Elliott Sober

Two Uses of Unification

1. Introduction

Carl Hempel[1] set the tone for subsequent philosophical work on scientific explanation by resolutely locating the problem he wanted to address *outside* of epistemology. "Hempel's problem," as I will call it, was not to say what counts as evidence that X is the explanation of Y. Rather, the question was what it means for X to explain Y. Hempel's theory of explanation and its successors don't tell you what to believe; instead, they tell you which of your beliefs (if any) can be said to explain a given target proposition.

At the same time that Hempel's problem developed outside of epistemology, another project gained a footing within that discipline, one that also laid claim to the concept of explanation. This is the Peircean idea of abduction – of "inference to the best explanation" – which was put forward as a solution to the problem of understanding how nondeductive inference works.[2] To evaluate competing theories in the light of a body of observations, we are supposed to ask which of those theories is most explanatory of the observations. Or, more generally, we need to ascertain which of the competing theories, if added to our stock of beliefs, would lead to a total system of beliefs with the most explanatory power.[3] "Peirce's problem," as I will call it, is the epistemological problem of specifying the rules that govern inference to the best explanation.

In addition to the fact that these two discussions of the concept of explanation are located on opposite sides of the fence that separates epistemology from the rest of philosophy, the two projects differ in further respects. Hempel's question is *yes* or *no*. Does X explain Y, or not? However, for the Peircean proponent of abduction, the fundamental question is comparative – does X explain Y better than Z does? An additional contrast derives from the fact that Hempel wanted to explicate the concept of an *ideally complete scientific explanation*. Although he recognized that the discoveries of science can be explanatory without satisfying his requirements – he termed these fragments *explanation sketches* – he once again set the problem that most of the subsequent literature addressed. The task is not to understand what allows one proposition to *help explain* another;[4] rather, the goal is to understand what it means for one proposition to *provide a complete explanation* of the other. Students of abductive inference, on the other hand,

F. Stadler (ed.), The Vienna Circle and

usually take it as a given that the competing hypotheses one needs to assess rarely provide complete explanations in themselves.

I have distinguished Hempel's problem and Peirce's problem because I want in this essay to describe how the concept of unification pertains to each. I'll do this by evaluating two principles that address the problem of how one should explain two "similar" observations. Giving a full description of what similarity amounts to in this context would be difficult, but for our purposes we can restrict ourselves to the following sufficient condition: two observational propositions are similar when they apply the same predicates to different individuals.[5] The observation that a is green and the observation that b is green are similar in this sense, as are the observation that c runs faster than d and the observation that e runs faster than f. Here, then, are two principles about how the similar observations O_1 and O_2 should be explained:

(P) If you want to judge which of the hypotheses U and D is more plausible in the light of the observations O_1 and O_2, the fact that U provides a unified explanation of O_1 and O_2, while D gives a disunified account, is evidence in favor of U.

(H) If U and D are both true, and U provides a unified description of phenomena O_1 and O_2, while D provides a disunified description of those phenomena, then U may explain O_1 and O_2, but D does not.

Principle (P) is an epistemological claim that addresses Peirce's problem; it says that unification is confirmatory. Principle (P) does not rule out the possibility that the only true explanations of the observations are disunified; rather, it says that in deciding what to believe, it is a point in favor of U that it provides a unified account. Principle (H), on the other hand, is not epistemological; it addresses Hempel's problem, asserting that similar phenomena must receive a unified explanation. The point of (P) is to help you figure out which of U and D to believe; in contrast, the point of (H) is to help you decide which of the propositions you already believe provides an explanation of two target observations.

My thesis will be that these two uses of the concept of unification should receive very different assessments. I'll argue that unification has no objective connection with Hempel's problem. Whether we embrace or reject (H) depends on what sort of explanation we seek; there are perfectly legitimate scientific goals that make unified explanation a virtue while other equally worthy scientific goals have the opposite effect. I'll argue that the status of (P) is different; there is an objective reason why unification is relevant to Peirce's problem.

