




1-1-2001

Contextual Unanimity and the Units of Selection Problem

Stuart M. Glennan

Butler University, sglennan@butler.edu

Follow this and additional works at: https://digitalcommons.butler.edu/facsch_papers

 Part of the [Philosophy of Mind Commons](#), [Philosophy of Science Commons](#), and the [Theory and Philosophy Commons](#)

Recommended Citation

Glennan, Stuart M., "Contextual Unanimity and the Units of Selection Problem" *Philosophy of Science* / (2001): 118-137.

Available at https://digitalcommons.butler.edu/facsch_papers/622

This Article is brought to you for free and open access by the College of Liberal Arts & Sciences at Digital Commons @ Butler University. It has been accepted for inclusion in Scholarship and Professional Work - LAS by an authorized administrator of Digital Commons @ Butler University. For more information, please contact digitalscholarship@butler.edu.

Contextual Unanimity and the Units of Selection Problem*

Stuart Glennan†‡

Department of Philosophy and Religion
Butler University

Sober and Lewontin's critique of genic selectionism is based upon the principle that a unit of selection should make a context-independent contribution to fitness. Critics have effectively shown that this principle is flawed. In this paper I show that the context independence principle is an instance of a more general principle for characterizing causes, called the contextual unanimity principle. I argue that this latter principle, while widely accepted, is erroneous. What is needed is to replace the approach to causality characterized by the contextual unanimity criterion with an approach based on the concept of causal mechanism. After sketching such an approach, I show how it can be used to shed light on the units of selection problem.

1. Introduction. Elliot Sober and Richard Lewontin's well-known argument against genic selection is based upon the claim that genic selection misrepresents the *causes* of evolution:

Although models of evolutionary processes conforming to [Richard Dawkins' and George C. Williams'] view of genic selection may permit computation, they often misrepresent the causes of evolution. The reason is that genic selection coefficients are *artifacts*, not causes of evolution. (Sober and Lewontin 1982, 108)

*Received May 2001; revised August 2001.

†Send requests for reprints to the author, Department of Philosophy and Religion, Butler University, 4600 Sunset Ave., Indianapolis, IN 46208; e-mail: sglennan@butler.edu.

‡This paper is based on a talk presented at a colloquium at the University of Maryland's Center for the History and Philosophy of Science. I am grateful to the faculty and students at Maryland, and especially to Lindley Darden, for feedback. Thanks also to Ken Waters and to two referees for comments on an earlier draft of this paper and to my colleagues Tom and Becky Dolan for their insights on malaria, moths, and other things biological.

Philosophy of Science, 69 (March 2002) pp. 118–137. 0031-8248/2002/6901-0006\$10.00
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

But how does one distinguish causes from artifacts? The principle that Sober and Lewontin use is that causes should make a context-independent contribution to bringing about their effects. The level of selection is the level at which the causes operate, and this level can be recognized by a context independence principle:

- (1) If the fitness of X is context-sensitive, then there is not selection for X; rather there is selection at a level of organization higher than X. (Sober and Lewontin 1982, 579)

Sober and Lewontin's critics, notably Sterelny and Kitcher (1988) and Waters (1991), defend genic selectionism by criticizing this principle. But while the critics do demonstrate flaws in principle (1), they do not thereby save genic selectionism.

Principle (1) is in fact an application of a general principle about causation which John Dupré has called the contextual unanimity principle. This principle is (roughly) that a cause should always increase the probability of its effect. I shall argue that this principle, while widely accepted, is false, and that it is because it is false that Sober and Lewontin's critique of genic selectionism fails. To substantiate Sober and Lewontin's claim that genic selection coefficients are artifacts rather than causes, what is needed is an alternative approach to causation that does not presuppose contextual unanimity. The main goal of this paper will be to argue that this need can be filled by an approach to causation based upon mechanisms (cf. Glennan 1996).

Section 2 of this paper recapitulates the argument between Sober and Lewontin and their critics, showing how their dispute centers around the contextual unanimity principle. In Section 3 I discuss the contextual unanimity principle more generally, showing both what is wrong with it and why, given how obviously false it is, philosophers have continued to espouse it. In Section 4, I describe some important aspects of the mechanistic approach to causation. Finally, in Section 5, I return to the units of selection problem, showing how a mechanistic approach allows one to make the case for higher-level selection in a way that is immune to the criticisms of the contextual unanimity principle.

Prior to beginning my discussion of the units of selection controversy, I should clarify which units of selection problem I will address. There are at least two. One problem concerns the nature and efficacy of group selection or selection for entities at a sub-organismic level of organization, as opposed to selection acting on individual organisms. Brandon (1982) calls this problem the 'level of selection problem'. It is centrally concerned with how one draws the boundary between differentially selected entities and their environments. Brandon reserves the term 'unit of selection problem' for the question of whether all selection can be understood in terms

of differential fitness of individual alleles, or whether fitnesses should be attributed to larger segments of the genome. To use David Hull's terminology, the first question concerns interactors, while the second question concerns replicators. Although their 1982 paper does not mention this distinction, Sober and Lewontin's argument clearly concerns the *units* of selection problem rather than the *levels* of selection problem. In the context of this debate, all parties acknowledge that selection for genotypes occurs indirectly via selection for phenotypes. The question at issue concerns how large or small a segment of the genome can be selected for in this manner.

Waters' critique of Sober and Lewontin makes it clear that Sober and Lewontin's position assumes an answer to the levels of selection problem—that is, for a given case it is not purely conventional how one draws the boundary between a selected entity (the interactor) and its environment. In particular, Sober and Lewontin assume for the case they discuss (sickle cell anemia) that the level of selection is that of the organism and argue on that basis that genes do not make context-independent contributions to an organism's fitness.

