

Journal of Philosophy, Inc.

Apportioning Causal Responsibility

Author(s): Elliott Sober

Source: *The Journal of Philosophy*, Vol. 85, No. 6 (Jun., 1988), pp. 303-318

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2026721>

Accessed: 01/03/2010 18:17

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=jphil>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

APPORTIONING CAUSAL RESPONSIBILITY*

IS this particle's acceleration due more to gravity or to electricity? Classical physics regards this question as well-conceived. We may answer by examining physical details of the system before us and using our knowledge of pertinent laws. Newton's law of gravity tells us how mass produces a gravitational force; Coulomb's law of electricity shows how charge generates an electrical force. Each of these *source laws* can be connected with the *consequence law* " $F = ma$ " to determine which force induces the greater component acceleration.¹

In this Newtonian case, two questions seem interchangeable: What contribution did gravity (or electricity) make to the particle's acceleration? What difference did gravity (or electricity) make in the particle's acceleration? These questions are simultaneously addressed by investigating two counterfactual questions: How much acceleration would there have been, if the gravitational force had acted, but the electrical force had been absent? How much acceleration would there have been, if the electrical force had acted, but the gravitational force had been absent? Classical particles obey John Stuart Mill's principle of the composition of causes.² The result of the two forces is just the sum of what each would have achieved, had it acted alone. We see what each contributed by seeing what difference each made in the magnitude of the effect.

For the accelerating Newtonian particle, apportioning causal responsibility is a local matter—we can assess the contributions of gravity and electricity just by discovering physical facts about the particle and the forces that affect it. In this context, the contribution a cause makes and the difference it makes seem to be one and the same issue. It would be natural to see here an example of some universal principle concerning how science apportions causal responsibility: natural, but unwarranted. For careful examination of how causal responsibility is apportioned in another context casts these questions in an entirely new light.

"Is Jane's height due more to her genes or to her environment?"

* I am grateful to Ellery Eells, Brian Skyrms, Peter Woodruff, and to the editors of this JOURNAL for their suggestions. I also thank the National Science Foundation for financial assistance.

¹ The distinction between source law and consequence law is elaborated and applied to evolutionary theory in my *The Nature of Selection: Evolutionary Theory in Philosophical Focus* (Cambridge: MIT, 1984), ch. 1.

² *A System of Logic, Ratiocinative and Inductive* (New York: Harper, 1859).

Biologists have been taught to regard this question as meaningless. The proper way to formulate a question concerning nature and nurture, they often say, is at the population level. Junk the question about Jane and replace it with something like the following: "In the population of U.S. adults, how much of the variation in height is explainable by genetic variation, and how much by environmental variation?"

The example of Jane's height shows that the question of how much a cause contributes to the effect and of how much difference it makes in the effect are two questions, not one. The latter is answerable, though not locally, whereas the former is not answerable at all. There is no such thing as the way science apportioning causal responsibility; rather, we must see how different sciences understand this problem differently, and why they do so. The particle's acceleration (E) is an effect of gravity (C_1) and electricity (C_2), just as Jane's height (E) is an effect of her genes (C_1) and environment (C_2). But this parallelism belies the following differences:

	Newtonian Particle (Gravity/Electricity)	Ontogeny (Nature/Nurture)
How much did C_1, C_2 contribute to E ?	locally answerable	meaningless
How much difference did C_1, C_2 make in E ?	locally answerable ³	answerable, but not locally
	Questions are equivalent	Questions are not equivalent

The Newtonian approach to apportioning causal responsibility, I have noted, is based on a *theory*. So, too, is the biologist's approach to the nature/nurture dispute. Or, more accurately, it is based on a *technique*. Biologists use a statistical method known as the *Analysis of Variance* (ANOVA) to say which of an array of causal factors explains more and which explains less of the variation in the effect property found in a population. First and foremost, the method is used to explain the variation in height found in the population that Jane inhabits. After describing how this procedure works, I shall argue that ANOVA can be brought to bear on Jane herself, not just on the population she inhabits. The suggestion will be that ANOVA

³ A proviso will be registered later to defend this entry—namely, that the relevant counterfactuals be nonbacktracking.

can be used to describe which causal factor at work in the *singleton case* made the largest difference in the effect. So my proposal will connect a *populational phenomenon* (e.g., variation in height among U.S. adults) with a singleton phenomenon (e.g., Jane's height).