2. ANTIREDUCTIONISM AND HEMPEL'S PROBLEM

Hilary Putnam has had a decisive influence on how philosophers think about reductionism. In fact, he has had that influence *twice*. The first time was his essay with Paul Oppenheim[6] in which he defended reductionism as an empirical thesis about the relationships that connect the different sciences. The second occasion was a pair of articles in which he argued against his earlier position.[7] The anti-reductionism that this more recent Putnam-time-slice espoused continues to be a powerful presence in philosophy; anti-reductionism is now more or less the "received view" among philosophers of science and philosophers of mind, and the standard argument for this standard view derives from Putnam. Putnam's argument appeals to the value of unification in explanation, and criticizes reductionism for failing to deliver the requisite sort of unification. It is ironic that he and Oppenheim earlier had seen reductionism as the vehicle for achieving unification among the sciences; reductionism changes from hero to villain in Putnam's thinking, but the value of unity remains constant.

The reductionism that Putnam aims to refute is micro-reductionism. This is the idea that the properties of wholes can be explained solely in terms of the properties (including relational properties) of parts. The social sciences reduce to psychology, psychology to biology, biology to chemistry, and chemistry to physics, or so the micro-reductionist claims, because the objects described at higher levels turn out to be composed exclusively of the objects described at lower levels. Putnam did not gainsay the thesis that societies are made of individuals, that brains are made of cells, and so on. Rather, he rejected the reductionist's claim about explanation – that every event explainable by a higher-level science can also be explained by a lower-level science.

Putnam challenges this reductionist thesis about explanation by way of a delightfully simple example. He asks the reader to consider a plank of wood that has two holes – a square hole that is one inch on a side and a circular hole that is one inch in diameter. He points out that a square peg that is 15/16th of an inch on a side will fit through the square hole, but not through the round one. Why is this so? Putnam claims that the right explanation for this fact is given by the macro-properties just cited. It is the geometric shapes and dimensions of the peg and the holes that do the explaining. Putnam then claims that it would be entirely unexplanatory to cite the positions and other properties of the molecules that compose the peg and the board. This welter of micro-details would miss the wood for the trees. The micro-story, says Putnam, is either a non-explanation, or a vastly inferior explanation. Since the macro-story explains something that the micro-story cannot explain (or cannot explain nearly as well), reductionism is false.

Putnam claims that there are objective standards that distinguish good explanations from bad ones (and from non-explanations); it is on this objective basis that he praises the macro-story and blames the micro-story. Good explanations should advert to details that are explanatorily relevant, and omit details that are explanatorily irrelevant, where "explanatory relevance" is supposed to be an objective notion, not just a function of our interests. Putnam says that one feature of explanatory relevance is the idea that good explanations are *general* – they apply not just to the specific system under study, but to others. This is why the macro-story is superior to the micro-story; there are lots of other peg and board systems that have the dimensions specified, while there are relatively few (perhaps none) that have the exact molecular configuration that the micro-story describes.

The imperative to find explanations that are *general* is equivalent to the imperative to find explanations that are *unified*. To see this, consider *two* peg and board systems, which differ in their micro-configurations but have in common the macro-dimensions that Putnam mentions. A unified explanation of why the pegs in both cases fit through one hole but not the other will appeal to the macro-properties that the two systems have in common. In contrast, the micro-details do not provide a unified explanation; from the micro point of view, the two systems are different. This is where principle (H) makes its appearance. In the search for explanation, macro-accounts are preferable to micro-accounts because they are more general, and generality is just unification by another name.

Putnam never says much about what he means by explanatory relevance. In the example he describes, it may seem obvious that a detailed description of all the many molecules does not enhance our understanding. But surely it isn't always true that micro-details are explanatorily irrelevant. For example, suppose we observe that

Smith and Jones both have lung cancer.

A unified explanation of this pair of observations might be that

(U) Smith and Jones both smoked cigarettes.

However, suppose that Smith and Jones smoked different brands of cigarette, and that these brands contain different carcinogenic ingredients. This allows us to formulate the following disunified explanation:

(D) Smith smoked brand X, which contained carcinogenic x-particles, while Jones smoked brand Z, which contained carcinogenic z-particles.

