DEBORAH G. MAYO

UNDERSTANDING FREQUENCY-DEPENDENT CAUSATION

(Received April 24, 1985)

Causal effects often vary with the relative frequency of the causal property in the population of interest. (For example, "being a female candidate for high office causes a stir" only because such events are *rare*.) Such causal effects are called *frequency-dependent*, and they pose interesting problems for both philosophy of causality and philosophy of biology.[1] Elliot Sober[2] argues that certain cases of frequency-dependence "constitute counterinstances" of the causal model proposed by Ronald Giere.[3] To solve the problem Sober raises, Giere's model must provide a satisfactory solution to the question: *How can a property causally act on all members of a population when the property is effective only when it is rare*?

In response to Sober, Giere[4] has recently proposed a special extension to his model to accommodate frequency-dependent claims. John Collier's[5] proposed defense of Giere's original model restricts the domain of frequency-dependent claims. In this paper I argue that neither Giere's nor Collier's proposals adequately answers the question posed by Sober's problem. More importantly, I shall argue that neither proposed revision is needed; that Giere's original model adequately answers the above question – provided frequency-dependent causation is correctly understood. While my primary aim is to arrive at this solution, I shall also show how it sheds light on the appropriate methods for learning about frequency-dependent causal effects in science.

I

Giere offers a model of simple causal hypotheses for deterministic and stochastic systems. Both cases are covered in the following expression of Giere's *original model*:[6]

> (G): If C is a positive causal factor for effect E (in a population L) then the *mean value* of E,[7] if every member of population L

Philosophical Studies **49** (1986) 109–124.

were to have C, would exceed the mean value of E if no member of L were to have C.

On this view, "C is a positive causal factor for E" entails something about two *counterfactual populations*: the first, in which all members of the actual population were to have causal input C; the second, in which none were to have C. In particular, it entails that the mean value (e.g., average incidence) of effect E in the first would exceed the mean value of E in the second.

One of Sober's proposed counterinstances[8] to Giere's model concerns frequency-dependent selection in a population of butterflies of the species *Limenitis archippus*; which I abbreviate as *species L*. Species L is a cousin of the species of Monarch butterflies. Monarchs are distasteful to blue jays; so jays tend to avoid them. Members of species L are good-tasting to jays, but L's that mimic the appearance of the distasteful Monarchs may *deceive* jays into thinking they too are unpalatable. In this way, L's that can mimic are more likely to avoid being eaten by jays than L's that do not mimic. In the context of Sober's example, it is *given that* L's that can mimic enjoy this increased fitness just in case the ability to mimic is rare among members of L.

Consider Sober's own statement of the example and the problem to which it allegedly gives rise (with my emphasis on key phrases):

> Suppose that monarch butterflies are very common in a locale and that, within another species [i.e., L], there are both mimics and non-mimics. Let the mimics be comparatively rare. Now suppose that mimics produce 10 offspring per capita, whereas a nonmimic in the species [L] produces only 5. Selection favors the mimics, and it seems right to say that *mimicry causes differential reproductive success*. But the advantage provided by this sort of mimicry, we have noted, is frequency-dependent; so let's assume that, if the mimics were to become prevalent, their per capita offspring output would shrink to 5.
>
> *According to Giere's theory, to see whether mimicry causes individuals to have 10 offspring when mimicry is rare*, we must compare what would happen if 100% of the individuals were mimics with what would happen if 0% of the individuals were. At these two extremes, the per capita output of each individual is 5. *Giere's theory concludes that mimicry is not a causal factor in producing larger numbers of offspring when mimicry is rare* (p. 250)

The substance of Sober's argument may be sketched as follows:

(1) [Given] *When mimicry is rare* (in L) mimics produce an average of 10 offspring per capita, while nonmimics produce only 5 (i.e., "selection favors the mimics.").

(2) Premise (1) entails the truth of hypothesis *H*: Mimicry, when rare (in population L) causes increased fitness (i.e., expected increased reproduction rate).

(3) According to Giere's analysis (G), if *H* is true then the mean per capita offspring (average reproduction rate) among L's, if 100% of L's were mimics would exceed the mean number of offspring per capita if 0% of the L's were mimics.

(4) [Given] At these two extremes (i.e., 0% of L's are mimics, and 100% of L's are mimics) the mean number of offspring per capita is 5.