After explaining how the analysis of variance applies to the nature/nurture dispute, I shall try to say why apportioning causal responsibility proceeds so differently, depending on whether one analyzes an organism's ontogeny or a particle's acceleration. Why is a particle's acceleration decomposable in a way that Jane's height is not? This will lead back to Mill's principle of the composition of causes, which is much less central to the difference than might first appear. But, before considering how nature and nurture are disentangled, we must clarify the concept of locality itself.

I. WHAT IS LOCALITY?

The concept of locality requires more clarification than I shall be able to provide here. So, in the end, I shall have to rely on a somewhat intuitive grasp of the question at issue. Nevertheless, a few remarks may help direct the reader's attention to the right issues.

First, I should note that the problem of locality raised by apportioning causal responsibility is quite different from the one usually discussed in physics. The problem of locality in quantum mechanics has to do with what must be true of two token events if one causes the other. Very roughly, the question is whether a physical signal must proceed from one to the other, passing continuously through a series of intervening space-time points.

The question of causal magnitude is quite different from the problem of whether there can be action at a distance. Suppose, just for the sake of argument, that the physicist's question is correctly answered by some suitable thesis of locality.⁴ Even so, this would not answer our question about Jane. Granted, there is a continuous path linking her genes and environment, on the one hand, and her subsequent height, on the other. Locality in the physicist's sense is thereby assured, but the question of apportioning causal responsibility has yet to be addressed. We have yet to see what it means to assess the relative contributions of Jane's environment and genes to her height.

A second problem raised by the question of locality is that it must be formulated in such a way that it is nontrivial. Consider the humdrum fact that causes are rarely in themselves sufficient for their effects. The match's being struck caused it to light, even though it

⁴ There are several physical formulations which would have to be sorted out here, as John Earman shows in his "What Is Locality?" in Peter Achinstein, ed., *Theoretical Physics in the Hundred Years Since Kelvin's Baltimore Lectures* (Cambridge: MIT, 1986).

would not have ignited if the match had been wet or if there had been no oxygen in the air. This is the trivial sense in which locality fails; whether one event will cause another depends on features of the world external to them both.

In the nature/nurture controversy, we acknowledge this familiar point when we say that genes are not in themselves enough. For Jane to reach a certain height, she must be raised in an appropriate environment. Genes are no good, unless supplemented by numerous meals. Nor is environment in itself sufficient, since there are genetic configurations that will impede Jane's growth, no matter how much milk she drinks. The cliché is that development is the result of a gene/environment interaction.

Be this as it may, our problem is not thereby resolved. Neither genes nor environment are themselves sufficient. But, when they conspire to produce Jane's height, how are we to assess their relative contributions? In particular, is this matter resolvable by attending to features of Jane's physical development from zygote to adult? Or must we look to features outside of Jane's ontology and the genes and environment that produced it?

It seems clear that the problem of locality depends on how one carves up nature into discrete physical systems. If Jane's sequence of environments and genes comprise the system in question, then locality will be refuted if we find that factors external to them play a role in determining which mattered more. If from the first we think of the population in which Jane resides as the relevant unit of inquiry, however, then we may find that locality is verified. The causal facts about Jane may turn out to depend on the population, but the facts about the population may not depend on anything external to it. If so, whether locality is true or false will turn on whether we pose our question about Jane or about her population. As emphasized earlier, the principal question of interest here concerns "singleton" physical systems (Jane, the particle), not populations of such.

