NOÛS 38:2 (2004) 197–232

Would "Direct Realism" Resolve the Classical Problem of Induction?

MARC LANGE
University of North Carolina at Chapel Hill

I

Recently, there has been a modest resurgence of interest in the "Humean" problem of induction. For several decades following the recognized failure of Strawsonian "ordinary-language" dissolutions and of Wesley Salmon's elaboration of Reichenbach's pragmatic vindication of induction, work on the problem of induction languished. Attention turned instead toward confirmation theory, as philosophers sensibly tried to understand precisely what it is that a justification of induction should aim to justify. Now, however, in light of Bayesian confirmation theory and other developments in epistemology, several philosophers have begun to reconsider the classical problem of induction.

In section 2, I shall review a few of these developments. Though some of them will turn out to be unilluminating, others will profitably suggest that we not meet inductive scepticism by trying to justify some alleged general principle of ampliative reasoning. Accordingly, in section 3, I shall examine how the problem of induction arises in the context of one particular "inductive leap": the confirmation, most famously by Henrietta Leavitt and Harlow Shapley about a century ago, that a period-luminosity relation governs all Cepheid variable stars. This is a good example for the inductive sceptic's purposes, since it is difficult to see how the sparse background knowledge available at the time could have entitled stellar astronomers to regard their observations as justifying this grand inductive generalization. I shall argue that the observation reports that confirmed the Cepheid period-luminosity law were themselves "thick" with expectations regarding as yet unknown laws of nature. The observations that confirmed the Cepheid law could not have been made without being fortified by these expectations, and these expectations, in turn, made the observations count as evidence for the Cepheid law.

"Direct realism" has recently become popular in epistemological accounts of perception. In section 4, I shall investigate whether direct realism, if correct, would help to support my view regarding the ground of an inductive leap like Leavitt's and Shapley's. Ultimately, I shall argue that direct realism prompts the rejection of the Humean principle, captured perhaps most explicitly by Bayesian confirmation theory, that our observations alone are always insufficient to confirm our predictions regarding unexamined cases. In other words, the rejected principle says that, to ground a given inductive confirmation, our observations must be supplemented either (1) by some ampliative rule governing all correct inductions, where the details (not to mention the justification) of such a rule remain mysterious, or (2) by certain independent background opinions relevant to this particular induction, which apparently would require their own empirical support and thereby launch an infinite regress.

By rejecting this Humean principle, the approach prompted by direct realism accounts for the curious ease with which a grand inductive generalization (such as Leavitt's and Shapley's) can manage to become justified even when (or perhaps partly because) the relevant background knowledge is very meager. In section 4, I shall argue that observations whose reports are "thick" in the relevant sense are, paradoxically, among the sorts of observations that are most readily made in a new area of investigation, where there is a dearth of relevant background knowledge. In section 5, I shall further investigate the astronomers' warrant for the expectations needed to sustain the confirmation of the Cepheid law. I shall argue that these expectations were necessary not only for classifying certain variables as Cepheids, but also for discovering any variable stars in the first place.

As a reply to the inductive sceptic, this approach may be accused of merely replacing the problem of induction with the problem of understanding how one becomes qualified to observe (under certain circumstances) that a certain sort of fact about the "external world" obtains. Yet the latter problem would be with us in any case, should direct realism be correct. So to eliminate the problem of induction, despite retaining the problem of the qualified observer, would represent progress.

There is inevitably a certain hubris in reporting with qualified optimism on some approach for resolving the problem of induction, perhaps the most infamous sinkhole in all philosophy. But the problem's infamy derives from its perennial significance and stubbornness—the very features that ought to drive philosophers to resume their assault upon it. It is scandalous that according to some philosophers, concern with the problem of induction betrays a want of philosophical talent, tact, or taste.

II

Since the demise of the ordinary-language and Reichenbachian research programs for resolving the problem of induction, epistemological naturalism and externalism have gained some popularity. They have inspired a response to the problem of induction (Brueckner 2001, van Cleve 1984, Sankey 1997; see also Kornblith 1993). According to this response, an inductive argument from the frequent success of past inductive inferences to the likely success of some new inductive inference suffices to justify that inference. This may appear circular, but the naturalist-externalist holds that if inductive reasoning from true premises does, in fact, tend to lead to the truth, then an inductive argument can justify its conclusion even if the reasoner has no non-circular basis for believing that it does.

To my mind, this approach simply fails to engage with the traditional problem of induction. It does not set out to persuade the inductive sceptic that she has a good reason to believe that a given inductive argument will likely yield the truth regarding unexamined cases. Rather, this approach changes the subject. Suppose the externalist can persuade the sceptic that to be justified in some belief is to arrive at it by reliable means. Then the sceptic is persuaded that *if* induction is actually reliable, the conclusion of an inductive argument (from justified premises) is justified. She is also persuaded that *if* induction actually is reliable, an inductive reasoner is justified in her belief (arrived at inductively, from the frequent success of past inductive inferences) that induction will continue to be reliable. Nevertheless, the externalist has *not* persuaded the sceptic that induction *is* reliable.

Bayesian confirmation theory has also become much more prominent since the demise of the ordinary-language and Reichenbachian research programs for resolving the problem of induction. Suppose it could be shown—perhaps through a Dutch Book argument or an argument from calibration (Lange 1999b)—that rationality obliges us (in typical cases) to update our opinions by Bayesian conditionalization (or some straightforward generalization thereof, such as Jeffrey's rule). This result would still seem to leave us far from having a justification of induction. Whether Bayesian conditionalization yields induction or counterinduction, whether it underwrites our ascribing high probability to "All emeralds are green" or to "All emeralds are grue," whether it leads us to regard a relatively small sample of observed emeralds as having any bearing at all on unexamined emeralds—all depend on the prior probabilities. So a justification of induction requires some kind of solution to the problem of the priors. Otherwise, it remains unexplained why we ought to reason *inductively*.

This argument has been challenged in several ways. For example, in a recent book devoted entirely to Hume's problem, Colin Howson (2001) argues that a justification of induction should explain why *we* ought to reason inductively. Thus, it can appeal to *our* prior probabilities; Bayesian conditionalization, acting on these particular priors, underwrites recognizably *inductive* updating. The "initial assignments of positive probability . . . cannot themselves be justified in any absolute sense" (2001, 239). But never mind, Howson says. Inductive arguments are in this respect

> like sound deductive arguments, they don't give you something for nothing: you must put synthetic judgements in to get synthetic judgements out. But get them out you do, and in a demonstrably consistent way that satisfies certainly the majority of those intuitive criteria for inductive reasoning which themselves stand up to critical examination. (2001, 239, cf. 171)

All we really want from a justification of induction is a justification for *updating* our beliefs in a certain way, and that is supplied by arguments showing Bayesian conditionalization to be rationally compulsory. Ian Hacking addresses the problem of induction in the final three chapters of his new survey of induction and probability (another sign of resurgent interest in the "Humean" problem), and he puts the argument thus:

> At any point in our grown-up lives (let's leave babies out of this) we have a lot of opinions and various degrees of belief about our opinions. The question is not whether these opinions are 'rational'. The question is whether we are reasonable in modifying these opinions in light of new experience, new evidence. (2001, 256)

Well, the question in which I am interested, the traditional question, is whether by reasoning inductively, we arrive at knowledge. If knowledge involves justified true belief, then the question is whether true beliefs arrived at inductively are thereby justified. And if an inductive argument, to justify its conclusion, must proceed from a prior state of belief that we are entitled to occupy, then the question becomes whether we are justified in holding those prior opinions, and if so, how come.

I said that Bayesian conditionalization can underwrite reasoning that is intuitively inductive, but with other priors plugged into it, Bayesian conditionalization underwrites reasoning that is counterinductive or even reasoning that involves the confirmation of no claims at all regarding unexamined cases. However, it might be objected that if hypothesis h (given background beliefs b) logically entails evidence e, then as long as $\mathrm{pr}(h|b)$ and $\mathrm{pr}(e|b)$ are both non-zero, it follows that $\mathrm{pr}(e|h\&b) = 1$, and so by Bayes's theorem, we have $\mathrm{pr}(h|e\&b) = \mathrm{pr}(h|b)\mathrm{pr}(e|h\&b)/\mathrm{pr}(e|b) = \mathrm{pr}(h|b)/\mathrm{pr}(e|b) > \mathrm{pr}(h|b)$, meaning that e confirms h. On this objection, then, Bayesian conditionalization automatically yields induction. But this confirmation of h (of "All emeralds are green," for example) by e ("The emerald currently under examination is green") need not involve any *inductive* confirmation of h—roughly, any confirmation of h's *predictive accuracy*. For example, it need not involve any confirmation of g: "The next emerald I examine will turn out to be green." Since g (given b) does not logically entail e, $\mathrm{pr}(e|g\&b)$ is not automatically 1, and so $\mathrm{pr}(g|e\&b)$ is not necessarily greater than $\mathrm{pr}(g|b)$.[1]

Howson contends that a justification of the requirement that we update in accordance with Bayesian conditionalization would suffice to justify induction. But Howson recognizes that Bayesian conditionalization can underwrite non-inductive ampliation just as easily as induction. So his "justification" of induction would embrace counterinductive reasoning (for example) as well as induction. A justification of induction, at least of the sort I seek, must privilege inductive reasoning over the alternatives to it. Hence, Howson's argument would need to be supplemented with an account of how prior probability distributions that lead (by Bayesian updating) to recognizably inductive reasoning come to be justified, in contrast to priors leading to (say) counterinductive reasoning.

Why does Howson say otherwise? He insists that it is no part of the justification of induction to justify the choice of priors, just as deductive logic does not concern itself with justifying the premises of deductive arguments (Howson 2000, 2; also 164, 171, 239; cf. Howson and Urbach 1989, 189–90). To my mind, this parallel between deduction and induction is not apt. It presupposes that prior probabilities are the premises of inductive arguments—are, in other words, the neutral input or substrate to which is applied Bayesian conditionalization, an inductive rule of inference. But, as Howson rightly emphasizes, it is Bayesian conditionalization that is neutral; anything distinctively "inductive" about an episode of Bayesian updating must come from the priors. Consequently, a justification of induction must say something about how we are entitled to those priors.