(5) [Since (4) negates the consequent of (3), it follows that] Giere's analysis concludes that mimicry does not cause differential reproductive success *when mimicry is rare*; i.e., it negates *H*. This contradicts (2).

Before giving my own response to Sober's argument, there is much to be learned by probing the strategies proposed by Giere and by Collier.[9]

II

In responding to Sober, Giere does not deny that the contradiction in (5) would follow were his original model (G) applied to frequency-dependent cases. But he denies that his model (G) was intended to apply to such cases, claiming that "Sober has not provided a 'counterinstance' to my analysis. He has merely exhibited examples to which the above simple model [i.e., (G)] does not apply" (387). Giere proposes an "enrichment" of his model (G) for cases of frequency-dependence. In particular, Giere suggests dealing with the factor of mimicry (in L) by treating the proportion of L's that are mimics, denoted by r, (and the ratio of Monarchs to the cousin species L, denoted by R) as part of the *environment*. Causal model (G), Giere claims, is then to be relativized to this environment (390):

> This we do by considering not the whole population of potential mimics [i.e., population L], but only a small *subpopulation*. Thus, changes in the ratio of mimics in the subpopulation leave the value of r approximately constant. For any nonextreme ratios, R and r, my interpretation of the causal hypothesis has the implication that the average fitness in the subpopulation will be greater if all members of the subpopulation are in fact mimics than if none are.

By generalizing this, Giere's enriched version of (G), which I denote by (G: Frequency-Dependence), asserts the following:

> (*G: Frequency-Dependence*): If (a) C is a positive (frequency-dependent) causal factor for effect E in population L, whose

effectiveness depends on the proportion of C'S in L (i.e., on the value of r), *and*
(b) L-sub is a subpopulation of L where changes in the proportion of C's in L-sub leave the proportion of C's in the full population (r) approximately constant, *then*
(c) the mean value of effect E in L-sub would be greater were all members of L-sub to have C, than if none were to have C.

The idea is for the subpopulation to be small enough (relative to the entire population) that the proportion of C's in the entire population is insignificantly altered, both when 100% of the subpopulation is imagined to have factor C (i.e., the first counterfactual population) and when 0% is imagined to have C (i.e., the second counterfactual population).

Notice that (with L finite) *no* subpopulation leaves the overall proportion of C's in the population unchanged throughout all changes in the proportion assigned C in the subpopulation.[10] At best, the proportion of C's in the overall population (i.e., r) remains approximately constant, thereby satisfying condition (b). Still, (a) and (b) fail to imply (c), contrary to (G: Frequency-Dependence), because not all subpopulations satisfying (a) and (b) need also satisfy (c). In some subpopulations the factor C may fail to cause any positive difference in the mean value of effect E. At the same time C may be a positive causal factor for E in the population at large. Giere[11] himself draws attention to such a possibility in discussing a non-frequency-dependent case:

> Taking the overall population as given, vaccination is a negative causal factor for polio – for that population. This, however, is compatible with their being some *other* population for which the reverse is true. And indeed, that other population might be a subgroup of the original.

Just such reversals, I claim, come back to haunt his later model (G: Frequency-Dependence).

(G: Frequency-Dependence) holds, at least approximately, if the causal effectiveness of C, in *any* subpopulation satisfying (b), is in the same *direction* as in the full population.[12] As long as there may be variability in the direction in which C effects E, the most (a) and (b) entail is that some *proportion of the possible subpopulations* would satisfy (c). Giere does not indicate now the subpopulation would be selected from population L, if one wanted to imagine carrying out this hypothetical experiment. By adding that it is to be

selected randomly (or by some other method of probability sampling), it would often be possible to calculate the proportion of subpopulations satisfying (c). This would yield a *probabilistic reconstrual* of (G: Frequency-Dependence) of the form: If (a) and (b), then with some specified *probability*, (c).[13]