Waters argues that the decision about where to draw the boundary between the selective domain and the environment is essentially pragmatic and conventional, and it is on that basis that he argues against Sober and Lewontin's position. My concern will be to show that *even if* Sober and Lewontin are right that this boundary is non-conventional, they still can't use contextual unanimity to individuate a unit of selection relative to this boundary. In the last section of this paper I shall defend Sober and Lewontin's principle about the non-conventionality of the boundary between selective domain and environment. Ultimately any argument against the universal applicability of genic selection models requires a successful defense of this principle. But even those who do not find my defense of this principle persuasive will perhaps grant that the mechanistic approach to causation shows how one can, relative to a choice of environmental boundary, identify the unit of selection.

2. The Contextual Unanimity Approach to the Units of Selection. Sober and Lewontin's critique of genic selection is based upon what Ernst Mayr has called the principle of genetic relativity—that is, the principle that the contribution of a gene to an organism's fitness depends upon the organism's other genes. Phenomena such as dominance, heterozygote superiority, and epistasis all illustrate this principle. The principle of genetic relativity implies that genic fitnesses are context dependent, and hence, by principle (1), that there cannot be selection at the genic level. The specific example Sober and Lewontin use to demonstrate their claim is the phenomenon of heterozygote superiority that occurs at the locus implicated

in sickle cell anemia. To recount briefly the standard textbook account, sickle cell anemia, which is a disease that virtually eliminates the possibility of producing offspring, is caused by being homozygous for a single mutant allele a . While it is clearly deleterious to be homozygous for the sickle cell allele, possession of a single sickle cell allele has benefits for the organism. The heterozygote suffers few effects of anemia and gains increased resistance to the malaria parasite. Thus, at least in environments in which malaria is a threat to survival and reproduction, the heterozygote Aa is superior to homozygotes for either the sickle cell (aa) or wild-type allele (AA). Sober and Lewontin conclude that one cannot say, in isolation, that possessing the sickle cell allele either enhances or diminishes fitness, so selection cannot occur for (or against) the individual allele. Allelic selection coefficients for a population can be obtained by averaging over the coefficients for the genotypes (AA , Aa , aa), but it is only the genotypic coefficients that represent the forces of selection.

The idea that context sensitivity of effect can be used as a criterion to rule out an entity as a unit of selection is not original to Sober and Lewontin (1982). The principle was explicitly introduced to the philosophical literature by Wimsatt (1980, 1981), who in turn was inspired by earlier work of Mayr (1963) and Lewontin (1974). The novel contribution of Sober and Lewontin's paper was to propose that the context sensitivity criterion was in fact a criterion for distinguishing causes from artifacts.

The difficulty with the context sensitivity criterion that has been raised both by Sterelny and Kitcher and by Waters is that even the fitnesses of the larger segments of the genome, in this case the allele pairs occurring at a locus, have only a context-sensitive effect on fitness. The fitness of any genotype varies with environment. In the sickle cell case, for instance, the heterozygote is only more fit than the wild-type homozygote in environments where malaria is a threat.

Sober and Lewontin of course recognize that the fitness of a genotype varies with environment, and their response to this point would be to say that a different set of selective forces will be acting on populations where malaria is a threat than on populations where it is not. The differing selection processes will produce the differing evolutionary trajectories of these isolated populations. However, as Sterelny and Kitcher argue, this kind of variation in fitness may affect organisms that we would consider to be in the same population. For instance, though an increase in melanic coloration of moths in England has been an adaptive response to environmental changes caused by industrialization, it certainly must be the case that some melanic moths will spend time in parts of their environment where melanic coloration actually will increase their visibility to predators.¹ As these moths are not geographically isolated from moths that

1. Recent critics of Kettlewell's experiments (cf. Majerus 1998) have raised serious

happen to spend more time on trees in which melanic coloration provides protection (and as the two sets of moths may interbreed) it seems problematic to treat those moths as part of a different selection process.

Another way to try to defend Sober and Lewontin is to construe the context independence principle as being relativized to a fixed environment. Of course, the fitness of a genotype depends upon the environment of organisms possessing that genotype, but (the argument goes) the fitness of individual genes will depend upon more than the environment. For instance, the fitness value of a sickle cell allele will depend not only upon the environmental risk of malaria, but, most importantly, on the other allele that is found at that locus. But, as both Waters, and Sterelny and Kitcher, point out, the genic selectionist will characterize the genetic context as *part* of the allele's environment. Once Sober and Lewontin grant (as they must) that the fitness of an entity varies with its environment, then they must offer some principled reason to draw the boundary between the selected entities and their environment at the boundary of the organism. Sterelny and Kitcher and Waters, however, see no such reason, and conclude from this that there is no privileged unit of selection.

I shall argue in Section 5 of this paper for the claim that in the sickle cell case there are principled reasons to draw the selected entity/environment boundary at the level of the organism. But even if my claim is correct, it will not immunize Sober and Lewontin's position from concerns about context sensitivity. Suppose, for the sake of argument, that a segment of the genome is a unit of selection so long as, relative to a specific ecological environment, that segment makes a context-independent contribution to an organism's fitness. Even if we grant this, we cannot say that the diploid genotype is the unit of selection in the sickle cell case. Sober and Lewontin argue as if the fitness of the alleles A and a depend upon genetic context, while the fitness of the genotypes AA , aa , and Aa do not depend upon genetic context. But this is clearly not the case, for the fitness of the organism will depend (non-additively) upon genes at other loci. To take the most obvious example, if an organism carries a lethal allele at another locus, the alleles at the sickle cell locus will not affect the organism's fitness in any way. Mayr's principle of genetic relativity does not stop at the individual locus. The consequence of this is that, if one requires context independence for an entity to be a unit of selection, the unit of selection will inevitably be the entire genome.

doubts about the standard explanation of the selective value of melanism in the pepper moth (according to which melanic moths were selected for because industrialization darkened tree bark and destroyed lichens on trees whose coloration provided camouflage for the wild type). However, these doubts do not undermine Sterelny and Kitcher's conceptual point.