I shall endorse the idea that a justification of induction is concerned exclusively with the inductive leap; it does not need to explain how the observations from which induction proceeds manage to acquire their positive justificatory status. Accordingly, the positive justificatory status of the posterior $\mathrm{pr}(e)$, acquired directly via observation, falls beyond the scope of a justification of induction. Nevertheless, by characterizing inductive arguments as having the same "logical form" as various arguments that we would not pre-theoretically characterize as inductive, and then suggesting that a justification of induction need address only that form of reasoning, Howson's view essentially defines the problem of induction out of existence. Wherever we finally locate the difference between induction and other (less savoury) ampliative reasoning (whether in their logical form, or in their premises, or somewhere else), we must say something about why we are entitled to employ the elements distinguishing inductive reasoning and not similarly entitled to employ those distinguishing induction's rivals. Anything less is not enough.[2]

Some personalists about probability would argue that if anything distinctively "inductive" about Bayesian updating must come from the priors, then so much the better for resolving the problem of induction, since we are automatically entitled to adopt *any* probability distribution as our priors. This view is prompted by the notorious difficulties (associated with Bertrand's paradox of the chord) attending any principle of indifference for adjudicating among rival priors. In yet another recent treatment of the problem of induction, Samir Okasha writes:

> Once we accept that the notion of a prior distribution which reflects a state of ignorance is chimerical, then adopting any particular prior distribution does not constitute helping ourselves to empirical information which should be suppressed; it simply reflects the fact that an element of guess work is involved in all empirical enquiry. (2001, 322)

Okasha's argument seems to be that a prior state of opinion embodies no unjustified information about the world since any prior opinion embodies *some* information. But the inductive sceptic should reply by turning this argument around: Since any prior opinion strong enough to support an inductive inference embodies some information, no prior opinion capable of supporting an inductive inference is justified.

In other words, Okasha's argument seems to be that there are no objectively neutral priors, so if the inductive sceptic accuses our priors of being unjustified,

> we need only ask the sceptic 'What prior probability do you recommend?' [...] It does not beg the question to operate with some particular prior probability distribution if there is no alternative to doing so. Only if the inductive sceptic can show that there *is* an alternative, i.e., that 'information-free' priors do exist, would adopting some particular prior distribution beg the question. (Okasha 2001, 323)

But there *is* an alternative to operating from a prior opinion strong enough to support ampliative inferences. If the sceptic is asked to recommend a prior probability, she should suggest a distribution that makes no probability assignment at all to any claim about the world that concerns logically contingent matters of fact. By this, I do not mean the extremal assignment of *zero* subjective probability to such a claim. That would be to assign it a probability: zero. Nor do I mean assigning it a *vague* probability value. I mean making no assignment at all to any such claim. As Okasha rightly emphasizes (2001, 308–9), the sceptic is not making merely the boring point that induction is fallible. Her point, rather, is that regarding predictions about unobserved cases, there is *no* degree of confidence to which we are entitled.

Admittedly, the sceptic's prior distribution violates the requirement that the domain of a probability function be a sigma algebra. For example, it may violate the additivity axiom [$\mathrm{pr}(q \text{ or } \sim q) = \mathrm{pr}(q) + \mathrm{pr}(\sim q)$] by assigning to ($q$ or $\sim q$) a probability of 1 but making no probability assignment to q and none to $\sim q$. Some Bayesians would conclude that the sceptic's "pr" does not qualify as a probability function. However, the sceptic is not thereby made vulnerable to a Dutch Book. She is not thereby irrational or incoherent. What is the worst thing that can be said of her? That she shows a certain lack of commitment. That characterization will hardly bother the sceptic! It may well be overly restrictive to require that the domain of a probability function be a sigma algebra. (See, for instance, Fine (1973, 62) or any paper discussing the failure of logical omniscience.[3])

Though an observation's direct result may be to assign some probability to e, the sceptic's prior distribution fails to support inductive inferences from our observations (since it omits some of the probabilities required by Bayesian conditionalization or any generalization of it). But that is precisely the inductive sceptic's point. There is no alternative to operating with a prior distribution that embodies information about the world, as Okasha says, *if* we are going to carry out ampliative inferences from our observations. But to presuppose that we are justified in carrying out such inferences is obviously to beg the question against the inductive sceptic.

In view of the fact that the same observation (i.e., posterior pr(e)) will have a radically different confirmatory impact through Bayesian conditionalization, depending on the rest of the probability distribution (e.g., pr($e|g$)), it is often said that Bayesian confirmation theory nicely captures Hume's insight that our observations alone, without a theoretical context, are unable to confirm or to disconfirm anything. As Elliott Sober says:

> If my beliefs about my present and past environments [e.g., "The sun has risen each day that I have made an observation"] are to justify the predictions and generalizations I believe [e.g., "The sun will rise tomorrow"], then I have to assume something about the *relationship* between [the beliefs at the two levels]. Perhaps the principle of the uniformity of nature ("the future will resemble the past") is an example of such a bridge principle. Or some other bridge principle might be proposed. But my present perceptions and memories are simply not enough, taken all by themselves.... Not only do lower level statements fail to deductively imply high level statements; I also want to claim that lower level statements, all by themselves, aren't enough to provide a justification for higher level statements. Lower level statements, taken by themselves, don't even provide good evidence for higher level statements. (Sober 1995, 199; see also Howson 2001, 180)

Of course, the Humean sceptic now intervenes to demand our warrant for the probabilities we assigned to various "bridge principles" h', probabilities that enabled e to confirm h inductively. If we justify those probabilities by appealing to other observation reports e' that confirmed h', the sceptic demands our warrant for the probabilities we assigned to various other bridge principles h'' in virtue of which e' was made relevant confirmation-wise to h'. And the regress ensues.

Sober contends that although "we aren't justifying a belief at one level strictly in terms of beliefs that are on lower levels" (1995, 201), the regress does not set in. What counts as being justified depends on the intended audience. If both you and your target audience believe the bridge principle h', then you can justify yourself to them by appealing to h' (1995, 202). Sober recognizes that if your intended audience is an inductive sceptic, then the problem is "insoluble," but "this does not mean that more familiar problems of rational justification are, too" (1995, 202).

To me, this is giving up on what we wanted, which was precisely a reply to the inductive sceptic. But even for "more familiar problems of rational justification," this is too quick. I can justify myself by appealing to h' only if I am, in turn, entitled to believe h'. That I and my target audience both believe h' does not mean that either of us is *entitled* to believe it. Of course, someone who believes h' presumably also believes herself entitled to this belief. But she may be mistaken in believing herself so entitled.

"Entitled in what sense?" Sober might ask. "For what target audience?" I have argued as if there were entitlement *simpliciter*, not merely for a given audience, and Sober might see this as begging the question. But, I would ask Sober, what is it to "justify" yourself to a given target audience? Suppose it is to use your observations to justify your probability assignment through Bayesian updating from a prior probability distribution that your target audience is willing to grant. Then what does your belief's being "justified" for a given target audience have to do (according to you) with the likelihood of your belief's being true? Why are justified beliefs (in this sense) worth having? As we have seen, a Bayesian will recognize that for target audiences who are willing to grant different priors, the same observations will have different confirmatory impacts; given the same observations, different beliefs will be "justified for" different target audiences. If I should aim to hold justified beliefs, then justified *for whom*? Clearly, I cannot believe that for *every* conceivable target audience, my beliefs' being justified "for them" increases the likelihood of my beliefs' being true. Why should I care whether or not my beliefs are justified for a certain audience—unless justifications are merely dramatic performances that I offer for practical purposes, to smooth my relations with an audience about whom I care for non-epistemic reasons? But then we have left behind the problem of understanding induction's *epistemic* justification.

Sober's view is thus subject to the same objection as the "ordinary-language" view. Sober explains that "[a]ccording to Strawson, it is entirely rational to use inductive methods to formulate our beliefs about the future, even though we can offer no good reason for expecting that the method will lead to true beliefs.... I find Strawson's argument unconvincing." (1995, 194) Strawson's view is that what we *mean* by having a good reason is having an inductive reason. The standard objection insists that to support this identification, we need some good non-circular reason for defining "good reason" in terms of induction rather than in terms of some other rule for ampliation—that is, for supposing that induction will often lead us to the truth (Salmon 1967, 51). Likewise, Sober's own proposal is that in a given context, what we *mean* by having a good reason is having a reason relative to some background implicitly being taken for granted. But to support this identification, we need in this context to have some good non-circular reason for defining "good reason" in terms of this particular background rather than any other.

Confirmation theory has come a long way from the models of Hempelian instance confirmation, Goodmanesque projection of observed uniformities onto unexamined cases, Reichenbach's straight rule, enumerative induction, and so forth. The Bayesian account leaves plenty of room for background knowledge but assigns no special place to the belief that unexamined cases will be *like* examined cases. Rather, an inductive argument depends on "a vast store of background information" (Okasha 2001, 309). Consequently, we should not seek to justify induction by trying to justify some general principle of the uniformity of nature that is supposed to underwrite all inductions. Instead of thinking about our task as the justification of induction as a whole, it would perhaps be more enlightening to uncover the basis of some particular inductive inference.

This will inevitably lead us to investigate the relevant background beliefs. What, in turn, is their ground? It might be thought that our problem is now essentially solved, since (as Hacking says) "[a]t any point in our grown-up lives (let's leave babies out of this) we have a lot of opinions," so for every background opinion whose warrant is demanded, we should be able (at least in principle) to identify some evidence warranting it in light of some yet prior opinion. However, this means that the "answer" to each challenge presumes that there is warrant for some yet prior opinion. Until we reach the end of this regress, none of these challenges has been met.