Is it really so obvious that (U) is the uniquely correct explanation – that (D) is either no explanation at all or a terrible explanation? On the contrary, I think it is obvious that (D) may have considerable explanatory interest. Perhaps the different carcinogenic ingredients give rise to different types of tumor, and maybe these different tumors are best treated by different therapies.

Putnam values generality but accords no explanatory importance to detail. It would be equally one-sided to value detail only, with no recognition of the virtue of generality. Rather, I suggest that which type of explanation is better depends on your interests; (U) is better if your goal is generality and unification, but (D) is better if your goal is detail. Both goals are important in science. I suggest that the same is true of Putnam's peg and board example. Perhaps the molecular details don't interest Putnam or most of the rest of us, but that doesn't mean that they can't have any scientific interest at all.

Is it possible that an account of explanatory relevance might be offered that shows that there is an objective basis for ruling out the micro-explanation in the case of Putnam's peg and board, but allows that the disunified explanation (D) is explanatorily relevant in the case of Smith and Jones? I am skeptical. The cases appear to be structurally the same, as the following causal diagrams indicate:

micro-constitution of the peg and board at time $t_1 \rightarrow$ macro-properties of the peg and board at $t_1 \Rightarrow$ the peg goes through the square hole but not the round one at t_2.

micro-constitution of the smoke and of Jones' lungs at $t_1 \rightarrow$ Jones inhales cigarette smoke at $t_1 \Rightarrow$ Jones gets lung cancer at t_2.

Here the relation of synchronic determination that connects lower-level to upper-level events is represented by the single arrow "$\rightarrow$" and the diachronic relation of causation is represented by the double arrow "$\Rightarrow$".

This is not to deny that there is (or may be) a difference between the two examples. Let us suppose that the macro-configuration of the board and peg *deterministically ensures* that the peg will fit through one hole but not the other, but that Smith's smoking merely *confers a probability* (that is less than unity) on his getting lung cancer. This difference entails a further difference – that the macro-configuration *screens-off* the micro-configuration from the effect in the case of the board and peg, but need not do so in the case of smoking and cancer. That is, the first probabilistic equality will be true, but the second need not be:

Pr(the peg fits through one hole but not the other | the macro-dimensions of the board and peg) =
Pr(the peg fits through one hole but not the other | the macro-dimensions of the board and peg & the micro-configuration of the molecules in the board and peg) = 1.0.

Pr(Smith gets lung cancer | Smith smokes cigarettes) =
Pr(Smith gets lung cancer | Smith smokes cigarettes & Smith inhales x-particles).

The screening-off relation will fail in the smoking example, if there are different types of cigarette that differ in how carcinogenic their ingredients are. The reason that screening-off cannot fail in the deterministic case, but can do so in the probabilistic case, is that in both cases the probability of the effect, given the macro-configuration, is an *average*. When there are n micro-configurations (i = 1,2,...,n) that have positive probability, given the macro-configuration described,

$$\Pr(\text{Effect} \mid \text{Macro}) = \sum_i \Pr[\text{Effect} \mid \text{Macro} \ \& \ \text{Micro}(i)]\Pr[\text{Micro}(i) \mid \text{Macro}].$$

Notice that if the left-hand side of this equality has a value of unity, each of the n probabilities of the form Pr[Effect | Macro & Micro(i)] must also have a value of unity; if a set of probabilities averages to one, all must have a value of one. However, if Pr(Effect | Macro) < 1.0, Pr[Effect | Macro & Micro(i)] and Pr[Effect | Macro & Micro(j)] (where i ≠ j) may have the same value, but they need not.