This is an unwelcome conclusion to anyone who thinks the units of selection question is a substantive one. Such a person would prefer a less degenerative conclusion in which selection operates simultaneously on units of different sizes. In some cases an allele at a single locus might be a unit of selection; in cases of dominance or heterosis, selection must be seen as operating on the pair of alleles at a locus; in cases of epistasis, the unit of selection would be a part of the genome consisting of multiple loci. Such a position can, I think, be sustained, but only if one gives up the requirement of context independence.

3. Contextual Unanimity as a Causal Concept. The context independence principle used by Sober and Lewontin is, as I suggested at the outset, an application to selection theory of a general principle about causation that Dupré (1984) has dubbed the contextual unanimity principle. Versions of this principle have been endorsed at times by Cartwright, Otte, Sober, Eells, and Humphreys. Dupré and Cartwright (1988) offer the following version, which they call "CL" for "causal law"

CL *Ps* probabilistically cause *Qs* if in every population homogeneous with respect to all other causally relevant factors, the presence of *Ps* raises the probability of the occurrence of *Qs*.

The variables *P* and *Q* are intended to be general terms of some kind—properties, event types or state types. For instance, if *P* and *Q* were the properties of smoking and having cancer, CL would provide a specification of the conditions under which smoking can be said to cause cancer. Broadly speaking, the point of CL is to distinguish correlation from causation. To say that smoking causes cancer is to say more than that smoking raises the probability of cancer. In particular, the import of Simpson's paradox is that there are situations in which a positive correlation between two factors in a population can be reversed when that population is partitioned by a causally relevant variable.² By stipulating that the cause must raise the probability of its effect in all possible partitions of causally relevant variables, the possibility of mistaking a causal for a non-causal correlation is avoided.

There are a variety of technical difficulties that can be raised with various versions of the contextual unanimity principle (Ray 1992, Glennan 1997), but the most important is this: There will be few if any candidates

2. The classic example of this concerns a *prima facie* case of sex discrimination in admissions at Berkeley's graduate school. The overall admission rate for women was significantly less than that for men, but when admissions data were partitioned by department, it was discovered that women applied disproportionately to departments with lower admissions rates, and that department by department, the admission rate for women was comparable to that of men. See Cartwright (1979).

for *P* and *Q* in which the contextual unanimity condition can be met. Thus, if CL provides the correct analysis of causal generalizations of the form '*Ps* cause *Qs*', almost all causal generalizations will be false. It will be false that smoking probabilistically causes cancer; it will be false that obesity probabilistically causes heart problems; it will be false that drug use by parents probabilistically causes drug use in their children. The reason such generalizations will be false is that an event of type *P* will only raise the probability of an event of type *Q* in the presence of suitable background conditions and in the absence of confounding factors. For instance, a condemned criminal who smokes his first cigarette moments before his execution will not raise his probability of getting cancer. Indeed, smoking might even decrease one's risk of cancer in unusual circumstances. Suppose an avid smoker heads out to buy cigarettes and suppose that while he is gone, an environmental accident occurs causing the release of powerful carcinogens near his home. If the carcinogenic effects of the accident outweigh those of the cigarettes, smoking would actually have been a counteracting cause of cancer.

Now it is possible to accept this consequence of contextual unanimity—to say that generalizations like the one that smoking causes cancer are false. Humphreys (1989) adopts just this view. But to adopt this approach would require a complete revision of our ordinary practices of making causal generalizations. Ordinarily, we would allow, for instance, that we have ample evidence that smoking causes cancer, even though it is not the case that every smoker has his/her probability of contracting cancer raised by smoking. Alternatively, we can reject contextual unanimity outright, as is suggested by Dupré (1984, 1993) and Dupré and Cartwright (1988). They argue that we ought to treat generalizations as representing "causal upshots," which is to say that we should treat them as saying that, in a particular population, *Ps* have on average increased the probability of *Qs*. While I find this the only plausible way to understand such claims, this interpretation implies that such claims are not really type-level claims at all, but are rather singular claims about particular populations.³

Contextual unanimity principles are meant to provide a definition (or at least a test condition) for causal laws, where causal laws are understood as relations between properties, or event types. I have showed that contextual unanimity is a very hard condition to meet. In light of this discussion, it should be clear that the objections raised to Sober and Lewontin's

3. Notice that there is a *prima facie* inconsistency between Cartwright's endorsement of contextual unanimity conditions in her 1979 and 1989 and her rejection of them in Dupré and Cartwright 1988. She attempts to resolve this inconsistency by suggesting that contextual unanimity applies only to what she calls "capacity claims" and not to what she calls "causal laws." See my 1997 for an argument that her strategy fails.

proposal by Waters, Sterelny, and Kitcher all amount to pointing out failures of contextual unanimity. Their examples of frequency-dependent causation, the differing microenvironments of melanic moths, etc., are all meant to show that there exist sets of causally relevant background conditions which reverse the supposedly positive effect of an allele, a diploid genotype, or any other larger segment of the genome. It is just wrong, applying the contextual unanimity principle, to say that having a particular genotype causes an increase in fitness.