It might be objected that in demanding such a regress-stopper, we have departed from the thoroughly realistic attitude that led us to try to uncover the warrant for some actual inductive episode. For such a regress-stopper would amount to nothing less than an argument against occupying the inductive sceptic's prior probability distribution, and this sceptical prior might be considered terribly artificial—entirely foreign to us grown-ups, with all of our opinions (as Hacking says), "very different from the [epistemic starting-point] in which we find ourselves in real life" (Okasha 2001, 319). Solutions to Sober's "more familiar problems of rational justification" might be thought not to require a response to the inductive sceptic.

I disagree, in the main. The history of science is filled with cases in which the scientists' relevant background knowledge was apparently extremely sparse at best. Nevertheless, scientists were somehow able to take certain uniformities among their observations and to project them justly onto a broad range of unexamined cases.

Here is an example. In 1908, Henrietta Leavitt (1908) announced her discovery of 1,777 variable stars in the Magellanic Clouds. By 1912, she had found 25 of the variables in the Small Magellanic Cloud (SMC) on enough photographic plates to allow her to determine their periods very precisely. On both occasions, she remarked that in the shape of their light curves, all of these variables (and, evidently, a majority of the SMC variables) "resemble the variables found in globular clusters, diminishing slowly in brightness, remaining near minimum for the greater part of the time, and increasing very rapidly to a brief maximum" (1912, 1; also 1908, 107). She noted a "remarkable relation" (1912, 1; also 1908, 107): the brighter the star, the longer its period between maxima. She took this evidence "to warrant the drawing of general conclusions," a "law" (1912, 2). Of course, the relation Leavitt graphed (figure 1) was merely between the period and brightness of certain stars *as seen from Earth*. But all of the stars in the SMC are about the same distance from Earth; the SMC's depth was presumed to be a negligible fraction of its distance from Earth. Therefore, Leavitt argued (1912, 3), if one SMC star is brighter than another as seen from Earth, then "probably" the former's real luminosity is greater than the latter's. Thus Leavitt took her evidence as suggesting a relation between period and *intrinsic* luminosity: "their periods are apparently associated with the actual emission of light, as determined [in an unknown manner] by their mass, density, and surface brightness" (1912, 3).

In 1913, Ejnar Hertzsprung took Leavitt's evidence as justifying the application of her period-luminosity "law" to the 13 variables that exhibit light curves of the type Leavitt had described and that have the best determined motions relative to Earth. (In note 5, I explain why he chose these particular stars.) As had become customary, he termed the stars with such light curves "Cepheid" variables after the star delta Cephei, a well-known example. He used these stars and the "law" to compute the SMC to be 30,000 light-years away, the largest astronomical distance yet estimated. In 1915, Harlow Shapley extended Hertzsprung's work by projecting the period-luminosity law to all known Cepheids in the Milky Way (our galaxy). He then estimated their distances from Earth by comparing their brightness as seen from Earth to their actual luminosities (ascertained by applying the period-luminosity law to their periods). Most famously, Edwin Hubble in 1924 projected the period-luminosity relation to 40 Cepheids he found in the nebula M31 in the constellation Andromeda. He thereby estimated the nebula's distance from Earth to be (as we would say) astronomical: about a million light-years. (He discovered that the nebula actually lies outside the Milky Way; it is now known as the Andromeda galaxy.)

Yet there was scant background knowledge on which to ground these projections. It is difficult to see what *evidence* existed that Hubble's 40

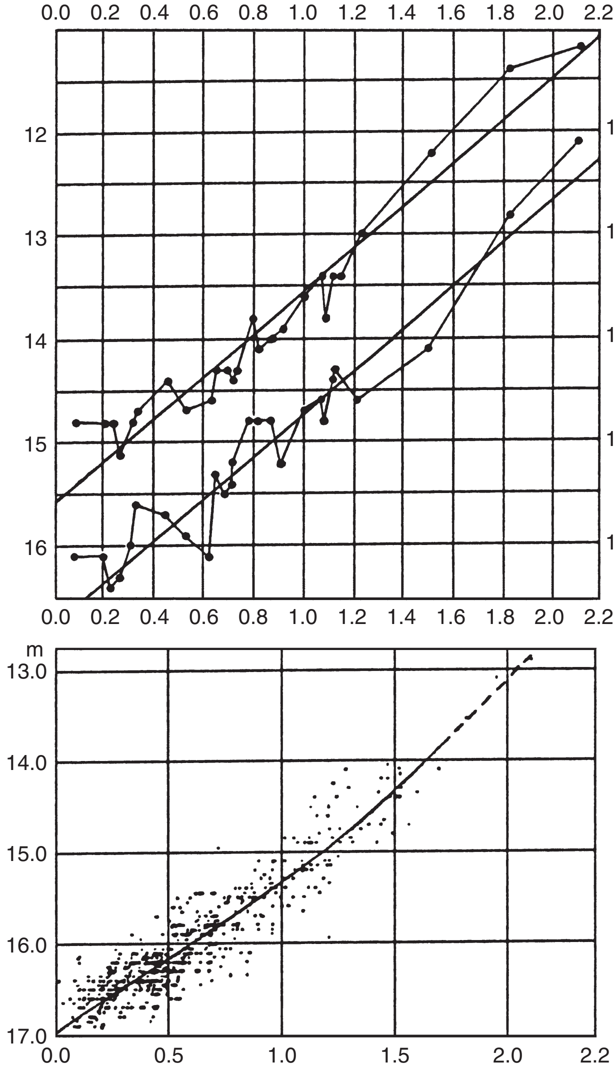

Figure 1. Upper diagram: Leavitt's original period-brightness curve for the SMC (Leavitt 1912, 3), showing each star's maximum and minimum brightness as a function of its period of variation. Lower diagram: The period-brightness relation for mean magnitudes, plotted with all of the SMC Cepheids measured by 1946 (Hoffleit 1993, 4). In both graphs, the y-axis is apparent magnitude (brighter stars have lower magnitude numbers) and the x-axis is the logarithm of the period (in days).

Cepheids in M31, for example, were like Hertzsprung's 13 Milky-Way Cepheids or Leavitt's 25 "cluster variables." For example, there was no real understanding of the mechanism by which any of these stars varies in its brightness.[4]

Of course, background knowledge of a minimal sort played an obvious part in supporting the generalization of Leavitt's relation to all Cepheids.[5] For example, astronomers knew the constellation boundaries to be wholly arbitrary. So astronomers would have been unjustified in regarding Cepheids in the Andromeda nebula as having greater bearing on other Cepheids in virtue of their also lying in the constellation Andromeda rather than in some neighboring constellation. Nevertheless, in such a theoretically impoverished context, there does not seem to be sufficient background knowledge to justify Hertzsprung's, Shapley's, and Hubble's grand inductions. Why couldn't Cepheids be a heterogeneous category? For instance, why couldn't Cepheids located near a galactic core differ intrinsically from Cepheids located near a galaxy's periphery?

To justify induction, we must understand precisely how scientists in theoretically impoverished contexts nevertheless manage to become entitled to suspect the existence of various sorts of grand uniformities in nature—how, for instance, stellar astronomers were entitled to suspect certain sorts of uniformities among all Cepheid variables despite knowing so little about variable stars. It seems as if the empirical work I needed to carry out last academic quarter in order to justify believing that a given student cheated on his exam far *exceeded* the empirical work that astronomers needed to undertake a century ago in order to justify believing that Cepheids in M31 obey whatever period-luminosity relation is obeyed by SMC Cepheids. To explain this remarkable difference, it clearly does not suffice merely to gesture toward "background knowledge."

Of course, I am presuming that Leavitt, Hertzsprung, and the rest *were* actually justified in their inferences. However, this should not be understood to mean that the evidence supporting these inferences was utterly overwhelming. Greater caution and even outright dissent would also, for a considerable period at least, have been justified stances to take. Indeed, Hertzsprung himself offered a good reason to oppose co-classifying the "cluster variables" with the other alleged Cepheids (see note 12), and Heber Curtis (Hoskin 1976, 179) wondered whether stars in galactic spiral arms may behave differently from stars in denser concentrations (such as globular clusters and the SMC), since variables appear to be much more common where stars are denser.[6] In short, rationality permits us considerable latitude; frequently, both proponents and critics of a scientific innovation are taking reasonable positions, at least for a time. (This fact tends to drive each side's search for more powerful arguments.)

But while acknowledging that there were good reasons for hesitating to accept the inferences made by Leavitt, Shapley, and the rest, we must not lose sight of the fact that there were also good reasons for placing considerable trust in these inferences. Many astronomers in 1913 would doubtless have agreed with the eminent astrophysicist Henry Norris Russell, who wrote to Hertzsprung:

> I had not thought of making the very pretty use you make of Miss Leavitt's discovery about the relation between period and absolute brightness. There is of course a certain element of uncertainty about this, but I think it is a legitimate hypothesis. (Smith 1982, 72)

Here, from another astronomer, is a typical 1924 comment regarding Hubble's application of the period-luminosity relation to M31:

> The applicability of the relation is, of course, in doubt but probably no more so than in some other cases where it has been applied. (Smith 1982, 118)

Though one could try to argue (without appealing to wholesale inductive scepticism) that astronomers were simply unjustified in inferring that there exists a Cepheid period-luminosity relation, their inference seems to me quite typical of the sort that scientists often make despite occupying theoretically impoverished contexts. The evidence had to be taken seriously. Critics generally offered *arguments* against the existence of a Cepheid period-luminosity relation (rather than merely calling for more evidence), thereby implicitly recognizing the prima-facie bearing of Leavitt's evidence on every Cepheid. (Sometimes explicitly. For example, Curtis (1921, 203) acknowledged that unless there is some reason for supposing SMC Cepheids to behave differently from Cepheids in galactic spiral arms, Leavitt's data provide a good reason for believing that Cepheids not just in the SMC, but also in our galaxy's spiral arms, obey a period-luminosity relation.)