The question of locality presupposes that descriptors of an individual can be divided into ones that are "intrinsic" and ones that are not. Although the predicate 'x lives in a genetically homogeneous population' is true of individuals, it nonetheless expresses a nonlocal property of them. I do not believe that the distinction needed here can be drawn syntactically; nor should it be, since the conclusions reached should be language independent. Nor do I have an informative semantic criterion to suggest. As in other problem areas (e.g., how the thesis of determinism should be formulated), we must rely on a somewhat intuitive grasp of what it is for a property to be local.

II. THE ANALYSIS OF VARIANCE

Why do biologists think the nature/nurture question is meaningless unless formulated in terms of population variation? To be sure, it is silly to think of Jane's six-foot stature as decomposing into two feet due to genes and four feet due to environment, as if the genes built Jane from the navel up, while the environment took care of what lies below. But one silly suggestion does not show that the question of causal magnitude must be understood nonlocally.

To see why biologists think that apportioning causal responsibility requires a populational analysis, we need to consider the basics of the ANOVA technique. So as to avoid the macabre prospect of experimentally manipulating human beings, let us switch from Jane to a corn plant. It has a certain height; we want to know how genes and environment affected this outcome. As its name suggests, the analysis of variance understands this problem at the population level. Instead of focusing exclusively on the singleton case, we consider different possible corn genotypes (G_1, G_2, \dots, G_n) and different possible environmental conditions (E_1, E_2, \dots, E_m)—amount of fertilizer, for example. There are then nm possible "treatments." We might divide a field into different plots, the corn plants on each plot receiving one of the treatments (G_i & $E_j, i = 1, 2, \dots, n$ and $j = 1, 2, \dots, m$). Each plot of land would contain the same number of plants; within a plot, the plants would have identical genotypes and would receive identical amounts of fertilizer. We then would record the average height within each plot and enter the result into an n by m ANOVA table:

		Environmental Variable				
		E_1	E_2	\dots	E_m	Averages
	G_1	x_{11}	x_{12}	\dots	x_{1m}	$M_{1.}$
Genetic Variable	G_2	x_{21}	x_{22}	\dots	x_{2m}	$M_{2.}$
	
	
	
	G_n	x_{n1}	x_{n2}	\dots	x_{nm}	$M_{n.}$
Averages		$M_{.1}$	$M_{.2}$		$M_{.m}$	

The marginal averages simply record the average height of plants with the same genotype which receive different amounts of fertilizer (the various $M_{i.}$ s) and the average height of plants receiving the same amount of fertilizer, but having different genotypes (the various

$M_{.s}$). One final number, not shown in the above table, is the grand mean M , which is simply the average height across all treatments.

It is a matter of arithmetic that the difference between any cell in the above table and the grand mean ($x_{ij} - M$) must equal the sum of three terms: the so-called *environmental main effect* ($x_{ij} - M_i$), the so-called *genetic main effect* ($x_{ij} - M_j$), and the so-called *gene-environment interaction I* (this last term is a fudge factor, which merely insures that the sum comes out right):

$$(x_{ij} - M) = (x_{ij} - M_j) + (x_{ij} - M_i) + I.$$

Rearranging the above equation a little, we can show how the average height within any treatment plot must be related to the averages defined on the entire field of plots:

$$x_{ij} = M + (x_{ij} - M_j) + (x_{ij} - M_i) + I.$$

What I have set forth so far is quite unobjectionable.⁵ But what use is this partitioning when it comes to the task of causal explanation? I begin with a negative claim: ANOVA does *not* explain the occurrence of singleton effects (i.e., why a given plant has a height of x_{ij}).⁶ In the next section, something positive will be defended concerning the relevance of ANOVA to explaining singleton phenomena.

It seems quite clear that the height attained within a given plot does not causally depend on the heights attained in other plots. This is true even though the height within a plot is the sum of terms that represent that plot's deviations from various means (and a corrective factor, I). I would supplement this causal claim with an explanatory one: If you want to know why the corn plants in a given treatment attained the height they did, do not cite the information codified in the analysis of variance formula. The analysis of variance does not identify causes or explain the upshots found in individual cells.