While I believe that Waters, Sterelny, and Kitcher have indeed succeeded in showing that contextual unanimity, and with it, principle (1), is never met, I do not think this criticism undermines the larger point of Sober and Lewontin's position—which is that some fitness values (like those of single alleles in the sickle cell case) represent artifacts, while others (like those of the diploid genotypes) represent causes. The problem with Sober and Lewontin's approach is that they try to distinguish causes from artifacts by means of *type-level* claims generated by CL (or similar principles), but no such type-level claims are available. However, when one claims, for instance, that a diploid genotype fitness represents a cause while an allele fitness represents an artifact, one is really making a remark about the causes of evolution in a *particular* evolutionary process—that is, one is making a token-level (or singular) causal claim. Even if one accepts a principle like CL as being the correct sort of principle for marking type-level causal claims, it is simply inapplicable to the problem of determining the correctness of singular causal claims.

Many philosophers will doubt my suggestion that the causal claims made in evolutionary biology are singular claims, for at first glance this position seems at variance with the population-oriented and statistical character of evolutionary theory. Fitnesses it is said should be attributed to genotypes, not to genotokens. Evolutionary theory, as Sterelny and Kitcher so elegantly put it, abstracts from “the thousand natural shocks that organisms and natural populations are heir to” to divine the “central tendencies” (Sterelny and Kitcher 1988, 593). True enough—but central tendency claims are still singular claims. They are claims about the average contribution of some property to an outcome *in a particular population*. In short, they are a species of causal upshot claim in the sense of Dupré and Cartwright. A population may consist of hundreds, thousands, or millions of organisms, but this population is itself a complex individual of which these organisms are its many parts. Change the population, and the average contributions will change.⁴

4. A referee suggests that, even if in reality populations are complex individuals with non-uniform parts, treating them as classes is often a useful modeling technique. I would agree that such techniques are *useful*, but this is not the same as providing a realistic account of the actual causal processes involved.

But, if the causal claims of evolutionary theory are singular, then we must look to another principle besides contextual unanimity to distinguish causes from artifacts. For this, we need a theory elaborating the truth conditions for singular causal claims. The best known theory is that of David Lewis (1973). Lewis' theory grounds the truth of causal claims in the truth of certain counterfactual conditionals. Although there are clearly things to recommend this theory, there are serious difficulties with using it to determine the truth of particular claims in evolutionary theory. The truth conditions for counterfactuals stipulated by Lewis (which have to do with the truth of material conditionals in similar possible worlds) are not presented in a way that makes clear what empirical tests should be applied to test particular counterfactual claims. I have elsewhere developed an alternative approach to singular causation according to which an event is connected to another event if those events are appropriately connected by an intervening mechanism (Glennan 1996). The theory that I have developed does not ultimately eliminate the dependency of singular causal claims on the truth of counterfactuals, for the definition of a mechanism I offer itself appeals to counterfactual principles. However, as a practical matter, this theory has the great advantage that it offers an epistemic principle for discerning the truth of causal claims which is consonant with the practices of science: The best way to establish that a certain event was a cause of another event is to show that there exists a mechanism connecting the first event to the second.

In brief, the position I shall defend is this: Sober and Lewontin's critique of genic selectionism, and with it their claim that genic fitnesses typically represent artifacts rather than causes of evolution, can be sustained by showing that the mechanisms responsible for enhancing the fitness of individual organisms generally depend upon larger segments of the genome than the single allele.

4. The Idea of Mechanism. Before considering a mechanistic approach to the units of selection controversy, it is well to consider the nature of mechanisms in general.⁵ In earlier work (Glennan 1996) I have argued for an analysis of mechanisms according to which a mechanism is a complex system whose parts interact to produce a behavior. I use the term "behavior" loosely to refer to "what the mechanism does," that is, to the set of dispositions which the system has in virtue of the arrangement of and

5. There has been a spate of recent work on the concept of mechanism, including Machamer, Darden, and Craver 2000, Thagard 1999, and a symposium at the 2000 PSA meetings. While there are some important variations in these accounts, I believe that the basic view of the nature of mechanisms is similar to the one I propose in Glennan 1996.

connections between its parts. It is crucial to notice that mechanisms, as we describe them, always do or are for something—the behavior. A description of the behavior of a mechanism constitutes the explanans, while a description of the parts of the mechanism and their interactions constitutes the explanandum.⁶

Products of human engineering provide excellent examples of mechanisms. Let us consider an example whose principles are well understood by most everyone—the modern toilet. A toilet can be analyzed as a mechanism for removing water and waste from the bowl and replacing it with fresh water. To characterize this mechanism (or as we might say, construct a model of it) we begin by giving a description of the mechanism's behavior. This behavior is actually dispositional, for what a toilet bowl does is to replace the water and waste within the bowl with fresh water, *if* the toilet's lever is depressed. We complete the model by describing the various parts of the mechanism and the way in which these parts interact to produce the mechanism's behavior. In this case the parts of the mechanism include the bowl, the lever, the tank, the flapper valve connecting the tank to the bowl, the chain connecting the lever to the float valve, the intake valve, and the float. The mechanism works about like this: Depressing the lever pulls the chain, which in turn momentarily opens the flapper valve. When the flapper valve opens, the water in the tank rushes into the bowl, forcing the waste and water then in the bowl down the drain and replacing it with the fresh tank water. At the same time, the drop in water level in the tank causes the float to fall, opening the intake valve and causing the tank to refill—a process which continues until the water level raises the float, closing the valve and returning the toilet to its original state.⁷

I have deliberately chosen a simple system—an artifact of human engineering in which both the behavior of the artifact and the mechanism that produces this behavior are well understood. It should be evident, however, that not all mechanisms are artifacts. In the first place, many interesting mechanisms (especially biological ones) are products of selec-

6. The analysis of mechanisms I propose resembles in many respects what Cummins (1975) calls functional analysis, and the mode of explanation what Kaufmann (1970) calls "articulation of parts explanation." It is tempting to use 'function' in place of the term 'behavior'. I avoid this because (1) the function of an entity is often construed as its role within a larger system, and (2) it is also sometimes assumed that entities can have functions only if they have this function as a result of design. For more on the explanatory role of mechanisms, see Craver 2001 and Glennan forthcoming.