III

Let's now examine this historical episode more carefully and try to uncover the inductive inference's ground.

Cepheids were among the classes of variable stars that had been recognized by Leavitt's day.[7] Each class was associated with one or two prototype stars—the "exemplar" or "model" of the "class", or the "head" of the "family" (Clerke 1903, 319, 327, 328). For the Cepheids, these were delta Cephei and epsilon Aquilae (the first two Cepheids discovered—in 1784, long before any others were). For three other classes, the prototypes were Algol (beta Persei), beta Lyrae, and Mira (omicron Ceti), respectively. "[E]ach of these stars examplifies a certain type or law of variation," wrote the distinguished astronomer Simon Newcomb (1902, 99). In other words, the basis for the classification was the "light curve" depicting the variation in the star's brightness over time. For example, "[w]hen we say that a star is of the Algol type, we mean that it varies in the same way as Algol does..." (Newcomb 1900, 214). Algol-type stars "are invariable in brightness during the greater part of the time, but fade away for a few hours at regular intervals" (1902, 104). (See figure 2.) In contrast, a light curve of

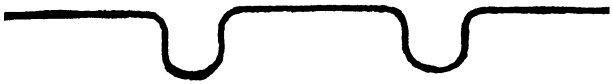

Figure 2. "The law of variation [for Algol type stars] is expressed by a curve of [the above] form" (Newcomb 1902, 102). The y-axis is brightness; the x-axis is time.

the beta Lyrae type alternates between two unequal minima, the star's brightness varying continuously between them. (See figure 3.) We might then expect each Cepheid's light curve to display the shark-fin shape that is so striking in the light curves of delta Cephei and epsilon Aquilae. (See figure 4.) This would not only set the Cepheids distinctly apart from the other types of variables, but also explain why trained astronomers could recognize a Cepheid-type light curve on sight. For example, in his 1924 letter to Shapley announcing his discovery of a Cepheid in M31, Hubble enclosed the star's light curve, which (he commented), "rough as it is, shows the Cepheid characteristics in an unmistakable fashion." He obviously expected Shapley to agree that the star is a Cepheid simply upon seeing the curve (Hoskin 1982, 161). In the ensuing paper, Hubble presented four light curves from among the 40 Cepheids that he had identified in M31, commenting that their "Cepheid characteristics are obvious" (Hetherington 1996, 77).[8]

But this view leads to a dead end. If the Cepheids are merely the variables with light curves possessing certain intrinsic characteristics, then how did astronomers justify deeming "plausible" the hypothesis that Cepheids "are comparable wherever found" (Shapley 1918, 116), which grounded their projecting regularities among a few examined Cepheids onto all other Cepheids? What reason did astronomers have for believing that stars alike in this one respect are alike in other respects?[9] Historians of astronomy are curiously unhelpful here. Hetherington writes:

> Next, [Hertzsprung] assumed that the galactic Cepheids are similar to Cepheids in the Small Magellanic Cloud. Necessity might have been a justification for this assumption of uniformity, had anyone thought to question it, but it did not seem unreasonable. (1996, 40)

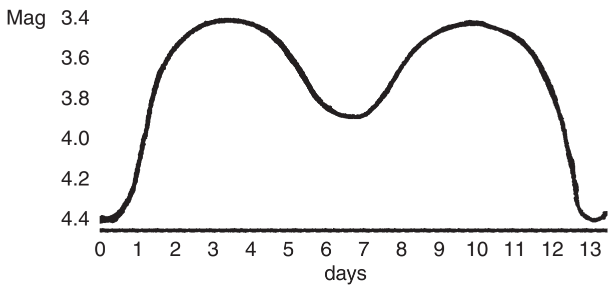

Figure 3. Light curve of beta Lyrae (Clerke 1903, 337).

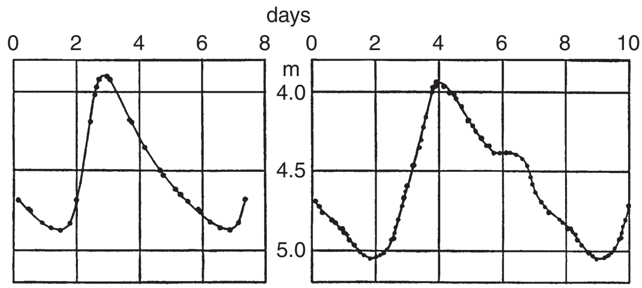

Figure 4. Light curves of delta Cephei, at left, and epsilon Aquilae, at right (Campbell and Jacchia 1946, 63).

Hoskin writes:

> Such a uniformity assumption may appear arbitrary and unjustified. . . . The justification lies in the possibilities opened up by such an assumption. (1982, 6)

This is unconvincing. That astronomers needed to presuppose some such uniformity in order to infer stellar distances from their observations does not explain why astronomers were justified in believing in some such premise, much less this particular one. That a given uniformity assumption would lead to interesting conclusions is a good reason to investigate whether or not it is true, but not a good reason to believe it true.[10]

What evidence, then, did astronomers have that all Cepheids are basically alike? Let's consider a prior question: What range of stars would this evidence have to concern? How was the Cepheid class demarcated?

Although astronomers could tell just from looking at a star's light curve whether or not the star was a Cepheid, the Cepheids' light curves form a varied group. (See figure 5.) While many have steep, narrow maxima (like delta Cephei and epsilon Aquilae), some maxima are quite broad or even double. Although the rise to maximum is often much more rapid than the decline (as Leavitt noted), some curves are fairly symmetric, even sinusoidal. In sum, one could not catch on to the Cepheid class simply by being told that Cepheid light curves are those "like" delta Cephei's and epsilon Aquilae's. The range of stars over which astronomers were prepared to generalize was not understood as the class with light curves displaying a shark-fin shape or some other purely formal, intrinsic characteristic. No set of formal, intrinsic, necessary and sufficient conditions for qualifying as a Cepheid-type light curve was ever proposed in the period's scientific literature.

Shown two Cepheid light curves, an untrained observer might not automatically deem them alike. But between any two Cepheid light curves, however dissimilar, one can construct a smooth progression of light curves

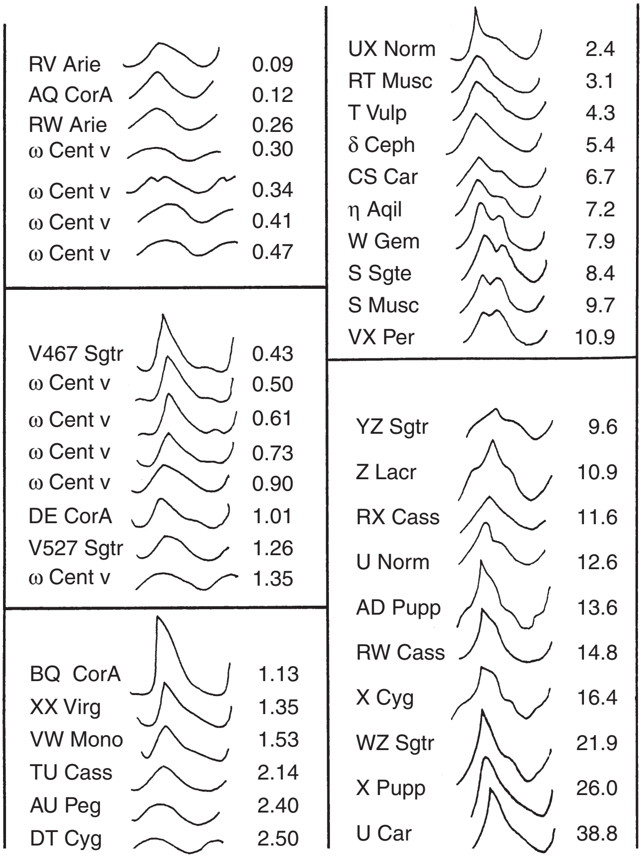

Figure 5. Some Cepheid light curves with periods given in days, organized to reflect the resemblances among them (Burnham 1978, I, 588; after Campbell and Jacchia 1946, 67).

belonging to other variable stars (all of which are Cepheids). Figure 5 displays several such progressions. For example, between light curves displaying a single sharp maximum and those displaying a broad, double-humped maximum, there are light curves with plateaus or small bumps on the falling branch. In contrast, there aren't variable stars with light curves forming a continuous series between a Cepheid light curve and an Algol-type light curve, for example. Between these, there is a wide gap in the stellar population. Consequently, after being exposed to a host of variable stars' light curves, Cepheids' and non-Cepheids' alike, astronomers learned how to recognize a light curve as belonging to the Cepheid class (or not) by acquiring a sense of the extent of the cluster surrounding the prototype's light curve.[11] An astronomer developed a "good eye." Astronomers were also able to settle possible disputes over a given star's Cepheid character by showing that its light curve is linked to those of well-established Cepheids (such as the prototypes) by a smooth series of intermediate cases, or that a gap in the stellar population exists between them.[12]

Astronomers, then, characterized the range of the Cepheid class ostensively, by pointing to some light curves. But their understanding of the Cepheid class could not *simply* have been as encompassing all stars with light curves that are shaped like those of the prototypes or even as encompassing all stars with light curves that are shaped like the examples in figure 5. These characterizations would clearly have led astronomers to classify as non-Cepheids certain stars that they actually deemed to be Cepheids. Astronomers must have understood the relevant sense of "like" as extending from the prototypical Cepheids out to wherever the nearest significant gap appears in the distribution of stellar light curves. Furthermore, they must have understood that no stars are to be classified as *somewhat* Cepheid, or *quasi* Cepheid, or Cepheid *to some degree*, or *borderline* Cepheid, else they would surely have so classified various stars whose light curves are rather unlike those of the prototypical Cepheids (though able to be connected to them by a smooth series of intermediate cases).