What I claim for this probabilistic reconstrual is the virtue of rendering (G: Frequency-Dependence) true (despite varying causal directions); not that of providing an adequate model for frequency-dependent causation. For it does not. Satisfying part (b) often forces the size of the subpopulation to be so small that the probability of selecting a subpopulation having the difference in mean effect (required in (c)) is low. And, the better (b) is satisfied, the smaller the subpopulation must be; thus, decreasing the proportion of possible subpopulations where C's effect is close to what it would be in the full population. To get a feeling for how this occurs, note that satisfying (b) (r approximately constant) precludes subpopulations with many more members than the number of C's actually in L to begin with. For suppose the size of the subpopulation exceeds the number of C's in the actual population L. Then when this subpopulation is imagined to have 100% C's (i.e., the first counterfactual population) the overall proportion of C's must necessarily exceed the original proportion of C's in L.[14] Hence the rarer the property C in the full population, the smaller the subpopulation, and the less likely it will reflect C's effect in the full population – particularly when this effect is small. And even when this would not be a problem, I think the probabilistic reconstrual of (G: Frequency-Dependence) should be rejected by a Gierean as too weak.[15]

III

What other avenues might be taken to avoid Sober's criticism? One can attempt yet a different revision of (G) for frequency-dependent cases, or one can argue that applying Giere's original model (G) to such cases does not lead to the alleged contradiction in statement (5) of Sober's argument. And one way of following the latter of these two routes is simply to deny that the frequency-dependent causal claim *H* holds true for the entire population L (i.e., deny premise (2)). This appears to be the strategy recommended by Collier.[16] Collier denies that the failure of Giere's original analysis (G) to affirm certain frequency-dependent claims shows that (G) should be aban-

doned or revised. On the contrary, Collier claims, it shows the *value* of (G) "by ruling out certain populations from having the effect in question" (625). Collier presents a non-biological example (which he attributes to Sober) to illustrate how, in his view, (G) accomplishes this:

> There is no doubt that standing up is a positive causal factor in having a better view [at a baseball game]. We must be careful, however, to be clear concerning the population or populations of which this is true. If we take the relevant population to be the attendees at the game, then Giere's analysis fails. Giere must hold that for the population of attendees at the game, standing up is not a causal factor in having a better view... There are however, several smaller populations for which standing up *would* provide a better view. For these populations, standing up is a positive causal factor... (For example, we could tell only the people in the back row to stand up.) (p. 624)

According to Collier, the failure of Giere's analysis (G) to affirm hypothesis:

> *B-full*: Standing causes a better view in the full population of attendees

(since the view is no better with everyone standing than with everyone sitting) might just mean that the true hypothesis is the restricted one:

> *B-sub*: Standing causes a better view in a subpopulation of attendees (L-sub)

like those in the back row. And Giere's original analysis

> (G): If hypothesis B-sub is true, then members of L-sub would have a better view were they all to stand than if none were to stand

would hold for this instantiation. For were members of a suitable subpopulation to stand, (barring blindness, etc.) the proportion standing among the full population of attendees would still be sufficiently small to cause them to see better than if they were sitting. But this same fact also entails (for the baseball case) the truth of Giere's revised account:

> (G: Frequency-Dependence): If hypothesis *B-full* is true, members of L-sub would have a better view were they all to stand than if none were to stand (where L-sub satisfies condition (b)).

My general criticism of (G: Frequency-Dependence), then is avoided in this case, and Giere would *not* have to deny that standing causes a better view for the full population of attendees (i.e., *B-full*). Giere's reason for considering

the subpopulations in (G: Frequency-Dependence) altogether, was to affirm, where analysis (G) could not, frequency-dependent hypotheses about the *full* population in question. (And in the baseball case, he succeeds.) Giere's aim is thwarted however, by a "defense" that requires restricting the hypotheses *themselves* to these subpopulations (as Collier's restriction of *B-full* to *B-sub* requires). Such a restriction would prevent many causal claims from being affirmed for populations over which they are normally taken to hold.[17] More importantly, Sober's criticism of (G) arises in cases where it *does* seem correct to hold that C is a positive cause of E in the *full* population L. And it is this problem that should be addressed. Thus far I have shown that neither Giere's extension of (G), nor Collier's restriction provide fully adequate solutions. In what follows I shall propose a solution that does not require extending Giere's original model (G), nor does it require restricting the populations over which hypothesis *H* holds.

IV

My solution rests on arguing that Giere's original model (G) escapes the contradiction in Sober's statement (5). But how? Granting the "givens" of Sober's argument, the only premise left to question is (3); and premise (3), I shall argue has been the culprit all along. My grounds for denying the truth of (3) is that it misconstrues the two counterfactual populations to which Giere's model (G) should be understood to refer. Furthermore, the responses of Giere and Collier harbor the same miscontrual of which I shall accuse Sober.