7. Some might object that this rough description of how a toilet works hardly constitutes a model; however, one could construct a model of a sort more familiar in the literature, by identifying a set of state variables to characterize the state of the various parts of the mechanism together with laws of succession to represent the various interactions between these parts. But to do this would be silly.

tion processes and are not the product of intelligent designers. Moreover, there are in fact many systems, e.g., our solar system, that are mechanisms in the sense of the definition above, even though they are not products of design processes of any kind.

Mechanisms in the sense I have described them here are *things*—that is, systems that are collections of parts. The term ‘mechanism’ is also, however, used in another, related sense. Consider, for instance the following question: What is the mechanism (or what are the mechanisms) which led to the outbreak of the First World War? The purport of this question is not to ask for a description of a thing, but rather a description of a *causal process*—the process that led to the outbreak of the First World War. We need to consider briefly the relation between these two concepts of mechanism, for it has implications for the distinction between singular and general causal claims that underlies my criticism of contextual unanimity.

A causal process is generally understood as a singular sequence of events, or, as it is sometimes called, a causal chain.⁸ When we ask about the causal process leading up to the First World War, we are inquiring as to what sequence of events led to the war. It is in virtue of the existence of such sequences—the causal nexus—that a particular event (like the assassination of Archduke Ferdinand) can be said to cause another particular event. It would be an abuse of usage to treat the “First World War starting mechanism” as a system, or a thing. The reason arises from the very singularity of the event sequence. Of course there were many things (complex systems) whose activities and interactions comprise the sequence of events leading to the start of World War I, but these things are not appropriately thought of as themselves forming a single larger system. The reason is that the particular sequence of events that led World War I to occur when and how it did is unique. It depended upon a variety of particular factors, like the travel plans of the Archduke and the particular political events occurring in Serbia, Austria, and elsewhere in June and July of 1914. There is just one instance of the event type “beginning of World War I.”

The case is rather different with toilets. The event of my depressing the lever on my toilet reliably flushes my toilet. It has done so many times in the past, and will do so many times in the future. Moreover, there are millions of toilets of essentially the same type which display the same reliable behavior.

8. The causal chain will not in general be single stranded. It may involve a network of interlocking events of arbitrary complexity. Note that some philosophers of science have taken path diagrams to represent causal chains, but these path diagrams generally represent type-level rather than singular causal associations.

The mechanistic viewpoint leads to a clearer understanding of the relation between type and token causal claims. Mechanisms like toilets reliably produce certain kinds of event sequences. A causal claim about an instance of this event sequence—for instance, that my depressing the lever on the toilet on 12 noon, January 1, 2000 caused the toilet to flush—is a singular causal claim. It is true because there exists a mechanism (*sensu* process) leading from the lever-depressing event to the flushing event. There is also a mechanism, *sensu* system, which produces this event. We call it a system because it is stable, and it is so arranged that it is causally responsible for many sequences of this type. It is the existence of such stable systems that permits us in these cases to make causal generalizations. I can say generally that depressing the handle of my toilet causes it to flush, because there are many singular causal processes of this type created by the stable arrangement of parts in my toilet. I can say, yet more generally, that depressing toilet handles causes toilet flushings, because there exist many toilet mechanisms of similar design whose stable configurations induce appropriate singular causal sequences. Note that while these causal generalizations about toilets are (I take it) true, there is not a contextually unanimous relation between lever depressing and flushing. Toilets can break in certain ways—the drain can become clogged, the chain connecting lever to flapper valve can become detached, etc.—and when they do break the connection between lever depressing and flushing breaks down. These constitute causally relevant background contexts in which lever depressings don't increase the probability of toilet flushings.

As a more scientifically interesting example, consider the claim that smoking causes lung cancer.⁹ As we saw, the contextual unanimity approach fails to work because there are instances in which smoking does not raise, or perhaps even lowers, the probability of cancer. On the mechanistic analysis, we begin with the examination of singular claims, looking at the mechanism whereby ingestion of cigarette smoke leads to the growth of tumors. Carcinogens operate by mutating genes in the cell nucleus into oncogenes, genetic variants of normal genes that in various ways stimulate cell growth. However, there are two reasons that exposure to carcinogens does not invariably lead to cancer. In the first place, carcinogens do not invariably produce mutations. In the second place, even if a mutation occurs to produce an oncogene, cells generally have other genes that suppress tumor growth. Thus, it is only in the case where mutations occur both to an oncogene and possible suppressor genes that a cancer begins to grow. It is only by knowing in detail the mechanism by which carcinogens change the structure of cellular DNA and by which these changes

9. This discussion is abbreviated from Thagard's discussion of cancer mechanisms (1999, 31–34).

accelerate cell growth that we can establish beyond doubt that in a particular instance the carcinogens in tobacco smoke caused cancer.¹⁰

Only after the mechanism of carcinogenesis has been established in particular cases is there a basis for making the causal generalization that smoking causes lung cancer. The crucial step is to recognize that within the population about which the generalization is made (namely human beings), lung cells have very similar structures, and that these structures reliably tend to produce cancers.

While it is necessary to establish singular causal connections before making generalizations, the criteria involved in making causal generalizations are in part epistemic and pragmatic. For a causal generalization to be true there must be a reliable mechanism, but there are no non-pragmatic criteria for deciding how reliable is reliable enough. Similarly, we generalize by segregating individuals into classes and generalizing over those classes, but there are various ways to create these classes and our choices are certainly influenced by pragmatic concerns.