Even if screening-off holds in the case of Putnam's example, but is violated in the smoking example, it would be a mistake to embrace screening-off as the key to understanding explanatory relevance. For one thing, in a Markov chain that links a distal cause (C_d) to a proximate (C_p) and this to an effect (E), the proximate cause screens off the distal cause from its effect, but this should not lead us to conclude that the distal cause is not explanatorily relevant to the effect.[8] I propose that precisely the same point holds when we consider the relationship between the macro-cause (C_p), its supervenience base (S), and the effect term (E):

$$\begin{array}{ccccc} C_d & \Rightarrow & C_p & \Rightarrow & E \\ & & \uparrow & & \\ & & S & & \end{array}$$

Even if the macro cause (C_p) screens off its supervenience base (S) from the macro-effect (E), it does not follow that S is explanatorily irrelevant to E.[9] Suppose the different brands of cigarette are equally carcinogenic (and so screening-off holds); it still may be explanatorily relevant to say that Smith's lung cancer traces back to his inhaling x-particles; going beyond the weaker claim that he "smoked cigarettes" is not irrelevant babble.

Although Putnam does not provide an account of what the relation of explanatory relevance amounts to, there is a natural candidate that is worth mentioning. If an explanation must describe the *causes* of the event one wishes to explain,[10] and if the difference between cause and non-cause is objective, then some proposed explanations are objectively false. It is false to say that the peg fit through one hole but not the other because the board was painted red, and it is equally false to say that Smith got lung cancer because his fingers had nicotine stains on them. However, I know of no plausible reading of the concept of cause that entails that the board's micro-constitution was like its color, or that Jones' inhaling x-particles was like his having yellow fingers. Putnam, like Jaegwon Kim, thinks one has to choose between the micro- and macro-accounts.[11] Kim argues that the causal action is to be found solely at the micro-level; Putnam contends, as we have seen, that the explanatory action is at the macro-level alone. In fact, both of these monolithic positions are mistaken; there is no need to choose. Both micro and macro provide true descriptions of the causal facts, and both thereby provide true causal explanations.[12] The two types of description have complementary virtues; one offers detail while the other offers generality and unification.[13]

3. MODEL SELECTION AND PEIRCE'S PROBLEM

"Model selection" is the name for a family of statistical inference problems in which the hypotheses considered contain adjustable parameters (a concept I'll explain below). There are a variety of criteria now on the market for choosing among such models. Since my goal here is to show how unification is relevant to Peirce's problem, I won't attempt to provide a complete survey and critical assessment of the options. It will suffice to describe briefly one of the most influential and interesting approaches – that due to Hirotugu Akaike;[14] like other proposals, it applies to a context in which unification is epistemically significant and thereby provides an explication of principle (P).

As an example, let's consider the X- and Z-cigarettes that Smith and Jones smoked. Suppose you want to infer how the number of cigarettes smoked affects an individual's chance of getting lung cancer. Although there are many different mathematical models that might be considered, for the sake of simplicity I'll restrict my attention to linear models; these models maintain that each additional X-cigarette smoked increases one's chance of getting lung cancer by a constant amount, and that the same is true of each additional Z-cigarette. However, I will not assume that an X-cigarette and a Z-cigarette are equally risky. Deciding whether this is so is precisely the model selection problem I want to describe.

Suppose we examine a large number of smokers and ascertain, of each individual, how many X- and Z-cigarettes he or she smoked, and also whether he or she has lung cancer. With enough individuals in our sample, we can describe the frequency of lung cancer within different categories of smokers, ranging from those who have smoked only a few cigarettes in their lifetimes to those who have smoked many. We also may be able to compare the incidence of lung cancer among individuals who have smoked the same number of cigarettes, but who have smoked different brands. There will be an upward trend in this data set – the more cigarettes smoked, the higher the frequency of lung cancer.

There are two linear models that I want to consider. The first says that it is merely the total number of cigarettes smoked that determines one's risk of lung cancer:

(U) Pr(S will get lung cancer | S has smoked n X-cigarettes and m Z-cigarettes) $= \alpha_1 + (n+m)\alpha_2$.

This model is unified because it assumes that the effect of smoking an additional X-cigarette is precisely the same as the effect of smoking an additional Z-cigarette. Smoking is smoking, this model says; it doesn't matter which brand you smoke, as far as the probability of getting lung cancer is concerned. An alternative model, also linear in form, allows for the possibility that the different brands may pose different risks:

(D) Pr(S will get lung cancer | S has smoked n X-cigarettes and m Z-cigarettes) $= \beta_1 + n\beta_2 + m\beta_3$.