Consider the fact that a given treatment (G_i & E_j) can be embedded in different experimental designs. For example, I might investigate plants that have the same genotype and place them in different environments. In this case, I shall find that the genetic main effect is zero. As biologists say, all the variance in height is explained by

⁵ For the statistical details of this method, see Robert Sokal and F. James Rohlf, *Biometry: The Principles and Practice of Statistics in Biological Research* (San Francisco: Freeman, 1969). I have omitted mention of an "error" term; this plays a role in ANOVA inference, but does not affect the points about causality to be made in what follows.

⁶ The following arguments are explored in somewhat more detail by Richard Lewontin in "The Analysis of Variance and the Analysis of Causes," *American Journal of Human Genetics*, xxvi (1974): 400-411; and by myself in *The Nature of Selection*, *op. cit.*

environmental variance. Likewise, I might see how plants of different genotype fare in a single environment. In this case, the environmental main effect will be zero, and all the variance in height will be explained by genetic variance. This shows how misguided it is to interpret the analysis of variance as explaining why a plant in a given treatment cell attained the height it did. If I run the experiment one way, I shall say that the genes played no role; if I run it in the other way, I shall conclude that the environment played no role. And, if I vary both treatments, as I did initially, I shall doubtless conclude that genes and environment both played a role.

An additional reason for rejecting the analysis of variance as a device for identifying the causes that contributed to a singleton effect is provided by considering what this procedure says about characteristics that are universal in a population. Let us switch from corn-plant height to the human characteristic of having a single head. If we consider a population in which everyone has just one head, the analysis of variance will tell us that the genetic and environmental main effects are zero. It would be a mistake to conclude from this, however, that Jane's environment and genes played no role in providing her with a head.

The calculations derived from an ANOVA table do not allow one to deduce what the causes of a given singleton effect are. But this is not to deny that ANOVA can be used to describe the relative contributions that different factors make to explaining the pattern of variation of a trait in a population. Although Jane's genes obviously contribute to her having a head, it may or may not be true that genetic variation helps explain the way that phenotype varies in the population. This is perhaps why biologists have thought of apportioning causal responsibility as fundamentally a population level problem, not one that can be meaningfully addressed for the singleton case. I now want to suggest, however, that the proper use of ANOVA at the population level has relevance to a question about singleton events. One can say which causal factor at work in a singleton event made the greatest difference by embedding that singleton effect in a population, and then judiciously applying the ANOVA to that population.

III. THE COUNTERFACTUAL TEST

I suggest that causal responsibility for Jane's height can be apportioned between her genes and environment by considering two counterfactuals. How tall would Jane be if she had different genes, but the same environment? How tall would Jane be if she had a different environment, but the same genes? Environment is a more powerful influence than genes, if changing her environment would lead Jane's height to depart more from its actual value than would changing her genes.

To answer these two counterfactuals questions, we must elaborate. If Jane were raised in an environment different from the actual one, what would her environment be like? And, if Jane had a different complement of genes, what genes would she in fact have?⁷ My suggestion is that the way we fill in the details here is often nonlocal. When this is so, the question of causal magnitude is also a nonlocal one.

When we consider what genes Jane would have had if she had been differently endowed, we do not consider gazelle genes or Martian genes. In some sense, we go to a more "similar possible world." We may consider the other genes that were available from Jane's parents. If these are the same as the ones she actually possesses (i.e., Jane and her parents are all homozygous for the genes in question), we may go back to Jane's grandparents. Or perhaps, rather than tracing the lineage backwards in time, we may look at the other alleles present in the local population that Jane inhabits.⁸ If there are only two alleles in this local population and Jane has one of them, then it is natural to say that she would have had the other one, if she had had a gene different from the one she actually possesses. If there are several, then we might calculate Jane's expected height as a weighted average of the heights she would have had with each alternative gene, the weighting being supplied by the genes' population frequencies.