It is beyond the scope of this paper to give a full account of the pragmatics of causal generalizations. For present purposes, it is sufficient to point out that, whatever the details of this account, singular causal claims are conceptually prior to causal generalizations. Causal generalizations are true because there exist multiple instances of reliable mechanisms (*sensu* systems) of certain types. Reliable but not unbreakable mechanisms give rise to causal generalizations that are robust but not contextually unanimous.

The distinction between mechanisms as systems and mechanisms as processes is connected to another distinction that may be implicit in Sterelny and Kitcher's description of evolutionary theory as being concerned with "central tendencies," and is explicit in Sterelny (1996); this is the distinction between what Sterelny calls robust-process and actual-sequence explanations. An actual-sequence explanation "identifies *the particular possible world that we inhabit*" while the robust-process explanation "*compares our world to others*" (Sterelny 1996, 195). An actual-sequence explanation identifies the detailed sequence of events that leads to some other event in the actual world, while a robust-process explanation identifies events which are invariant across many nearby possible worlds. In this parlance, it would appear that the explanation of World War I in terms of the assassination of the Archduke is an actual-sequence expla-

10. I am not suggesting that epidemiological studies showing high correlations between smoking and cancer do not give good reason to believe that smoking causes cancer. I am merely suggesting that the causal claim is ultimately a claim that there exists a mechanism linking smoke to carcinogenesis and that epidemiological studies provide at best indirect evidence of this claim.

nation, since the particular sequence of events leading up to and following upon the assassination seem highly contingent. It was by no means inevitable (i.e., true in all nearby possible worlds) that the Archduke was going to be assassinated and that this assassination would lead to the outbreak of global war. On the other hand, the explanation of toilet flushing is a robust-process explanation.¹¹

Sterelny's main point in his 1996 is that the same biological event is susceptible to two different kinds of explanations. It is also clear, in light both of this discussion and of Sterelny and Kitcher's (1988) claim that evolutionary theory is concerned with the "central tendencies" of the evolutionary process, that Sterelny believes that evolutionary explanations are robust-process explanations. I am in agreement with both of these claims, but I believe that where Sterelny falls short is in characterizing the conditions that give rise to robust processes. Like Lewis' approach to the truth conditions for counterfactuals, Sterelny seeks to characterize robust processes in terms of actual sequences in similar possible worlds. But the concept of a similar possible world is fraught with epistemic problems, and at any rate, is not a concept that a biologist would appeal to in explaining the robust character of certain evolutionary processes. Much more congenial is the concept of a mechanism (*sensu* system), since mechanism is a concept frequently appealed to by biologists and reliable mechanisms give rise to robust processes. And, while I agree with Sterelny and Kitcher that evolutionary explanations should be concerned primarily with robust processes, when we look for the mechanisms that underlie these processes we shall find reasons to support Sober and Lewontin's suspicions of genic selectionism.

5. Applying the Mechanistic Approach to the Units of Selection Problem.

Sober and Lewontin's attack on genic selectionism is actually two tracked. First, they claim that genic selection coefficients are artifactual, because single genes don't make contextually unanimous contributions to viability and fitness. Second, they claim that they are artifactual because they ignore "the biological facts." Since the first of these attacks has failed, I propose to concentrate on the second. The biological facts to which Sober and Lewontin allude are, in fact, facts about the mechanisms responsible for the production of phenotypic traits that differentially affect viability and fitness.

Let us consider such a strategy as it applies to the sickle cell case. The

11. Interestingly, both Sterelny and I use the assassination of the Archduke as an example of an actual-sequence explanation, though I had written this passage before coming upon Sterelny's example. Whether this choice of examples is the product of some unseen robust process or a case of pure chance I can only speculate.

distribution of alleles *A* (wild type) and *a* (sickle cell) is a product of selection. The question is, is selection occurring for each allele (*A* or *a*), or for allele combinations (*AA*, *Aa*, and *aa*)? The mechanistic approach suggests the latter alternative, because developmental mechanisms produce three different phenotypic traits—namely, three different types of red blood cells—depending upon which of these three combinations appears in the genetic material of the individual organism. The *AA* genotype leads to the production of red blood cells containing only hemoglobin *A*.¹² These cells are the most efficient of the three at oxygen transport, as well as the most flexible and least fragile (i.e., least prone to hemolysis). The *aa* genotype leads to the production of red blood cells containing hemoglobin *S* in place of hemoglobin *A*. The altered conformation of the protein diminishes the capacity of red blood cells to carry oxygen and increases rigidity and fragility of cells. This increased rigidity and fragility leads to sickling, which leads to blocking of capillaries and hemolysis (release of hemoglobin into plasma). These cellular traits lead to anemia with consequences for growth, resistance to infection, and ultimately the probability that a person survives and reproduces. In the heterozygote, both the wild-type allele and the sickle cell allele are used by red blood cells to produce hemoglobin, leading to the presence of both normal and *S* hemoglobin. Red blood cells of this kind are intermediate in terms of oxygen-carrying capacity, rigidity, and fragility. However, because red blood cells are “over-engineered,” in most circumstances sickling does not occur, and red blood cells are able to deliver adequate supplies of oxygen to other systems in the body. Bearers of the sickle cell trait have somewhat reduced life expectancy, but the trait does not seem to have a significant negative impact on either viability or fecundity. On the other hand, as noted above, the red blood cells containing some hemoglobin *S* are resistant to infection by the malarial parasite, with a consequent increase in the viability and fecundity.

I have told the story of how the sickle cell allele affects fitness in somewhat more detail than is usual, because it is precisely these “biological facts”—the mechanisms by which genes produce traits affecting fitness—that indicate what we should take to be the units of selection. The claim I am making in the sickle cell case is that the traits relevant to explaining the differential fitnesses of organisms are produced by diploid types rather than single alleles. The causal story of the production of these traits makes reference to both of the homologous alleles.