This example recalls a point made familiar by Wittgenstein. In order for an ostensive definition of some term to succeed in teaching you how to apply the term, you must already understand something about what is being defined; any act of ostensive definition presupposes some stage-setting, some prior background understanding of the respects of similarity to be noted among the objects pointed out. How, then, did astronomers understand what was being defined—what the relevant sense of similarity was supposed to be? What was the background understanding that enabled astronomers to go on together from the prototypical Cepheids and apply the term "Cepheid" to the range of light curves they did?

The answer, I think, is not hard to find. Astronomers understood from the outset that the Cepheids were supposed to form a *natural kind* of variable star, a *species* alongside the Algol class, the Beta Lyrae type, and so forth.[13] Let's consider what that means.[14] By virtue of being a natural *chemical* kind, for instance, a chemical compound or element must figure in various sorts of natural laws. These laws specify certain *characteristic properties* that all samples of that substance must possess. (Such a law might say that all samples of the substance are liquids, with such and so optical density and such and so refractive index, if placed under certain conditions of pressure and temperature.) The same *sorts* of characteristic properties (such as boiling point under standard conditions, ionization potential, and reactivity with other chemicals under certain sorts of conditions) must figure in each substance's laws. The characteristic properties, for each substance, are all determined (via natural law) by that substance's *essential properties*, which are possessed *uniquely* by that substance. Since the substance's characteristic properties arise (via natural law) from its essential properties, which the substance shares with no other substance, the substance's characteristic properties are (with relatively rare exceptions) also unique to it.

Analogous considerations apply to other sorts of natural kinds. For instance, a natural kind of sub-atomic particle (e.g., the electron) must figure in various natural laws specifying its characteristic properties (e.g., its charge, mass, gyromagnetic ratio). Each of these properties is shared by all particles of that kind and (by and large) is possessed only by particles of that kind.[15] By the same token, a natural kind of variable star would have to figure in natural laws specifying various characteristic properties possessed by all stars of that kind. These intrinsic properties would all have to follow (via natural law) from the essential properties of that kind, which would involve the mechanism responsible for its light variation. (A variable-star class's "distinguishing features," says Newcomb (1902, 95), "show some radical difference [from the other classes] in the causes on which the variations depend"; "it is but natural to conclude," says Roberts (1895, 284), "that some common cause must operate in producing this common type.") Just as no atom is intermediate between two elements (at least not for very long, at any rate), no molecule between two compounds, and no sub-atomic particle between two particle species, so membership in a natural class of variable star is not supposed to be a matter of degree, and no star is supposed to belong to more than one species. Hence, if a variable star's type is supposed to be ascertainable by observing its light curve, then there must be appreciable gaps between the light curves of stars belonging to different classes.

Why must membership in one of these kinds be all or nothing? A natural kind *F* of variable star would have to figure in a host of strict natural laws of the form "All *F*'s are *G*," and these laws would make no reference to stars that are *F* merely to some degree or other. The laws themselves (whatever they turn out to be) would refer only to "All Cepheids . . . ," "All Algol-type variables . . . ," and so forth. As I have said, the laws governing different kinds tend by and large to demand incompatible characteristic properties, since those properties follow from incompatible essences. (For instance, perhaps the spectral lines of all Cepheids must shift their positions towards the blue as the star is brightening and to the red as the star is dimming, whereas all Algol-type variables must display the opposite pattern of line shifts.) Thus no star can belong to more than one kind.

Astronomers went on from the prototype Cepheids to apply the Cepheid category as they did only because astronomers came to the ostensive definition already understanding that "Cepheid" was supposed to denote a natural kind of variable star. Without this understanding, they would have gone on differently from the prototypes, observing to be "Cepheid" a far different range of light curves—perhaps designating certain stars "somewhat Cepheid" or "Cepheid-like" rather than extending full Cepheid-hood all of the way out to the gap in the stellar population. At this point, we are beginning to move away from the "thin" conception of observation reports that animates the problem of induction. An astronomer's report that a given star belongs to the Cepheid class was "thick" in that it embodied expectations regarding as yet undiscovered facts, such as that there exist certain sorts of laws governing all Cepheid variables.[16]

This "thickness" talk is obviously borrowed from philosophical discussions of ethical concepts (such as "courageous") that are both descriptive and evaluative (Williams 1985, 129 and 140–3). The key feature of these ethical concepts is that their descriptive and evaluative elements are inextricably linked. From a series of examples, one cannot catch on to the thick concept, going on to apply it correctly to new cases, unless one has its evaluative import in mind. This is not some sort of learning disability or psychological quirk on our part, keeping us from "getting" the ostensive definition except under special conditions. Rather, this is a normative, epistemological matter: one cannot *justly* apply the concept correctly to new cases without having its evaluative import in mind. The way it is correct to go on from examples, in light of the concept's evaluative significance, would be arbitrary, gerrymandered, and unmotivated taken independent of these evaluative considerations.[17]

I am arguing that in identifying a star as a Cepheid variable, astronomers purported not only to be describing a particular thing they were seeing, but also to be recognizing certain facts as inductively relevant to various possible generalizations. As in the case of thick ethical concepts, I am arguing that these two elements are inseparable: astronomers justly applied "Cepheid" to the range of light curves they did, extending outward from the prototypes, only because of their commitment to the Cepheid category being a species of variable star, and hence to the existence of various sorts of natural laws covering exactly the Cepheids.[18] This commitment underwrote the use that astronomers made of their Cepheid observations in confirming inductively various candidates for these laws, such as the Cepheid period-luminosity relation.[19]

I am not trying to argue that the inductive relevance accorded by these astronomers to reports of Cepheid variables is built into the meaning or content of "Cepheid" *as opposed to* resulting from various important background beliefs about Cepheids that these astronomers held. To say that the inductive relevance of some Cepheid observation-report e depends on background beliefs b is not to say that this background is an utterly independent variable, in the way that the Bayesian's "$\mathrm{pr}(h|e,b)$" notation suggests. I am arguing that the Cepheid observations had to be accompanied by certain background beliefs.

Accordingly, when Leavitt, Shapley and their colleagues noticed a simple relation among the periods and luminosities of a few examined SMC Cepheids, they automatically possessed a theoretical background that was sufficient (in view of the absence of any other relevant background knowledge) to underwrite the suspicion that this relation holds of all Cepheids. Their background included the belief that every Cepheid is covered by the same laws (of the form "All Cepheids are *G*'s") specifying various properties *G* that are bound up with a star's luminosity and its variation. That is because the astronomers' background included the belief that Cepheids, as a natural kind of variable star, all have the same distinctive mechanism of variation; some intrinsic property, possessed by all Cepheids and only by Cepheids, is responsible for their changes in luminosity by virtue of that property's role in certain natural laws. Possession of this property is a "common cause" of every Cepheid's light curve, whether the Cepheid has been examined yet or not. The astronomers' belief in this common cause committed them (in the absence of any further relevant information) to taking the period-luminosity relation displayed by examined Cepheids and regarding it as relevant confirmation-wise to every unexamined Cepheid.[20]

In short, the astronomers' observations of Cepheids, in the absence of any further, *independent* background, sufficed to justify the astronomers in projecting the period-luminosity relation onto unexamined Cepheids. That the observations of Cepheids can do it alone runs contrary to Hume and his Bayesian heirs (such as Sober, in a passage I quoted earlier), who contend that past observation reports cannot confirm empirical predictions without some independent background opinions expressing the relationship between actual past and possible future observations. My argument denies the *independence* of the requisite background. Its alleged independence was crucial, since that was what raised the spectre of a regress, inviting us to demand additional observations to justify the background (inductively . . .).

Earlier, I explained the inductive sceptic's prior probability distribution, which was supposed to enable her to gather the data but not to project anything from them. In other words, the inductive sceptic's prior probability distribution was supposed to enable her to observe that various stars are Cepheids and to recognize that they conform to a certain period-luminosity relation, but was supposed to preclude her from regarding this fact as evidence that unexamined Cepheids will do likewise. I have just argued that this prior distribution is logically uninhabitable, because to observe that certain stars are Cepheids, an astronomer must already have the resources for going beyond those observation reports under certain conditions. The traditional problem of induction concerns how to go beyond our observation reports; it takes the reports themselves for granted. In contrast, I am arguing that there is no room to ask, "Granted that all Cepheids examined so far obey this period-luminosity relation, but what reason do we have for expecting unexamined Cepheids to do likewise?"

This question, on my view, is no more open than "Granted that the firefighter's action was courageous, but what reason do we have for regarding it as morally praiseworthy?" To grasp what is common to all courageous actions in virtue of which they are courageous, we must approach them from an evaluative standpoint; a purely naturalistic perspective will not suffice to make salient their distinctive character. Analogously, astronomers could classify the customary range of stars as Cepheids only by construing all of those stars as variables of the same physical type, and thus only by being prepared (in the absence of any other relevant information) to regard examined Cepheids as confirmationally relevant in certain respects to every other Cepheid.

IV

I have suggested that to uncover the basis of an inductive leap like Leavitt's and Shapley's, we must look behind the observation reports to various background beliefs that must accompany them. An inductive sceptic, on this view, must therefore abjure even from making the observations from which the inductive leap is launched.

It might be objected: Let's concede that astronomers could not have identified a given star as a Cepheid without committing themselves to the existence of (as yet undiscovered) natural laws of certain sorts covering all and only Cepheids, and thus to examined Cepheids being (under certain circumstances) confirmationally relevant (in certain respects) to unexamined Cepheids. But then astronomers could not ever really have *observed* a given star to be a Cepheid. Without reasoning inductively, astronomers could not have known that there are Cepheids—that the "Cepheids" form a natural kind of variable star. The more that is alleged to be packed into an observation report, the more remote that report becomes from what is truly accessible directly by observation.