Similarly, when we ask what environment Jane would have inhabited had she not inhabited the one she in fact did, we do not imagine her floating weightlessly in outer space, nourished by a diet of candy bars. This is not a "similar possible world." Since human beings inhabit such an incredible variety of environments, it often will be impossible to say which environment Jane would have inhabited, had she pursued a different way of life. If so, a natural strategy is to do what we did above for the case of multiple alleles. We assign different possible environments a probability and then calculate Jane's expected height as a weighting over these.⁹

⁷ Like all counterfactuals, the two at issue here at times may be vague or indeterminate. To the degree that this is so, I suggest that the question of causal magnitude also is vague or indeterminate.

⁸ From this point on, I shall imagine that Jane is a haploid organism, so that I can talk of her having a given gene (not two genes) at a locus. This is simply to simplify the exposition.

⁹ The probabilistic formulation of the counterfactual test advocated here for apportioning causal responsibility is basically the same as the one I describe in "Causal Factors, Causal Inference, Causal Explanation" [*Proceedings of the Aristotelian Society, Series B* (1986): 97–113] for inferring which of the causal factors present before an effect was its token cause.

It seems intuitive that, in both these problems, we elaborate our counterfactuals by attending to features of the world which are extrinsic to Jane’s actual genes and environment and the height she actually attained. It is no intrinsic feature of Jane which dictates what genes she would have had, if she had had a different genome. The same holds true of her environment.

We can illustrate this idea in a thought experiment akin to Hilary Putnam’s story about twin earth.¹⁰ Jane lives in environment E_2 and has genotype G_2 . We need to consider what her environment would be, if it differed from E_2 , and what her genotype would be, if it differed from G_2 .

We shall consider two possible contexts. In the first (I), Jane lives in a population in which the only alternative to E_2 is E_1 , and the only alternative to G_2 is G_1 . In the second (II), Jane inhabits a population in which the only alternative to E_2 is E_3 , and the only alternative to G_2 is G_3 . We shall see how this contextual difference leads to contradictory assessments of how we should apportion responsibility for Jane’s height between her genes and environment.

The following table describes the height that Jane would have in each of the nine possible gene/environment combinations:

		Environments		
		E_1	E_2	E_3
Genes	G_1	I x_{11}	x_{12}	x_{13}
	G_2	x_{21}	x_{22}	x_{23}
	G_3	x_{31}	x_{32}	x_{33} II

Jane’s actual height is given by x_{22} . If she lives in a population of the first type (I), then we assess the relative impact of her genes and environment by comparing x_{12} and x_{21} . If she lives in a population of the second type (II), we must compare x_{23} and x_{32} . Suppose that $x_{22} = x_{12} = x_{23}$. This means that, in situation I, her genes make no difference to her height; if she had had a different genome, she would have had the same height. In situation II, however, it is her environment that makes no difference; if she had had a different

¹⁰ See “The Meaning of ‘Meaning’ ” in *Mind, Language and Reality: Philosophical Papers*, vol. 2 (New York: Cambridge, 1975), pp. 215–271.

environment, she would have had the same height. So, whether genes matter more than environment depends on features of the world external to Jane's own environment and genes. It follows that the question of causal magnitude is nonlocal.

The above table is just what one would obtain by pursuing an analysis of variance on causal factors antecedently identified. One would clone Jane and place her in different environments and place genetically different individuals in the environment she actually inhabits. Or, more precisely, since ethics prohibits this manipulation, one would try to assess what the consequences would be if these things were done. When we make counterfactuals actual, so to speak, we use the analysis of variance to apportion causal responsibility.¹¹

IV. COMMENSURABILITY AND LOCALITY

In arguing that the analysis of variance correctly apportions causal responsibility, I mean to suggest that it can identify the *difference* that various causes make in the observed effect. This is quite different from assessing how much each contributed. This latter question I interpret as a local one, which the facts of the matter about nature and nurture render unanswerable. Why this is so we can see by inventing a science fiction scenario in which the second question is quite intelligible.