One objection to this analysis is as follows: If it is supposed that we

12. For simplicity I shall assume that the only two hemoglobins are *A* and *S*. There are other types that may be present, but they are not relevant to the discussion here. See Wagner and Mitchell 1965, 615–622.

must include in the genotype all those genes that are implicated in the causal processes giving rise to a trait, one must include *much more* than the genes at a single locus. The production of hemoglobin of any variety relies upon a complex developmental mechanism involving many genes, and in the absence of these genes, none of the character variants will manifest themselves.

This objection is really an instance of a much more general concern about causal claims. Commonly we assert that some single event or state of affairs was "the" cause of some other event, when in reality it was only the supposed cause in conjunction with a variety of other events or states of affairs that would be sufficient to bring about the effect. For instance, we might say that a fire was caused by a short circuit, but evidently the short circuit would not have been sufficient to bring about the fire in the absence of other factors, such as a nearby collection of flammable material, oxygen, etc. J. L. Mackie (1980) has provided an analysis of this kind of usage that explains and justifies isolating a single factor as the cause. According to Mackie, singular causal claims are made against the background of a causal field, which is a set of conditions that are presumed to hold. When one claims that *A* was the cause of *B* relative to a causal field, one is claiming that, relative to the causal field, it was *A* that made the difference between *B*'s occurring and its not occurring. Applying this analysis to the case at hand, the wide variety of genes whose presence plays some role in the development of red blood cells will be part of the causal field. It should be admitted that to a certain degree the question of what falls within the causal field is a pragmatic issue, and to admit as much is to raise certain doubts about the claim that there is a unique unit of selection. However, there are objective features of the particular scenario which largely determine the choice of the causal field. Most notably, any genes that are required for the development of red blood cells generally will be shared by all viable human beings. Thus, when we ask, what caused this person to have the sickle cell trait, we say that it is the presence of the two recessive sickle cell alleles, since it is this pair of genes that is sufficient relative to the causal field to produce sickling red blood cells.

A second objection to the proposed analysis of the sickle cell is that, contrary to what I have said, single alleles do produce traits that differentially affect fitness. Having the sickle cell allele is necessary to produce *S* hemoglobin, while having the wild-type allele is necessary to produce *A* hemoglobin. On what basis does one claim that there are three different traits, rather than two?

The problem here, of course, is that how one counts traits depends upon the level of organization at which those traits are described. Within the cell, both *A* and *S* hemoglobin are produced, so there is codominance, while, at the cellular level, the sickle trait red blood cell is an intermediate,

so there is incomplete dominance. But why prefer the cellular level for a description of the phenotype? The reason, I suggest, is because the cell is the lowest level at which the phenotypic traits have a reliable effect on the fitness of an organism.

As I suggested at the beginning of the paper, the genic selectionist will here argue that I have begged the question, by simply assuming an ecological as opposed to a genetic environment concept. Waters in particular has argued that one can provide an equally adequate description of the same causal process at the genetic level, simply by arguing that that net changes in gene frequencies for the sickle cell allele are the sum of two selection processes, one in genetic environments containing A as the homologous allele and another in environments containing a (Waters 1991, 562). Waters takes this case as exemplifying a more general principle, that there is generally more than one way to correctly describe a given causal process. About this general claim, I cannot disagree. Indeed, the problem of deciding how to draw the boundary between selective domain and environment is an instance of the general problem of how to draw the boundary between a mechanism and its environment—a problem which I have suggested elsewhere (1992) has pragmatic dimensions.

But while I grant the general claim, I must take issue with the claim that we can't judge the organismic description of the sickle cell case superior to the genic description. Viewed as a singular explanation of the viability of a single allele, there is not much to choose between the accounts. The problem is that when one applies the gene's eye view to the description of a process involving the evolution of a population of organisms, the gene's eye view obscures important explanatory information. If each gene is within its own environment, the gene's eye view produces a massive proliferation of environments.

Sober and Wilson (1994, 1999) have argued that the key to understanding the levels of selection problem is the concept of common fate. A group of entities (of genes or organisms) fall within a single selective domain when they live or die together in virtue of some common property. In this case, human beings live or die as wholes, and consequently it is appropriate to model them as interactors facing a single environment. Once this fact is granted, one can ask what genes are implicated in the causal mechanisms that produce traits that affect the interactors' common fate.¹³

I have up until now articulated a mechanistic approach to the units of

13. Sober and Wilson directly repudiate Wimsatt's additivity criterion (1994, 543) and by extension the principle enunciated in Sober and Lewontin. Sober seems to realize that he must give up on this as an analysis of the units of selection, but he does not appear to recognize that it is symptomatic of more general problems with his approach to causation.

selection problem in the context of the particular case of the sickle cell allele. I would now like to generalize this to give a general definition of a unit of selection. To repeat, I am here considering only the units of selection question. I take it as given that the level of selection is the organism, and the unit of selection is some part of the organism's genome. The definition could be generalized to consider units responsible for the production of traits in interactors at other levels. In what follows, I use the term *genotype* to refer to some collection of loci, ranging from the location of a single allele to the collection of all loci in the genome. Each possible combination of alleles that can occupy the loci in the genotype I shall call a *variant* of the genotype. We have then the following definition:

A genotype G is a unit of selection in environment E if, and only if, instances of the different variants of G reliably produce phenotypic characters that have differential capacities to influence an individual organism's fitness in E , and these characters are not reliably produced by variants of any smaller segment of the genome.