This model is disunified, in that it provides separate representations of the effects of the two types of cigarette. It won't matter to our subsequent discussion if we assume that β_2 and β_3 must have different values; the point is that (D) does not constrain them to be the same. Notice that the unified model is simpler, in that it contains two adjustable parameters (α_1 and α_2) whereas the disunified model contains three (β_1 , β_2 , and β_3). Notice also that each model can be viewed as an infinite disjunction, each disjunct of which is formed by assigning values to the model's adjustable parameters.

How are these two models related to the data we have gathered? In each case, we can find the best-fitting member of the model by finding the maximum likelihood values of the adjustable parameters – i.e., the values of these parameters that maximize the probability of the data. These likeliest members of the two models, L(U) and L(D), contain no adjust*able* parameters; rather, their values have been adjust*ed*. It is pretty much inevitable that L(D) will fit the data better than L(U).[15] Does this mean that we should automatically prefer the more complex model? All of the model selection criteria now used in science answer this question with an emphatic *no*. Although the reasons they give are different, their bottom line is the same – there can be good reasons for preferring (U) over (D) even though L(U) fits the data worse than L(D) does.

Akaike's approach to this problem is especially instructive. Akaike suggested that the problem of model selection be conceived in terms of the goal of finding models that will be predictively accurate. This conception of the goal of model selection is what I mean by Akaike's "framework." Akaike also proposed a means for achieving that goal; he proved a theorem that describes how one can obtain an unbiased estimate of a model's predictive accuracy. This theorem is the basis for what has come to be called the Akaike information criterion (AIC). This separation of Akaike's framework from his criterion is important; there may be circumstances in which AIC is *not* the best criterion to use in model selection, even granting the goal of maximizing predictive accuracy.

Akaike's idea of predictive accuracy has to be understood in terms of a two-step process. Models that contain adjustable parameters make predictions in the following sense: first one draws a set of data from the true underlying distribution and uses those data to estimate the values of the model's parameters (by maximum likelihood estimation). One then uses that fitted model to predict a new data set drawn from the same distribution. In terms of our previous notation, we use a model M to make a prediction about new data by first using the old data to find L(M) – it is L(M) that makes a definite prediction. The predicted values may be close to the new data, or far away (as measured by the Kullback-Leibler distance measure). Imagine using a model repeatedly in this two-step process; there will doubtless be some variation among these repetitions in terms of how well the fitted model predicts new data. The *average* performance of the model is what defines its predictive accuracy. The predictive accuracy of M is the *expected* likelihood of L(M).

Having models that are predictively accurate may be a desirable goal, but how is one to ascertain how predictively accurate a model is apt to be? Akaike's (1973) remarkable theorem provides an answer:

> An unbiased estimate of the predictive accuracy of model M ≈Log-likelihood[L(M)] - k.

One takes the logarithm of the likelihood of the fitted model and subtracts k, the number of adjustable parameters.[16] Complex models, when fitted to the data, tend to have higher log-likelihoods than simpler ones, but they also incur a larger penalty because of their complexity. For a complex model to have a higher AIC value than a simpler one, it isn't enough that the complex model fit the data better; it must fit the data better by a sufficient margin to overcome the fact that it is more complex.

In terms of our problem of comparing the unified and disunified models (U) and (D), the Akaike criterion entails that if L(U) and L(D) fit the data about equally well, we should prefer the unified model. Unified models aren't always preferable to disunified models, but unification nonetheless is epistemically relevant. Unification is evidence of superior predictive accuracy when the unified model contains fewer adjustable parameters.

As mentioned earlier, Akaike's criterion for model selection is not the only game in town. However, AIC and the other criteria now in use agree that high likelihood of L(M) is good news for a model, whereas a large number of parameters is bad. The criteria differ in how they *weight* the importance of these two factors. However, for present purposes, their points of agreement are more important than the things they disagree about. AIC is typical of these other criteria in that it embodies a defeasible presumption in favor of unification.