Suppose height were the result of the accumulation of height particles, which organisms could obtain from their environment and also from their genes. Imagine that an individual's height is some increasing function of the number of height particles obtained from all sources. If so, we could look at local facts about Jane and say whether her genes or environment contributed more.

Of course, there are no such things as height particles, but this thought experiment suggests the following (somewhat vague) conjecture: For it to make sense to ask what (or how much) a cause contributes to an effect, the various causes must be commensurable in the way they produce their effects. Height particles would provide this common currency.


The idea of height particles shows the inadequacy of one diagnosis of why locality fails in the nature/nurture case. In the Newtonian example, we can ask what would happen if no gravitational force acted or if no electrical force were at work. In the nature/nurture

¹¹ When there is a univocal answer to the question of how tall Jane would be if her genes were different, probabilities need not intrude; when there are multiple-alternative possible genomes, a probabilistic weighting is needed. In both these cases, however, there may be disagreement over whether there is a uniquely correct answer to the question of causal magnitude, or whether the "choice of a reference class" is inherently arbitrary. My argument is neutral on this, since, in either case, the thesis of locality fails.

case, no significant answer is obtained by asking how tall Jane would have been if she had had no genes or no environment. The intelligibility of these questions is not essential, however. If there were height particles, the question of whether genes or environment contributed more would make sense, even though an organism requires genes and an environment of some sort if it is to exist at all.

Nevertheless, this difference between gravity/electricity and nature/nurture is not without its significance. I have claimed that the difference made by gravity and by electricity is a local matter; intrinsic features of the physical system under study determine what would happen if either force had acted alone. But we may embellish the details of this example to show that the issue of locality is, even here, not so straightforward.

We so far have considered a physical system made of a particle and the gravitational (G) and electrical (E) forces that affect its acceleration (A). Suppose we embed this physical system into a larger context. Let us imagine that some cause C —a switch, say—insures that a third force T of given magnitude would have come into play if and only if there had been no electrical force (E). The set-up is as follows, with a box drawn around the factors that comprise what I have so far called “the physical system”:



There now is an ambiguity in the question of what would have happened if gravity had acted, but electricity had not. We could interpret this, as we have so far, as meaning that gravity would have acted *alone*. Or we could reason that, if electricity had been absent, then the third force (T) would have acted in its stead.


That is, we must decide whether or not we allow counterfactuals to “backtrack.”¹² I believe that both interpretations have their place and that it is contextual clues that settle which reading is appropriate. If so, there is no univocal answer to the question of whether the difference made by gravity and electricity is locally determined. If

¹² I here use the vocabulary, if not the full-blown theory, developed by David Lewis in his “Counterfactual Dependence and Time’s Arrow,” *Noûs*, XIII (1979): 455–476.

we forbid backtracking, then the question of what difference gravity and electricity made reduces to the question of what each would have achieved if it had acted alone. This is a local matter. If we permit backtracking, however, then the question becomes nonlocal. We must consult physical facts extrinsic to the particle and the gravitational and electrical forces acting on it. Thus, the lower-left entry in the table presented in the first section requires that counterfactuals not backtrack.

No similar ambiguity can arise in the nature/nurture case. The reason is that it makes no sense to imagine an individual developing without any environment at all and that an individual will not develop at all, if it has no genes. The ambiguity in the gravity/electricity case was made possible by the fact that we can ask what would have happened if only one of those forces had been present. This provides the nonbacktracking and local reading of the question of what would have happened if one had been present and the other absent.

This local interpretation is not available in the nature/nurture case; we are forced to consider the wider context which dictates what alternative environmental and genetic factors would have been present, if the actual ones had been absent. Here the question of how tall Jane would have been, if she had had a different complement of genes (or a different environment) must be understood in terms of backtracking counterfactuals. Jane's environment E_2 and genes G_2 produce her height (H). Some causal factor (C) presumably determined that she would develop in environment E_2 , rather than in E_1 . If we reason that Jane would have developed in E_1 , if she had not developed in E_2 , we are backtracking; we are imagining that the cause that determined that she would grow up in E_2 and not in E_1 would have been different, if she had failed to grow up in E_2 . The causal facts are shown below:



I conclude that the question of what difference two causal factors made may be local or nonlocal in the Newtonian case, but that it must be nonlocal when the issue concerns nature/nurture.