This definition clearly construes the unit of selection as a causal concept. The mechanistic element appears in the requirement that instances of variants of G reliably produce certain phenotypic characters. Reliable production of these characters is, as I have argued, the result of the operation of developmental mechanisms. On the singularist and mechanistic approach I recommend, the concept of unanimous effect of a type is replaced with the concept of reliable operation of a singular mechanism. Numerous instances of reliable mechanisms of a certain kind may give rise to unanimous or (more likely) near-unanimous effects. However, type causal relations are entirely derivative from token causal relations.¹⁴

6. Conclusion. Sober and Lewontin claim that genic fitnesses are in general artifacts, while genotype fitnesses are not. In fact *both* are artifacts in some sense. Fitnesses in general are properties of types, but evolution ultimately occurs as a result of the differential reproductive success of individual organisms. The success (or lack thereof) of individuals in turn is caused by the various traits they possess. This fact is not meant to suggest that it is inappropriate to talk of evolution of genotypes. Evolutionary explanations are meant to explain changes in the frequency of genotypes and phenotypes. However, we must get the explanatory order correct. We ex-

14. While I shall not argue for the claim in detail here, it seems to me that my definition is really not far from the additivity definitions proposed by Wimsatt (1980) and Lloyd (1994). The largest difference is that, while the additivity definitions are based upon variations in fitness, mine is based upon the mechanisms giving rise to traits that produce variations of fitness.

plain the fitness of a genotype by reference to the reproductive success of individuals having that type, rather than explaining the reproductive success of the individual by reference to its instantiating a genotype. One genotype is fitter than another because on average the individuals that are instances of that type are more reproductively successful than another.

Although this paper has focused on the nature of causal claims about selection processes, the argument offered here (see also Glennan 1997) has general applicability to any domains (including most or all of the life and social sciences) in which causal generalizations are made about populations. If, as I have argued, causal claims are ultimately singular claims about mechanisms and mechanistic processes, we must reconceive the causal generalizations in those domains as causal upshots in particular populations. And to achieve better understanding why these generalizations are true, we must understand the mechanisms at work in the individuals that comprise these populations.

REFERENCES

- Brandon, Robert (1982), "The Levels of Selection", in Peter D. Asquith and Thomas Nickles (eds.), *PSA 1982*, Vol. 1. East Lansing, MI: Philosophy of Science Association, 315–322.
- Cartwright, Nancy (1979), "Causal Laws and Effective Strategies", *Noûs* 13: 419–437.
- (1989), *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.
- Craver, Carl (2001), "Role Functions, Mechanisms, and Hierarchy", *Philosophy of Science* 68: 53–74.
- Cummins, Robert (1975), "Functional Analysis", *Journal of Philosophy* 72: 741–765.
- Dupré, John (1984), "Probabilistic Causality Emancipated", *Midwest Studies in the Philosophy of Science* 9: 169–175.
- (1993), *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- and Nancy Cartwright (1988), "Probability and Causality: Why Hume and Indeterminism Don't Mix", *Noûs* 22: 521–536.
- Glennan, Stuart (1992), *Mechanisms, Models and Causation*. Ph.D. Dissertation. Chicago: University of Chicago.
- (1996), "Mechanisms and the Nature of Causation", *Erkenntnis* 44: 49–71.
- (1997), "Capacities, Universality, and Singularity", *Philosophy of Science* 64:4: 605–626.
- (forthcoming), "Rethinking Mechanistic Explanation", *Philosophy of Science* 69 (Proceedings).
- Humphreys, Paul (1989), *The Chances of Explanation*. Princeton: Princeton University Press.
- Kauffman, Stuart (1970), "Articulation of Parts Explanation in Biology and the Rational Search for Them", in Roger C. Buck and R. S. Cohen (eds.), *PSA 1970*. Boston Studies in the Philosophy of Science, Vol. 8. Dordrecht: D. Reidel, 257–272.
- Lewis, David (1973), "Causation", *Journal of Philosophy* 70: 556–567.
- Lewontin, Richard C. (1974), *The Genetic Basis of Evolutionary Change*. New York: Columbia University Press.
- Lloyd, Elisabeth (1994), *The Structure and Confirmation of Evolutionary Theory*. Princeton: Princeton University Press.
- Machamer, Peter, Lindley Darden, and Carl Craver (2000), "Thinking about Mechanisms", *Philosophy of Science* 67: 1–25.
- Mayr, Ernst (1963), *Animal Species and Evolution*. Cambridge, MA: Harvard University Press.

- Mackie, John Leslie (1980), *The Cement of the Universe*. Oxford: Clarendon Press.
- Majerus, Michael E.N. (1998), *Melanism: Evolution in Action*. Oxford: Oxford University Press.
- Ray, Greg (1992), "Probabilistic Causality Reexamined", *Erkenntnis* 36: 219-244.
- Sober, Elliott, and Richard C. Lewontin (1982), "Artifact, Cause, and Genic Selection", *Philosophy of Science* 49: 157-180.
- Sterelny, Kim (1996), "Explanatory Pluralism in Evolutionary Biology", *Biology and Philosophy* 11: 193-214.
- and Philip Kitcher (1988), "The Return of the Gene", *Journal of Philosophy* 85: 339-360.
- Thagard, Paul (1999), *How Scientists Explain Disease*. Princeton: Princeton University Press.
- Wagner, Robert P., and Herschel K. Mitchell (1965), *Genetics and Metabolism*, 2nd ed. New York: John Wiley and Sons.
- Waters, C. Kenneth (1991), "Tempered Realism about the Force of Selection", *Philosophy of Science* 58: 553-573.
- Wimsatt, William (1980), "Reductionist Research Strategies and Their Biases in the Units of Selection Controversy", in Tom Nickles (ed.), *Scientific Discovery*, Vol. 2. Dordrecht: Reidel, 213-259.
- (1981), "The Units of Selection and the Structure of the Multi-Level Genome", in Peter Asquith and Ronald Giere (eds.), *PSA 1980*, Vol. 2. East Lansing, MI: Philosophy of Science Association, 122-183.