Begin by considering the hypothesis at issue in Sober's mimicry example. From Sober's passage we gleaned that hypothesis *H* (stated in premise (2)) is this:

> *H*: Mimicry, when rare (in population L) is a causal factor in producing increased fitness (i.e., increased reproduction rate per capita).

Next, what is the causal factor C? From the statement of *H*, it is clear that C is "the ability to mimic when that property is rare among population L." More simply, C is "the ability to deceive jays *advantageously*" by dint of resembling Monarch butterflies; where it is stipulated that members of L possess this "advantageous deception" just in case mimicry is rare in popula-

tion L. Substituting C in Giere's model (G) gives as the first counterfactual population the one that would result were all members of L to possess C: the ability to deceive jays advantageously; and the second, the one resulting were no L's to possess this ability. According to Sober (in premise (3)), the first counterfactual population is the one in which "100% of the individuals [in L] were mimics," which for Sober means 100% of L's resemble the color pattern of the mimicked species *simultaenously*.[18] I will refer to this as the *simultaneous construal* of the counterfactual population. But is the population of L's with 100% mimics the same as the one in which 100% of the L's have the ability to deceive jays advantageously?

No it is not. In the former counterfactual population (all L's mimic simultaneously) *none* has the *rare* property of mimicry, so none can advantageously deceive jays. For, if all L's are mimics, mimicry is not a rare property in population L; no L's have the rare property of mimicry, hence no L's deceive jays. This clearly differs from the latter population where all of the L's have the ability to deceive jays to their advantage. But it is this latter population that the first counterfactual population in Giere's model (G) should be taken to refer. It follows that Sober's premise (3) misconstrues this first counterfactual population.[19]

Moreover, the interpretation of causal factor C is required to fluctuate throughout Sober's argument. The phrase "when mimicry *is rare*" (or the equivalent) appears in premises (1) and (2), but is slid out in (3) so as to get the failure of increased reproductive success at the extremes asserted in (4). It is slid back in again in (5) to obtain the contradiction. If the causal factor in *H* is consistently interpreted, then premise (3) is false. A correct instantiation of Giere's model (G) is this:

(3)* According to Giere's model (G), if *H* is true, then the mean number of offspring, were 100% of L's able to deceive (having the rare property of mimicry), would exceed the mean number of offspring were 0% able to.

But how does replacing (3) with (3)* enable Giere's model (G) to avoid the contradiction in (5)? That is, how does it help avoid the absence of increased reproductive success in the two counterfactual populations – and so avoid conflicting with what (G) asserts hypothesis *H* entails? The answer must lie in showing the importance of distinguishing "all L's are mimics" from "all L's have the rare property of mimicry." The idea of all L's possessing the rare

property of mimicry may appear self-contradictory unless it just means all L's are mimics. But then (3)* is identical to (3), and the contradiction in (5) follows soundly. Just this identity is accepted in both Giere's[20] and Collier's responses. However, while it is contradictory for all L's to have a rare property *at the same time*, it is not contradictory for all L's to have a rare property. The solution I propose comes down to providing an interpretation of the latter. This, I think, is most effectively done by way of frequency-dependent cases *other* than those from evolutionary theory.[21]

V

A large class of familiar non-biological examples of frequency-dependent causation arise from properties involving exclusivity or superiority; properties involving comparisons with others in the population: An example is

> *J*: Being awarded a Nobel Prize (or other rare honor) is a causal factor in obtaining a salary increase (or other measure of professional success) in the population of U.S. scientists.

The first thing cases like *J* illustrate is that far from being exceptional or exotic, frequency-dependent (or, more generally, *context-dependent*) causal effects are commonplace. That such cases do not obviously alert one to this frequency-dependence further underscores the need for a single causal model that is applicable to both frequency-dependent and non-frequency-dependent claims. Secondly, it is clear that hypothesis *J* is *not* asserting that were the population of U.S. scientists to all (simultaneously) possess Nobel Prizes, that *the Prize* would afford them an average salary increase in excess of that were none to. Were all U.S. scientists to win a Nobel Prize, then *none* would have whatever advantage accrues (to salary) by dint of possessing such a *superior* record of achievement. Of course salary increases may differ in populations with 100% and 0% having the Prize, respectively. But this difference is quite distinct from that caused by winning the Prize in the U.S. (where it is comparatively rare). Giere[22] stresses that his model is intended to capture actual causal effects in the population and environment specified. The simultaneous construal of the counterfactual populations in (G) frustrates this. Analogously, the ability to learn about a context-dependent effect is frustrated when all experimental subjects are "treated" simultaneously; the effect of interest is altered. Examples of such poorly run studies are familiar.