4. CONCLUDING COMMENTS

I began this paper by asserting that unification has no objective significance in the context of Hempel's problem, but that it is objectively relevant to solving Peirce's problem. I expressed this contrast by saying that the relevance of unification depends on our interests in the former case, but not in the latter. However, in explaining Akaike's ideas, I first described the *goal* of predictive accuracy and then outlined how AIC provides a *means* for achieving that goal. Doesn't this mean that unification is equally subjective and interest-relative in the two contexts – that principle (P) has no more claim on our allegiance than principle (H)? No – the contrast between subjective and objective survives our recognizing that the difference between good and bad inference depends on what our goals are in drawing inferences. If you want to figure out which of several models will be most predictively accurate, then it is an objective fact that unification is relevant. However, if you want to figure out which of the propositions you already believe should be used to explain a body of observations, then whether you choose a unified or a disunified account depends on your interests. *Disunified explanation* is not a contradiction in terms, nor are disunified explanations inherently inferior to unified explanations. If we discover that smoking cigarettes increases one's risk of lung cancer, and that X-cigarettes and Z-cigarettes differ in their carcinogenic properties, we face a choice. We can explain Smith and Jones' cancer by describing what they have in common, or by describing a way in which they differ. They got cancer because they both smoked cigarettes, but it also is true that inhaling x-particles caused Smith's cancer while inhaling z-particles caused Jones'.[17]

NOTES

1. In Carl G. Hempel, *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press, 1965.
2. Gilbert Harman, "Inference to the Best Explanation." *Philosophical Review* 74, 1965, pp. 88-95; Gilbert Harman, *Change in View*. Cambridge: MIT Press, 1986; William Lycan, *Judgment and Justification*. Cambridge: Cambridge University Press, 1988; Peter Lipton, *Inference to the Best Explanation*. London: Routledge, 1991.
3. This more general formulation is needed if future events don't explain the occurrence of past ones, or if one effect of a common cause does not explain the other.
4. An exception to this pattern is Paul Humphreys, *The Chances of Explanation*. Princeton University Press, 1989, which provides a sustained discussion of the idea of an explanatory factor.
5. Or more precisely – one proposition can be obtained from the other by uniform substitution of names.
6. Paul Oppenheim and Hilary Putnam, "Unity of Science as a Working Hypothesis", in H. Feigl, G. Maxwell, and M. Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, Minneapolis: University of Minnesota Press, 1958, pp. 3-36.
7. Hilary Putnam, "Psychological Predicates", In W. Capitan and D. Merrill (eds.), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, 1967. Reprinted as "The Nature of Mental States" in *Mind, Language, and Reality*. Cambridge, England: Cambridge University Press, 1975; Hilary Putnam, "Philosophy and our Mental Life", In *Mind, Language, and Reality*. Cambridge, England: Cambridge University Press, 1975.
8. Wesley Salmon, *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press, 1984.
9. There is another reason to reject screening-off as an explication of the concept of explanatory relevance. Suppose Moriarty's death is caused by Holmes' shooting him through the heart. Presumably, the shooting is explanatorily relevant. However, the fact that Holmes shot Moriarty is screened off from Moriarty's death by a weaker, disjunctive specification of all the many events that would have killed Moriarty if they had occurred. For example, given that either Holmes shot Moriarty through the heart or Watson poisoned him with cyanide, Moriarty was bound to die, regardless of whether Holmes in fact poisoned him.
10. David Lewis, "Causal Explanation," In *Philosophical Papers*, vol. 2, Oxford University Press, 1986, pp. 214-240.
11. Jaegwon Kim (in *Supervenience and Mind*, Cambridge: Cambridge University Press, 1993) argues that countenancing both micro- and macro-causal explanations of the same event commits one to an objectionable form of overdetermination. However, overdetermination is usually understood as the occurrence of two causes that are independent of each other, with each sufficient for the effect. In this sense, Moriarty's death is overdetermined, if Holmes and Watson each simultaneously fire a bullet through his heart. Applying this point to the question of mental and physical causation that Kim considers, we should conclude that there is no genuine overdetermination here; if the macro-facts supervene on the micro, then micro and macro are not independent. See Elliott Sober, "Physicalism from a Probabilistic Point of View," *Philosophical Studies* 95, 1999b, 135-174, for discussion.