In reply to this objection, I shall begin by invoking direct realism. Roughly speaking, direct realism holds that at least some of our observation reports must concern the "external" world, not sense data, experiences, inner representations, sense impressions, colored surfaces, or the way things look or appear. Our knowledge of the outside world is not mediated by our knowledge of any such interpolated, reality-neutral, phenomenal objects, states, or episodes. According to direct realism, we do not *infer* from the way things around us appear to the way they really are. Rather, after appropriate training, we become qualified to observe (in certain conditions) that facts of certain sorts obtain. Although qualified observers are not infallible, "[w]e don't begin to hedge [by retreating to talk of appearances] unless there is some special reason for doing so" (Austin 1962, 142). A report of something *looking* a certain way, while not implying that it really *is* that way, presupposes that there are conditions where (at least in this respect) things are the way they look—and so where we can observe them to be that way (Sellars 1963, 147). Observations that things *look F* presuppose that some of our observations are that things *are F*. Though a given neuron in the cochlea of our ears may be directly sensitive only to a certain pitch, it does not follow that a trained musician hears various combinations of pitches rather than, say, that the oboe is out of tune (Hanson 1958, 17). That causal priority dictates epistemic priority is a version of what Sellars (1963) calls "the myth of the given." An observation report is distinguished by its epistemic role (a normative matter), not its causal role. Though a given belief about how things are brings with it various expectations regarding how those things will behave later or would behave in other circumstances, it does not follow that we cannot know how things are in this respect merely by looking now. Direct realism rejects the view that all we really know immediately is something about the here and now:

> There are vast numbers of things which I take for granted that a telephone won't do ... Must I try to eat it, and fail, in the course of making sure it's a telephone? (Austin 1962, 122)

No, of course not, the direct realist maintains. I can see that it's a telephone. Admittedly, if it was actually a flashlight in the shape of a telephone, then I did not see that it's a telephone. Having learned of my mistake, I would retreat to "It looks like a telephone." But this retreat does not show that originally, I *inferred* that it's a telephone from its looking like a telephone (or from any facts about the character of my subjective sensory experience). When we were mistaken about its being a telephone, there weren't special phenomenal objects or qualities, mediating between us and the world around us, about which we were not mistaken.

For the sake of argument, I shall take direct realism for granted; I shall not offer any argument for it.[21] My concern is to examine how far direct realism, if correct, could take us in elaborating our nascent reply to the problem of induction. If "seeing a bird in the sky involves seeing that it will not suddenly do vertical snap rolls" (Hanson 1958, 121), then the observation of a light curve as Cepheid-type can just as well be "thick" with commitments to the existence of (as yet unknown) natural laws, and hence to the confirmational relevance of this star's features. If we cannot make observations without undertaking commitments that go beyond them, then apparently the inductive sceptic must abstain from making *any observations at all*. Thus the problem of induction is circumvented, since it presupposes that we could be entitled to make observations without having any reason to infer inductively beyond them.

Let's consider this argument more carefully. Direct realism holds that at least some of our observation reports must concern how things in the "external world" really are, not merely how they now appear to us to be. Observation reports can be fallible; that a report is not certain does not make it incapable of possessing non-inferential justification. Therefore, direct realism precludes the following argument: since a Cepheid report sticks its neck out by implying the existence of various sorts of natural laws, a Cepheid report cannot be known for certain to be true, and so is incapable of being an observation report.

However, this argument does not supply everything that our response to the problem of induction needs.[22] A direct realist must say that in making an observation report "*x* is *F*," the observer takes a risk, since she undertakes commitments regarding various unobserved states of affairs. In observing that something is a telephone, Austin says, one predicts (in the absence of further relevant information) some consequences of trying to eat it. However, direct realism does not obviously imply that we can make an observation report "*x* is *F*" that implicitly involves commitment to the existence of various sorts of natural laws (through commitment to *F*'s forming a particular sort of natural kind), and thus commitment to various features of *x* being confirmationally relevant (under certain circumstances) to various features of other *F*'s. Even if our observation reports must be "theory laden," there is no evident reason why they must be laden with the sort of theory that would suffice to underwrite this confirmational relevance. So direct realism does not entail that the inductive sceptic must abstain from making *any* observations at all. There are only *certain sorts* of observations that she is prohibited from making, such as observations that various stars are Cepheids. She could observe that *x* is *F*, though this commits her to certain predictions about *x*'s future behavior, as long as this observation does not commit her to regarding further observations of *x* as confirmationally relevant to other *F*'s.

Admittedly, I see no reason to believe that every observation report "*x* is *F*" must involve commitment to certain features of *x* being confirmationally relevant to certain features of other *F*'s. On the other hand, an inductive sceptic who embraces direct realism cannot argue, simply from the risk involved in induction, that no claim involving inductive commitments (such as "*x* is a Cepheid") can be an observation report, since according to direct realism, some observations must be risky anyway. Why should inductive risks be specially prohibited?

But in the remainder of this section, I shall argue for a stronger point: that observations with inductive commitments ("taxonomic observations," we might call them) are among the sorts of observations that we can most readily make when we have very little relevant background knowledge.

Though an observation report has *noninferential* justification, direct realism says that the subject matter of observation reports is not of some special sort; observation reports are not distinguished by what they are about. In particular, an observation report's content does not automatically prevent the observer from justifying the claim *inferentially* to someone else. We saw, for example, that even if an astronomer observes a star to be a Cepheid by looking at its light curve, she may also be able to support that classification inferentially—for example, by arguing that the star's light curve is connected to the light curves of prototypical Cepheids by a smooth progression of intermediate cases. Some direct realists, such as Sellars, go further by arguing that *any* claim that you make as an observation report you *must* be able to justify inferentially to someone else or to yourself. Of course, that inference must proceed from other observations that you have made. As Sellars says, "empirical knowledge ... is a self-correcting enterprise which can put *any* claim in jeopardy, though not *all* at once" (1963, 170).

In particular, Sellars (1963, 167–70) contends that you are justified in observing "That is an *F*" only if you can justly argue, from your track record of making similar observation reports in similar circumstances, that your report on this occasion is probably true. In other words, Sellars holds that you must know that you have at various times (under various conditions) responded to your environment by maintaining some things to be *F*'s, that most of those things were indeed *F*'s, that you are now (under similar conditions) responding to your environment by maintaining some thing to be an *F*—and that in view of these premises, it is probably an *F*. Thus, you can step back and critically evaluate your conditioned responses, by adopting the very same perspective on them that another person could adopt. Sellars's view seems reasonable to me.[23] If you do not have a reason to trust yourself—if you are not in a position to infer the probable accuracy of your report—then on mature reflection, you should find that you are not in a position to endorse your report.[24] You should disavow it as having been made too hastily, merely as a kind of knee-jerk reaction.

This inference from your track record is obviously an inductive inference, as Sellars emphasizes.[25] How does the evidence for this inference suffice to confirm its conclusion? Once again, background knowledge is needed to connect the premise to the conclusion. But where is the requisite background knowledge to come from in cases where our relevant background is very meager?

We faced exactly this sort of question in the previous section. My answer then was that the observation report "That is a Cepheid" is inextricably linked to the background belief that Cepheids form a species of variable star, and that this background belief suffices (in the absence of further information) to render certain features of the observed star confirmationally relevant to certain features of any unexamined Cepheid. The same answer applies to the inductive inference from your track record. A premise of that inference is that most of the things that you have purported to identify as Cepheids from their light curves really are Cepheids, which requires that Cepheids really be a species of variable. But for Cepheids so to be, they must have in common certain (unknown) intrinsic properties that are essential to being a Cepheid, the combination of which generates (via the natural laws) the Cepheids' characteristic properties, such as their pattern of light variation. Since different species of variable star have different essences (i.e., different mechanisms producing their light variation), which are responsible for their characteristic properties, a species' characteristic properties are (with relatively rare exceptions) also unique to that species (just as one chemical species typically does not share with another its exact boiling point, ionization potential, and so forth). A star that produces a characteristically Cepheid light curve is therefore probably a Cepheid variable. Of course, there is no guarantee that this all will be so, even if the Cepheids form a natural kind of variable star (see note 16). But it is likely, if they are a natural kind, and this background belief enables your past accuracy in distinguishing Cepheids from non-Cepheids, by virtue of their light curves, to support your future accuracy in doing so. The best explanation of your past accuracy is not that it was a fluke—that you happened to check a non-representative sample of stars, which accidentally included only Cepheids with rather similar light curves and non-Cepheids with rather different light curves.

The induction from your track record is able to go through, despite the paucity of your background knowledge, because the theory that suffices to underwrite it is not independent of the induction's premise, but required by it. Taxonomic observations, such as "*x* is a Cepheid," are among the sorts of observations that scientists are most likely to be able to make when their background theory is impoverished. Let's illustrate this with another example. My family and I often visit the dog park. Suppose that on these trips, we begin to notice some dogs that look very much alike. Suppose we take ourselves to have noticed a breed of dog previously unfamiliar to us, and we make observations accordingly. ("Honey, I just saw another one of those dogs we were talking about; I still don't know what they're called, but they sure are cute.") I can give a reason for believing myself now capable of distinguishing dogs of this breed on sight: I have been pretty reliable in the past (as my family can attest), and as this is supposedly a genuine breed, each of its members must share considerable common ancestry, and no dog descended entirely from those ancestral individuals lies outside of the breed. In virtue of this shared history, all of the breed's members (at least in the relatively recent past and future), whether observed or as yet unobserved by me, probably look very much alike and rather different from other dogs. Any dog looking like these is unlikely to have been produced by a different lineage. Therefore, my past reliability in distinguishing members of this breed on sight is good evidence for my future reliability in doing so. The background knowledge connecting the premise and conclusion of this inference is implicit in the premise, since the premise is that these responses of mine have been pretty accurate, and their truth requires that we have here a genuine breed of dog.