One that is of special interest here satisfies each of the "givens" of Sober's mimicry example.

Researchers often study the effects of certain substances on mice by measuring changes in their blood pressure. Such research is vitiated if it confuses effects attributable to this substance with those due to the process of making experimental measurements. The blood pressure of mice, for example, is markedly effected by the changing composition of their cage, as mice are removed one by one to be examined.[23] This (frequency-dependent) effect may be stated as follows:

> K: Being taken from a cage, when most mice have already been removed, causes an increase in blood pressure (in a population of laboratory mice).

The failure of many researchers to recognize this effect is related to its frequency-dependency. Removing mice may be thought to have no effect on pressure because there is practically no difference in mean pressure when all mice are *simultaneously* outside the cage, as when all are normally in their cage. That is, by observing, in effect, the simultaneous construal of the counterfactual populations of mice one fails to affirm K, much as the simultaneous construal of model (G) fails to affirm the mimicry hypothesis H. Therefore, it is not unreasonable to expect a correct experimental strategy for detecting the effect in K to shed light on the correct way of applying (G) to hypothesis H.

VI

Let the population be the 500 mice in a given research lab; and suppose each cage forms a colony of 25 mice. The causal factor C in hypothesis K is "being taken from a cage after at least 13 of the mice have been removed." More simply, let C be "being traumatized" where mice are taken to have this property by dint of being stressed by the changing composition of the cage, without further unpacking its causes. Instantiating into Giere's model, I claim, yields:

(3)* [for K]: According to Giere's model (G), if K is true, then the mean blood pressure, were 100% of the mice "traumatized" (by being examined after the majority having already been removed), would exceed the mean pressure if 0% were.

But how does one observe 100% of the mice "being removed after most have already been removed (and thereby traumatized)"? Since each cage (20 in all) constitutes a mice colony, only 240 of the 500 mice could simultaneously be traumatized (in their respective colonies). Yet it is easy to see that each of the 500 mice can have its turn being traumatized. What makes it easy in this case is that here the causal factor can be assigned, and later *undone*. A mouse who is traumatized in one trial, can (after being given time to return to its normal pressure) serve as an untraumatized mouse in the next, and conversely.

Each of the mice can have its pressure taken at each of the 12 positions of traumatization (i.e., the 13th, the 14th...the 25th to be removed); and the arithmetic mean of these 12 pressures recorded. After this is done for each of the 500 mice, the 500 means are themselves averaged. This yields the mean blood pressure when each member of the population of lab mice has the property of traumatization. *This, then, is what it means to observe the mean effect in the first counterfactual population in model (G) (i.e.*, all mice traumatized). Similarly, 500 average pressures can be obtained for each mouse over the 13 positions of non-traumatization.[24] Averaging these yields the mean blood pressure in the population where non of the mice has property C; that is, in the second counterfactual population in (G). If hypothesis *K* is true, the mean pressure in the first counterfactual population (thus obtained) would exceed the mean pressure in the second. Thus (3)* (instantiated for *K*) is true, and the effect is detected by the trials described.

Although the mice example is exceptional in that the counterfactual populations can actually (quite accurately) be generated by a series of experiments, the same interpretation of the counterfactual populations, I maintain, is available when they are only hypothetical.