Thus, there are at least two differences between the Newtonian particle and the real world of nature and nurture. Causal contribu-

tion is well-defined and local in the first, but not in the second. In addition, the two cases part ways over how the difference a cause makes is determined. This can be understood either locally or non-locally in the Newtonian case, but it must be understood nonlocally in the case of nature and nurture.

It may be thought that the key to the difference between the Newtonian and the biological examples lies in the fact that the former obeys Mill's principle of the composition of causes. As mentioned before, source laws describe which physical properties of the particle system generate this or that force, and the consequence law " $F = ma$ " allows the separate contributions of gravity and electricity to be described in a common currency. This generic concept of force is indispensable, if the contributions of each cause are to be comparable. But it now is time to see that it is entirely irrelevant whether the system under study obeys the principle to which Mill called our attention.

Biologists would use the terms 'additive' and 'nonadditive' to draw Mill's distinction between systems which obey the principle of the composition of causes and ones which are, in his terminology, 'heteropathic' (*op. cit.*, p. 213). We now may elaborate our story about height particles to show why additivity is neither necessary nor sufficient for locality.

As long as such particles exist, the relative contribution of genes and environment to height will be locally assessable, regardless of the formula that translates number of particles into inches and feet. Suppose that height increases as the number of particles goes from 0 to 50 and declines thereafter. If 40 particles were obtained from the environment and 40 from the genes, we would assign an equal contribution to each source. And this would be correct even though an individual with 80 particles is not twice as tall as an individual with 40.

This shows why additivity is not necessary for locality. To see that it is not sufficient, we must realize that the analysis of variance has no problem describing an additive relation between genes and environment; the data displayed in an ANOVA table are additive, if moving across any row involves adding a certain quantity to one entry to obtain the next and the same holds for moving down any column:

$$\begin{aligned}x_{i,j+1} &= x_{i,j} + c, & \text{for some } c \text{ and for all } i,j. \\x_{i+1,j} &= x_{i,j} + d, & \text{for some } d \text{ and for all } i,j.^{13}\end{aligned}$$

¹³ This is a sufficient condition for additivity, not a necessary one; the definition of additivity is just that the interaction term (I) is zero and so the value within each cell is the sum of the grand mean and the two main effects.

The case of two forces (F_1 and F_2) acting together, where their joint effect is simply the sum of what each would have achieved if it had acted alone, is just a special case of this formula:

	not- F_1	F_1
not- F_2	0	c
F_2	d	$c + d$

The entries in an ANOVA table for nature and nurture might be additive in this sense, but this would not give local sense to the question of how much of Jane's height was due to her genes and how much to her environment. I conclude that additivity is neither necessary nor sufficient for locality.

The idea of height particles shows that the question of causal magnitude can be approached both locally and nonlocally with respect to the same physical system. Even if such particles existed, we still could pursue the analysis of variance to obtain a nonlocal answer to the question of causal magnitude. We also could do an intrinsic analysis of the relative contributions of Jane's environment and genes to her height. Indeed, we could obtain very different answers from these two procedures. If her genes produced 50 particles, but her environment only 10, we might judge from this intrinsic perspective that her genes contributed more. If, however, there is little genetic variation in Jane's population, the analysis of variance might conclude that environment made more of a difference than genes.

This shows why saying how much Jane's environment and genes contributed to her height can be a separate matter from saying whether her genes or environment made the greater difference. In the Newtonian case, these two questions are interchangeable and local (provided that counterfactuals are not allowed to backtrack). In the case of nature and nurture, they are different issues. How much difference Jane's environment and genes made to her height is answerable, though not locally. How much each contributed is an unanswerable question, since the empirical fact of the matter (I presume) is that there are no such things as height particles.¹⁴ Apportioning causal responsibility involves two issues, not just one.