In contrast, suppose that I can also tell, by gazing upon the culprit through our bedroom window, whether the dog barking outside is a neighborhood dog or a stray from afar, presumably lost. We can take my past accuracy in making these identifications as good evidence for my accuracy in some new case—but only if we have good reason to believe that my past accuracy was no fluke. For example, we must justly believe that I have already seen most of the neighborhood dogs, as would not be the case if (for example) there is rapid turnover in their population, and we must also justly believe that a potential stray is unlikely to look much like any of the neighborhood dogs. We would need *independent* empirical evidence for these premises before I could be qualified to observe whether or not a dog is from the neighborhood. In this way, it is generally much easier, in the absence of rich background knowledge, for someone to become qualified to observe that x belongs to "that new breed of dog we've been seeing" (a putative "natural" kind, of some sort) than for someone to become qualified to observe that x is one of the neighborhood dogs.

Accordingly, among the sorts of observations that tend to come first in a new field, where little background knowledge can be brought to bear, are often "taxonomic observations"—identifications of various things as members of various species.[26] Of course, there could be other kinds of variable stars, examples of which we have not yet encountered, that exhibit Cepheid-type light curves, just as there could be other breeds of dog that look like the one that we have just begun to recognize. But the different essences of different dog breeds or variable-star species usually demand different characteristic properties (see note 15). There is no such assurance in the case of neighborhood dogs.

V

But the question remains: Why were astronomers justified in believing that "Cepheids" form a natural kind of variable star[27]—or even that there are any such kinds having light-curve shape as a characteristic property? What reasons did astronomers have for believing it likely that some "common cause" is responsible for every "Cepheid" light curve? If the observations made at the outset of scientific work in some area, when there is a dearth of relevant background knowledge, must be laden with enough theory to get inductions going, then how do we ever become entitled to make those observations? Perhaps astronomers should have confined themselves to compiling the light curves of various stars, without grouping any of the stars together into species (and so without being able to propose—much less to confirm—candidate laws such as the Cepheid period-luminosity relation).

Let's think about this. To compile a star's light curve, astronomers must recognize the same star on several occasions—sometimes after not having seen the star for many nights (if the weather interfered), or for many months (if the star was hidden behind the Sun), or even for many years (if no one happened to look for it). On what grounds is a star re-identified? Its exact location on the celestial sphere does not suffice to distinguish it, since parallax, aberration, and a star's own motion relative to Earth may change its position (to a small but significant degree) from where it used to be seen.

Obviously, this problem is especially acute in the case of variable stars. A nonvariable star's magnitude and color will be constant, suggesting that we are now seeing the same object as we saw before at that approximate location. If one night, a star that looked like Rigel (a very bright, blue-white star in Orion) were found in place of Albireo (a medium-bright, double star in Cygnus, of unusual topaz and sapphire components) and vice versa, we might well be justified in hesitating to make any re-identifications. But a variable star *does* look like a different star on different occasions. In 1603, before any variable stars had been recognized, Johann Bayer in his famous catalogue labelled a star "omicron Ceti" without realizing that it was identical to a "nova" (new star) that had appeared in Cetus in 1596 and been observed to fade from view a few weeks later. It was not until 1639, after omicron Ceti had (again) disappeared and reappeared, that astronomers recognized it as a variable star. (Its period averages 333 days.) This kind of error was as common as it was understandable. More than two hundred years later, and despite occupying a prominent position in the sky, a bright variable in Orion was still being misidentified as a nova, having often escaped notice by an unlucky coincidence between its minima and the times of the year at which it was best situated for observation (its period being just about a year).

How, then, were variables discovered? Not from any constancy in their light or exact fixity in their position. (Omicron Ceti moves the not inconsiderable extent of 25 arc seconds per century.) The evidence was the repeated pattern in the observations. Because the same pattern of luminosity change was seen, the same object was believed to be responsible. For example, that the 1887 "nova" in Orion was preceded by another roughly 373 days before, in roughly the same location, and by another roughly 373 days before that, and each faded in a similar manner (figure 6), strongly suggested that these occasions all involved the repeated behavior of the same object (then designated U Orionis).

So before astronomers had amassed considerable background knowledge of variables, a variable was discovered *only together with* at least a rough pattern in its light curve being recognized. But what counted as such a "pattern"—as a "light curve" continuing *in the same way*, suggesting that the same object was involved throughout? Not all variables have light curves as regular as the Cepheids do (see figure 7), and even many Cepheids have undergone sudden changes in their periods (while their light curves have remained Cepheid-type). For astronomers to have justly regarded various stellar observations as probably falling on a single light curve—as probably all produced by the same type of mechanism, and hence probably by the same object—astronomers needed to depend on their sense of

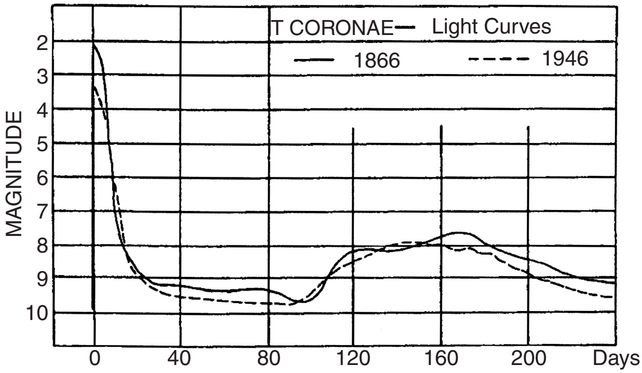

Figure 6. Light curves of T Coronae (a star of the same class as U Orionis) on successive maxima (Burnham 1978, II, 709).

what sort of differences are not drastic enough to suggest that a different type of mechanism may be responsible. In other words, astronomers needed to be thinking in terms of natural kinds of variable star even merely to justify drawing light curves, and hence to justify believing that there are variable stars.

As astronomers learned more about the types of variable star there are, they could more confidently (and on the basis of fewer observations) infer a variable's existence. Eventually, astronomers could infer a star to be variable from spectral and other symptoms, even before noting any change in its

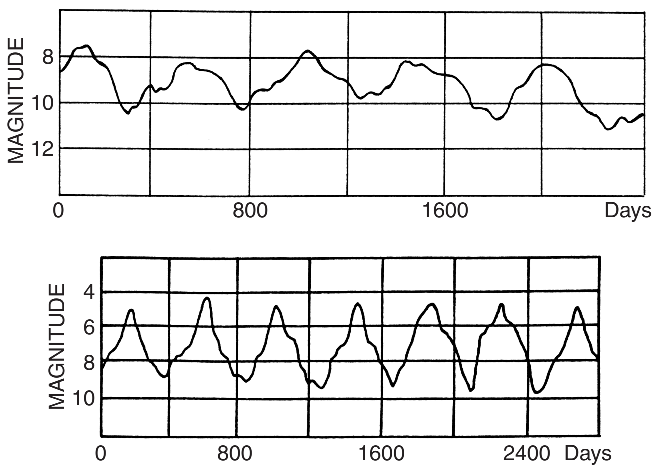

Figure 7. The light curves of S Cephei (above) and R Hydrae (below). Each star's own periods resemble one another, not unlike the family resemblance among all of the Cepheids shown in figure 5.

brightness over time. And long before that, astronomers could *observe* "That's omicron Ceti," re-identifying a variable on sight. But at the outset, astronomers had to rely upon their beliefs about when past and present observations formed a sufficiently consistent pattern that they could justly be ascribed to the same type of mechanism, and hence probably to the same object, and when they could not. For example, in May 1860, a seventh-magnitude star appeared to replace the nebula M80. Was the nebula now a star, the same object having changed (a "variable nebula"), or had a new star, distinct from the nebula, appeared and overpowered the nebula's feeble light? That the two were at roughly the same location in the sky, at least as seen from Earth, did not suffice to suggest that the nebula had become a star. The difference between a star and a nebula was too drastic for astronomers to be in a position justly to conclude that the star they were now observing was the same object as the nebula they had observed before. Astronomers could draw no conclusions—until the star began to fade, revealing that the nebula had remained there throughout.

The belief that there are natural kinds of variable star is not an optional addition to the discovery of variable stars in the first place. To have variables to classify at all, astronomers needed to recognize various sequences of observations as reflecting the persistence of the same type of behavior, and hence probably having the same object as a common cause. Astronomers therefore needed to be recognizing types of variable star, with characteristic light curves. Astronomers could not have confined themselves to assembling the light curves of various stars, since without the background beliefs enabling astronomers to recognize two stars' light curves as belonging to the same type, astronomers could not have recognized two sequences of observations as cases of seeing the same individual star.[28]

We have returned to our main theme: the intricate connections between observation reports and the background beliefs that underwrite inductive inferences from them. I am not claiming that all inductions in science are mediated by theoretical background that is implicit in the data. Observation reports are not *that* theory laden. Rather, in theoretically rich contexts, the background opinions responsible for a given observation's confirmatory impact on some hypothesis are often quite independent of that observation. Indeed, that background may well have been arrived at through considerable prior empirical work. The curious thing about grand inductions made in theoretically impoverished contexts, such as the induction to the Cepheid period-luminosity relation, is that they seem to proceed without any empirical work being done in advance to establish the requisite background beliefs. Grand inductions in theoretically impoverished contexts seem to be too easy to make. To account for this, I have suggested that the theoretical background required for the observations to be made in the first place may itself suffice to underwrite the observations' confirmatory power. The traditional problem of induction then disappears; our problem is no longer to explain how we can be justified in projecting beyond our observations rather than sticking merely with what we have observed. Nevertheless, there remains the task of explaining how we come to be qualified to make those observations, or any observations at all.

Notes

[1] For more on how to refine the notion of *h*'s *inductive* confirmation as the confirmation of *h*'s predictive accuracy, see my (1999a, 2000a).

[2] Thanks to Colin Howson for comments in response to an earlier draft of this section.

[3] Thanks to Samir Okasha, Stephen Glaister, and Alan Hajék for discussion of this point.

[4] The prevailing theory was that Cepheids were binary stars, one component eclipsing the other. Shapley in 1914 and Eddington in 1919 suggested (correctly) that Cepheids pulsate. Other hypotheses included starspots passing by as the stars rotate, meteor impacts, the star's passing through galactic clouds, and tidal disruptions. All were speculative at best.