For each member of the actual (finite) population, the causal factor C of interest would produce some value, or (for a stochastic case) some mean value of the effect E of interest. The mean of these values over all members of the population would give the mean value of the effect in the first counterfactual population. Similarly, the absence of C (defined relative to the causal process of interest) would produce some (mean) value of effect E for each member, and averaging over the entire population yields the mean value of E in the second counterfactual population in model (G). *Whether or not C is frequency-dependent*, if C is a cause of E in population L, it follows that the mean value of E in the first counterfactual population would exceed that of

the second – which is precisely what (on the present reading) Giere's model (G) asserts.[25]

The promised solution for the mimicry example drops out nicely from my general construal of (G). A brief sketch should suffice. In a population L of n butterflies there are $n/2 - 1$ "rare slots" that each butterfly can be envisioned to occupy and still be in a context where it has the property "mimicry when rare": each (in turn) could be a mimic when 49% are, when 48% are, and so on. Since the average offspring is likely to vary even at a given "rare mimic slot", here one is interested in the mean number of offspring each butterfly would produce at each of the possible permutations of the occupants of the remaining slots. The total average offspring (over the n L's, over the possible permutations) is what is meant by the "mean offspring per capita were each of the n L's able to deceive jays (i.e., were each to have the rate property of mimicry)". By the [given] (in Sober's premise (1)) the value of this mean would be 10. Analogously, the mean number of offspring in the second counterfactual population refers to the mean value that would result were each of the L's to independently occupy the slot designated "not a rare mimic". Whether this is taken to mean it is not a mimic, or that it is a mimic but it is not rare in L, the mean value would still be 5, by the "givens" of Sober's example. Therefore, premise (4) is false; the differential reproduction rate that the truth of *H* (in (3)*) entails, *is* satisfied by the two counterfactual populations stipulated by a correct construal of Giere's original model (G).

The understanding of frequency-dependent causal effects that emerges from my construal of the counterfactual populations in Giere's model (G) holds regardless of the hypothetical nature these will typically have. Most importantly, using this understanding indicates the sort of actual experiments that *would* detect the effect in question with some degree of accuracy (as with the mice example *K*), and which would not (e.g., simultaneously treating the subjects with a context-dependent property). Conversely, the value attached to certain experimental principles (and the disvalue to others) receives an explanation. The causal claims (of the sort considered here) are understood as assertions about a certain set of counterfactual populations, so an experimental strategy that gives information about these counterfactual populations is obviously valuable for testing and arriving at such claims. Because the specific construal of these counterfactual populations put forth here entails specific criteria for assessing the value (at least of some) potential strategies, it can be both utilized and tested by examining these assessments.

For instance, on my construal, (G) implies that the only thing that alters the mean value of the effect in the two counterfactual populations is the presence or absence of causal factor C. This explains the importance attached to experimentally creating "treated" and "control" groups that differ (with regards to the effect E) solely as a result of the presence or absence of the treatment (i.e., factor C), respectively. Why? Because in this way it is possible to observe, in effect, a sample from the two counterfactual populations of interest. And the rationale of certain methods of taking these samples[26] is that they enable inferences about these counterfactual populations to be sustained.

To sum up: I have proposed a construal of Giere's causal model (G) that accomplishes the primary task being addressed, namely, to disarm the criticism Sober raises. Emerging from this construal is an understanding of frequency-dependent hypotheses which accords with, as well as sheds light on, appropriate methodological strategies for investigating such hypotheses. According to Giere, what one seeks to show in support of empirical models of causal inquiry "is that they are useful in understanding a variety of causal hypotheses and methodological strategies" (385). On this criterion, my construal of Giere's causal model (G) has been shown to be of value. It stands ready to be applied to other problems and puzzles involving context-dependent causality,[27] to test if it flourishes or flounders.

ACKNOWLEDGEMENT

I am deeply grateful to Ronald Giere for the encouragement that led to this project, and for many valuable comments on earlier drafts. I would also like to thank Norman Gilinsky for numerous helpful remarks. A short version of this paper was presented at the Eastern APA meetings in December, 1984.

NOTES

[1] Those of recent philosophical interest stem largely from issues concerning the dynamics of natural selection, where frequency-dependency is seen to be prevalent. The best example is the problem of distinguishing individual from group selection, see fn 27. See especially Richard Lewontin, *The Genetic Basis of Evolutionary Change* (New York: Columbia, 1974), pp. 256–260.

[2] 'Frequency-dependent causation', *Journal of Philosophy* LXXIX, 5 (May 1982) pp. 247–253. All references to Sober are to this article.

[3] Giere develops this causal model in 'Causal systems and statistical hypotheses', L. J. Cohen and M. B. Hesse, eds., *Applications of Inductive Logic* (New York: Oxford,

1980), pp. 251–270, and in *Understanding Scientific Reasoning* (New York: Holt, Rinehart & Winston, 1979).