¹⁴ One could take the position that the questions of how much each factor contributes and of how much difference each makes are equivalent and nonlocal in the nature/nurture dispute when there are no local determiners of how causal responsibility should be apportioned. The fact that these questions come apart when we imagine a local determiner suggests, however, that the questions are better treated as separate ones from the start.

In the Newtonian case, the commensurability of gravity and electricity renders the question of causal contribution well-defined and local; in the nature/nurture case, we found neither commensurability nor locality. How much a causal factor contributes to its effect is, I would suggest, an inherently local question. But the question of how much difference a cause makes may be local or not, as is shown by the following example (suggested by Peter Woodruff).

Consider the fact that the distance traveled by a projectile shot from a cannon is influenced by both the muzzle velocity and the angle at which the gun is set. Suppose we fire a cannon and the shot goes half a mile. There is no saying *how much* each factor contributed to this outcome, nor which factor contributed more. The reason is that muzzle velocity and angle setting do not make their contributions in a common currency.

Yet, the question of *what difference* the two factors made may be locally or nonlocally assessable, depending on how we elaborate the story. Suppose the physical design of the gun shows that there are just two possible angle settings and two possible powder charges. If so, an intrinsic examination of the system allows us to fill out the relevant counterfactuals. Which factor made the larger difference is now a local matter.

On the other hand, suppose the gun has just one possible angle setting, though it may be packed with different amounts of powder. If the cannon maker designed several guns, each with its own fixed angle of fire, we would have to consider nonlocal facts to say what angle setting the gun would have had, if it had lacked the one it in fact had. Here the fixed setting is analogous to the fact that each of us has just one genetic endowment.

V. CONCLUSION

Causes are not necessary for their effects. It follows that one does not refute a causal claim by showing that the effect would have happened, even if the purported cause had not. Yet, when one shifts from the definition of causality to the problem of clarifying the concept of causal magnitude, something like the necessity thesis appears to be correct. If genes and environment are both causal factors influencing Jane's height, then genes have zero magnitude, provided that Jane would have had exactly the same height even if her genes had been different.

I therefore seem to find myself in the paradoxical position of saying that genes can be a cause of height, even if they are judged to have zero magnitude. But perhaps this air of paradox can be dispelled. It is not hard to fathom how causes can fail to be necessary for their effects. Suppose Watson would have shot Moriarty dead if Holmes had not. Holmes' pulling the trigger may have killed Mor-

iarty, even though Holmes' firing was not necessary for Moriarty's death. The point may be put by saying that Holmes' firing, in this case, "made no difference." Causes may make no difference, but they are causes nonetheless.

I also have argued that the relative contribution of a cause and the difference it makes in the effect are separate issues. In the case of our Newtonian particle, this may look like a distinction without a difference. Disentangling the roles of genes and environment, however, shows this distinction to be a real one. Indeed, it is not just the science of ontogeny which draws this distinction. Holmes' pulling the trigger may have made no difference, but he certainly made a contribution to Moriarty's death.

A paradox that is less easy to dispel concerns the question of locality. Waiving as we have the quite separate issues raised by quantum mechanics, we may insist that causality is a local phenomenon. Holmes' firing causes Moriarty's death because there is a continuous process leading from one to the other. The causal relation obtains in virtue of this local circumstance. Watson's standing in the wings is relevant to the question of necessity, but not to the question of causality.

I have argued, however, that, even if causality is local, the magnitude of causality need not be. Jane's own genes and environment locally conspire to produce her adult height. But apportioning causal responsibility is not, in this case, a local matter. This is because it is a factual question, though not a local one, as to whether her genes or environment made the larger difference in her height. And it is not a factual matter at all, much less a local one, as to how much each contributed.

ELLIOTT SOBER

University of Wisconsin/Madison