12. Elliott Sober, "The Multiple Realizability Argument against Reductionism." *Philosophy of Science* 66, 1999, 542-564.
13. Unlike Putnam, Philip Kitcher (in "Explanatory Unification and the Causal Structure of the World." In P. Kitcher and W. Salmon (eds.), *Scientific Explanation*. Minneapolis: University of Minnesota Press, 1989, pp. 410-505. 1989) does provide an account of explanation, rather than simply relying on an intuitive understanding of what explanatory relevance is. Kitcher rejects the idea that explanations must be causal and uses the idea of unification as his central idea. I won't discuss Kitcher's proposal here except to note that it seems not to deal adequately with the problem of explanatory asymmetry; on this, see Eric Barnes "Explanatory Unification and Scientific Understanding," *Proceedings of the Biennial Meetings of the Philosophy of Science Association*, 1, 1992, pp. 3-12, and Todd Jones, "How the Unification Theory of Explanation Escapes the Asymmetry Problem," *Erkenntnis*, 43, 1995, pp. 224-240. In an earlier publication, Kitcher (in "1953 and All That – A Tale of Two Sciences." *Philosophical Review*, 93, 1984, pp. 335-373) used Putnam's peg and board style of reasoning to argue that Mendelian genetics does not reduce to molecular biology, but without the apparatus of the unification theory of explanation that Kitcher later developed.
14. Hirotugu Akaike, "Information Theory as an Extension of the Maximum Likelihood Principle." In B. Petrov and F. Csaki (eds.), *Second International Symposium on Information Theory*. Budapest: Akademiai Kiado, 1973, pp. 267-281. See also Kenneth Burnham and David Anderson, *Model Selection and Inference – a Practical Information-Theoretic Approach*. New York: Springer, 1998; Malcolm Forster and Elliott Sober, "How to Tell when Simpler, More Unified, or Less Ad Hoc Theories will Provide More Accurate Predictions." *British Journal for the Philosophy of Science* 45, 1994, pp. 1-36; Elliott Sober, "Parsimony, Instrumentalism, and the Akaike Framework," *Philosophy of Science*, 2002, forthcoming.
15. To see why, consider a simpler problem: what are the maximum likelihood values for the parameters in (U) and (D) that pertain to a single category of smoker? For example, suppose that the frequency of lung cancer among those who have smoked two packs a day for thirty years is 0.001 for those who have smoked X-cigarettes exclusively and is 0.002 for those who have smoked Z-cigarettes exclusively. If these two types of smoker are equally represented in the sample, L(U) will say that the probability of cancer for such individuals is 0.0015, while L(D) will say that the probability of cancer for the X-smokers is 0.001 and the probability for the Z-smokers is 0.002. The data are more probable according to the latter hypothesis.
16. The formulation of Akaike's result that Forster and Sober (1994) recommend is that an unbiased estimate of the model's predictive accuracy per datum is $(1/N)\{\text{Log-likelihood}[L(M)] - k\}$, where N is the number of data.
17. Here I'm glossing over the problem of how facts about type- and token-causation are related, on which see Elliott Sober, "Two Concepts of Cause," in P. Asquith and P. Kitcher (eds.), *PSA 1984 – Proceedings of the Philosophy of Science Association*, E. Lansing, Michigan, vol. 2, pp. 405-424. The fact that Smith's smoking increased his risk of getting lung cancer does not mean that his smoking token-caused the cancer; perhaps a micro-examination of Smith's lungs would reveal that his cancer was due to asbestos exposure, not smoking. The statistical inference problem described above concerns claims about type causation, whereas the explanation of Smith and Jones' lung cancer presumably calls for facts about token causation.