[4] 'Causal models with frequency dependence', *Journal of Philosophy* LXXXI, 7 (July 1984): 384–391. All citations from Giere are to this article, except as noted in Note 11.

[5] 'Frequency-dependent causation: A Defense of Giere', *Philosophy of Science*, L (1983): 618–625. All references to Collier are to this article.

[6] Giere's original model also asserts the *converse* of the necessary condition stated in (G); however, it is to the entailment in (G) that the arguments considered here refer.

[7] It should be noted that I am allowing E to be a *qualitative* as well as a *quantitative* property. When the effect is qualitative, such as "overweight" or "not overweight", the mean of E refers to the *incidence rate* of E over members of L. With a quantitative effect, such as an individual's weight (or average weight), the mean of E is the *arithmetic* average of this quantity; the mean weight over L (or the mean of the average weights, one for each member of L). In claiming that Giere's model accommodates both types of effects, I am claiming it is more robust than even its author had hoped.

[8] While I shall not discuss the other examples Sober considers, the solution I propose may be applied to those cases as well.

[9] My sketch of Sober's argument already reflects the clarification afforded by both Giere's and Collier's responses. In particular, it is unclear whether Sober is referring to the rarity of mimicry within L alone, or its rarity over L plus the Monarch populations. After running through both possibilities, Giere (in a formal treatment) and Collier (less formally) show that only on the *first* construal of the population is Giere's model open to Sober's criticism.

[10] Note that r denotes the proportion of C's in L; a quantity that *varies* as does the number of L's with C. Let r_o be the *fixed* proportion of L's that have C in the actual population of interest. My claim is that if the number of C's in a subpopulation from finite population L is imagined to vary, then r will not remain equal to r_o throughout. Let N denote the number of individuals in population L (so, the number with C equals r_oN). The subpopulation from L consists of k of the N members of L; let k_c and $k_{\bar{c}}$ denote the number taken from those having C, and those not having C, respectively. The general equations for calculating r 1) under the assumption that all k members of the subpopulation have C (i.e., the first counterfactual population), and 2) under the assumption that none of the k have C (i.e., the second counterfactual population) are as follows: 1) $r = r_o + k_{\bar{c}}/N$, and 2) $r = r_o - k_c/N$. It is clearly impossible to have $r = r_o$ in *both* 1) and 2). [$r = r_o$ in 1) if all k members of the subpopulation are taken from L's that actually have C; for then $k_{\bar{c}} = 0$. But then r is less than r_o (by k/N) in 2).]

[11] Reply to Comments on 'Causal systems and statistical hypotheses', *op. cit.*, p. 288.

[12] For it to provide an *exact solution*, the causal effectiveness of C in any subpopulation satisfying (b) would have to be identical to (and not just in the same direction as) its effectiveness in the full population.

[13] Someone may ask whether the probability that I claim needs to precede claim (c) (for (G: Frequency-Dependence) to be generally true) is not already taken care of in the stochastic nature of the effect in (c). To assert, C increases reproductive success in L-sub, for example, is understood (by Sober and others) to refer, not to the reproductive output L-sub will actually produce; but to an increase in the probability (or mean value) of reproductive success C confers upon L sub. [Some express this as its increased *propensity* to reproduce.] That is true; but the additional probability to which I am referring is the probability of selecting an L-sub in which C brings about this probabilistic effect.

The three distinctions here are analogous to those drawn in statistics between (i) a sample average, (ii) the probability (or mean value) of the possible sample averages, and (iii) the probability of selecting a population with a given probability of producing various sample averages. (It should be noted that (iii) describes a legitimate probability experiment for a frequentist.)

[14] This is easily seen by referring to equation 1) in Note 10. If the size of the subpopulation (i.e., k) exceeds the number of C's in the actual population L, then k exceeds $r_0 N$. Then it is impossible for all k to be taken from those L's that have C in the actual population (as there are only $r_0 N$ of these.) So, $k_{\bar{C}}$ must be greater than 0, resulting in a value of r [given by the righthand side of equation 1)] that exceeds r_0.

This argument holds regardless of the composition of those not in the subpopulation.

[15] For it only enables the iterated statement spelled out in footnote 13. ("With some probability a subpopulation would show an increased mean effect rate...") Substituting the propensity notion of probability Giere favors weakens the statement to one of iterated propensities, which would not be very useful as a guide for testing causal claims. Moreover, one would rarely know whether to apply (G) or (G: Frequency-Dependence).

