controversy', in T. Nickles (ed.), *Scientific Discovery*, Vol. 2. Reidel, Dordrecht, pp. 213–259.
Wimsatt, W.: 1981, 'The Units of Selection and the Structure of the Multi-level Genome', *PSA 1980*, Vol. 2, Philosophy of Science Association, East Lansing, pp. 122–183.

University of California
Irvine
U.S.A.
E-mail: stanford@uci.edu