[16] Collier does not use this strategy, however, in dealing with the mimicry case. He claims, *op. cit.*, p. 622, that

> If 100% of the population were mimics, the jays would quickly reduce the population to the equilibrium level of the realistic case. If 100% of the population were non-mimics, the jays would reduce the population below this equilibrium level. Therefore, mimicry is a positive causal factor in reproductive success.

But this is just to deny one of the "givens" of Sober's example (i.e., premise (4)); namely, that at these two extremes the mean reproductive rates are identical (i.e., 5).

[17] Collier acknowledges that his view "may disagree with common assumptions and, consequently with common usage..." (625).

[18] In contrast, some biologists reserve the notion mimicry for the possession of an advantageous resemblance, which in this example requires the resemblance to be rare in L. However, my point is not that the rarity of the resemblance should be "packed into" the notion of mimicry, but that it should be "packed into" the factor C asserted to do the causal work in claim *H*.

[19] The second counterfactual population (no L's have C) would result if either the majority of L's were to be mimics, or if there were no mimics (Sober's construal). Since in either case the mean reproduction rate is given to be 5, either construal leaves the issue of the present discussion unaffected.

[20] It is my understanding, on the basis of private conversation, that Giere would no longer accept this identity, and that he generally endorses the solution put forward here.

[21] My strategy takes seriously Giere's insistence that proposed *empirical models* of causality, in contrast to proposed *explications* of the meaning of causality, are to be tested by applying them to actual scientific research. In this way I think it can be shown that the construal of (G), proposed in (3), in contrast to (3)*, leads to clearly erroneous experiments being endorsed. This calls for an application in which the counterfactual populations can be constructed, at least approximately; unlike those in testing *H*.

[22] See, for example, 'Causal systems and statistical hypotheses', *op cit.*, p. 265.

[23] See H. Magalhaes, ed., *Environmental Variables in Animal Experimentation* (Pennsylvania: Bucknell University Press, 1974) for an excellent discussion of this.

[24] This is not intended to describe any actual experiment that may have been carried out in affirming *K*. However, any additional variables required in investigating *K* could be handled by simply calculating a more complex set of averages.

[25] The converse of (G) holds as well. It should be emphasized that although what would be required (either on paper or in the lab) to obtain the means of the two counterfactual populations will vary from case to case, the measures to which these means refer, I claim, remain constant.

[26] I have in mind here methods of randomization and stratification. While these are given an important role in distinguishing spurious from genuine causal effects in the received theory of experimental design, nearly all philosophical discussions of statistics

deny that this role is justified. As Lindley (paraphrasing Harold Jeffreys) asks in *Bayesian Statistics; A Review* (Philadelphia: Society for Industrial and Applied Mathematics, 1972) p. 21, "How can the fact that a different result might have been obtained, but did not, influence you once the data is on view?" On Bayesian and Likelihood Principles, if strictly adhered to, the answer must be "It cannot." In contrast, it can matter a great deal if the goal is to compare a sample treated with C to (hypothetical) results that *would have* occurred were none to have been so treated. Standard frequentist (or Orthodox) methods of data generation and analysis, which are directed toward this goal, therefore receive an important justification from the counterfactual model of causal claims sketched here. I discuss this further in 'An objective theory of statistical testing', *Synthese* 57 (1983), pp. 297–340.

[27] For example, since the cause of differential group fitness is often a frequency-dependent factor, a suitable model for frequency-dependent causation should be expected to be relevant to the problem of characterizing group (as distinct from individual) selection. The counterfactual model in my construal of (G), I think, lends concreteness to the notion of causal forces proposed by Sober, for example in 'Holism, individualism, and the units of selection', in P. D. Asquith and R. N. Giere, eds., *PSA 1980*, vol. 2 (East Lansing, Mich.: Philosophy of Science Association, 1980): 93–121. It also seems to have the consequence that certain types of cases of frequency-dependent selection (as characterized by model (G)) appear as cases of group selection (on Sober's model of group selection). The further consequences of this, however, cannot be explored here.

Department of Philosophy,
Virginia Polytechnic Institute and State University,
Blacksburg, VA 24061,
U.S.A.