[5] Background knowledge played a further role. Quite apart from the question of how far Leavitt's relation among certain SMC variables can justly be inductively generalized, there is the vexed issue of how to calibrate Leavitt's relation. To compute a Cepheid's distance using the period-luminosity relation, one needs to know its period (from its light curve), from which the relation determines its intrinsic luminosity; by comparing its intrinsic luminosity to its brightness as seen from Earth, one can reasonably estimate its distance. (One must assume, for example, that negligible light is absorbed by dark matter between the Cepheid and us.) But Leavitt did not know the intrinsic luminosity of any SMC Cepheid, since she did not know the SMC's distance. Thus, she knew the period-luminosity relation's slope, but not its y-intercept. Leavitt had "hoped . . . that the parallaxes of some variables of this type may be measured" (1912, 3), and their distances thereby ascertained. (Note Leavitt's implicit claim: that these variables are all examples of the same distinct physical *type* of star, other examples of which may be near enough for their parallaxes to be measured.) But no Cepheid is near enough. To calibrate the relation, therefore, Hertzsprung and Shapley proceeded statistically. Here is roughly what they did. A star's relative speed towards or away from us can be determined spectroscopically, by exploiting the Doppler shift in the star's spectral lines. Hertzsprung and Shapley reasoned that, for the stars in any reasonably large sample, their *average* speed relative to Earth and *perpendicular* to our line of sight probably roughly equals their *average* relative speed *along* our line of sight. So for a sample of roughly a dozen Milky-Way Cepheids, Hertzsprung and Shapley could infer the stars' average relative speed perpendicular to our line of sight. A star's relative motion perpendicular to our line of sight may (given sufficient time) measurably change its position in the night sky. Thus, the star's relative speed perpendicular to our line of sight can be measured directly, in seconds of angular arc per century. (For this reason, Hertzsprung and Shapley selected the dozen or so Cepheids having the best determined relative speeds perpendicular to our line of sight.) From the selected Cepheids' average relative speed perpendicular to our line of sight, along with the average number of arc seconds per century that this relative motion has produced, Hertzsprung and Shapley could compute the stars' average distance. (There is exactly one distance at which a given length subtends a given angle.) Then from the stars' average brightness as seen on Earth, Hertzsprung and Shapley inferred their average intrinsic luminosity. This, taken with the stars' average period, gave Hertzsprung and Shapley a point with which to calibrate their period-luminosity relation. Hertzsprung's result was that the absolute magnitude $= -2.1 \log$ (period in days) – 0.6. (Bear in mind that the higher a star's "magnitude," the dimmer it is.) Shapley arrived at $M = -2.80 \log (P) - 1.43$. There was considerable disagreement over the plausibility of the calibration achieved by this method. (See Smith 1982.)

[6] There also remained plenty of logical space for an astronomer justly to believe that a period-luminosity relation holds for all Cepheids, but that no satisfactory method had yet been found to calibrate this relation (see note 5). I shall focus not on the controversial calibration method employed by Hertzsprung and Shapley, but just on the considerably less controversial view (in the wake of Leavitt's discovery) that a period-luminosity relation covers all Cepheids. That's the theoretically impoverished step—generalizing from Leavitt's small sample of SMC stars to all Cepheids.

[7] Although the earliest use in print (that I could find) of the term "Cepheid" is by Clerke (1903, 319), a "family likeness" (Roberts 1895, 284) among these light curves had been widely appreciated, and these stars were commonly characterized as "variable stars of the same class" (Campbell 1901, 95). Leavitt also referred to a "type" here; see note 5.

[8] By Leavitt's time, astronomers had also identified a few spectral features characteristic of Cepheids. However, they apparently did not need to discover that a given star possesses these spectral characteristics in order to be in a position to identify it as a Cepheid. Furthermore, since astronomers ascertained *inductively* that all Cepheids possess these spectral features, astronomers must already have justly concluded that examined Cepheids are relevant confirmation-wise to the others.

[9] Today, astronomers regard Shapley and his colleagues as having mistakenly grouped together several different kinds of stars ("classical" Cepheids, RR Lyrae stars, W Virginis stars) and hence as having arrived at estimates of stellar distances that, although far more accurate than previous ones, were several times too low. Yet astronomers today generally do not deny that Shapley and his colleagues were justified at the time in believing Cepheids to be uniform, and this justification is our concern.

[10] Hoskin (personal communication) says that the justification for the uniformity assumption is that doing something is better than doing nothing. However: (1) this principle does not obviously hold if doing something may well lead to false conclusions, whereas doing nothing would obviously avoid that result. Analogy: The only way to have a chance to win the lottery is to play, but it does not follow that we ought to play; it depends on the odds of winning and the cost of losing. (2) That doing something is better than doing nothing cannot justify doing one "something" rather than another. But surely, not all assumptions that would lead to stellar distance estimates were equally justified. (3) "Doing" something in the sense of pursuing a research program may indeed be better than doing nothing, but this "doing" does not oblige us to believe the conclusions yielded by that program until further research reveals evidence justifying them.

[11] Of course, any two light curves will have infinitely many respects of similarity and infinitely many differences. But some of these must be more salient to astronomers than others, else astronomers could never have caught on to "Cepheid" ostensively.

[12] For example, Loud (1907, 370) remarked that although many Cepheids rapidly rise to their maximum and then more gradually fall, "[a] few instances in which this peculiarity was deemed to be replaced by symmetry have been erected on that ground by some authorities [such as Clerke (1903, 328), though not Campbell (1901, 95)] into a separate species, having zeta Geminorum for a [proto]type." But Loud objected that there is a continuous series of cases between zeta Geminorum and more recognized Cepheids:

> [T]he usual asymmetry is in no way fixed in degree, varying in sundry instances much below the mean; thus in W Virginis the time of increase is to that of decrease as 46 to 54; while S Antliae . . . exhibits a corresponding ratio of 62 to 38, thus for once overpassing the limit of symmetry. These stars must then be regarded, at least provisionally, as merely aberrant members of the class represented by delta Cephei. On the other hand, the type represented by beta Lyrae is entirely distinct . . .

In contrast, Hertzsprung objected to co-classifying the cluster variables with the (other) Cepheids on the grounds that the latter had periods longer than a day, the former shorter, and there was a significant gap in the population with periods between 22 hours and 2 days

(Smith 1982, 72; Campbell and Jacchia 1946). Curtis evocatively called them "two different breeds of cats" (Smith 1982, 124).

[13] Astronomers generally do not use the philosophical term "natural kind." But they do frequently characterize some category as a "natural group," a "real physical group," or "a very well defined class of objects." (See also notes 7 and 12.)

[14] The account of a "natural kind" that I am about to sketch is hardly original. Similar ideas can be found in works from Jevons (1905, 675 and 708) to Kuhn (2000, 228–33).

[15] Or (to put it the other way around) that two alleged species are found to agree in some characteristic property is *prima facie* grounds for wondering whether they are indeed distinct species. For example, that the tau and theta-zero mesons were found (in the mid-1950's) to have similar mass was considered powerful evidence that they are actually the same particle (Pais 1986, 516).

[16] These are the expectations that (as we shall see) ground the inference to the Cepheid period-luminosity relation. But they are not the only expectations. In observing various stars to be Cepheids (or non-Cepheids) from inspecting their light curves, astronomers presupposed that a star *can* be identified as a Cepheid (or non-Cepheid) from its light curve alone. But Cepheids cannot really be so identified if, for instance, the gap that astronomers noticed between Cepheid and non-Cepheid light curves was merely an artifact of astronomers' having checked an unrepresentative sample of stars. (In that event, the gap might later have gotten filled in as astronomers discovered further variable stars and plotted their light curves.) Thus, in observing various stars to be Cepheids and non-Cepheids from inspecting their light curves, astronomers committed themselves to the view that the light-curve gap between examined Cepheids and examined non-Cepheids exists between all Cepheids and all non-Cepheids (and so will persist as further variable stars are taken into account, if their light curves are correctly plotted), since otherwise, observation of the light curve alone is not a good way to identify a Cepheid.

This is not to say that if the gap is merely an artifact of an unrepresentative sample, then there could not really be Cepheid variables, i.e., that the so-called Cepheids could not form a natural kind. The stars later discovered as filling in the gap might all be distinguished from the genuine Cepheids by their mechanism of variation and, hence, in other respects, such as spectroscopically. But then a Cepheid could not actually be identified from its light curve alone, even if astronomers did by sheer good fortune initially manage to distinguish some Cepheids in that way. Indeed, Newcomb (1902, 108–9) claimed that "[t]he [light-curve] gap between the variable stars of the Algol type and those of the Beta Lyrae type is at the present time being filled by new discoveries in such a way as to make a sharp distinction of the two classes difficult." But this is a merely epistemic problem, he said; the two classes remain distinct. He described how we could "have a star of the Algol type so far as the law of variation is concerned, yet, as a matter of fact, belonging rather to the Beta Lyrae type" in virtue of the mechanism responsible for the star's variation. Newcomb explained how no measurement then available would enable such a star to be properly classified.

The report that a given star is a Cepheid also embodies expectations regarding that star's future, since the report is not that the star is *currently* a Cepheid. The report presupposes that such stars, were they observed again, would nearly always be found still to be exhibiting a Cepheid light curve. A star's status as a Cepheid is, for practical purposes, one of its permanent features. (Of course, stars presumably change over *astronomical* time scales. RU Camelopardi is, to my knowledge, the only known example of a star that was once observed to be a Cepheid and is now no longer.)

[17] "Understanding why just those things belong together may essentially require understanding the supervening [evaluative] term" (McDowell 1981, 145); "We do not fully understand a virtuous person's actions—we do not see the consistency in them—unless we [have] a grasp of his conception of how to live (McDowell 1979, 346). See also Lange (2000c).

[18] This is an ironic result. From the 1880's to the 1950's, a corps of women (including Henrietta Leavitt) was employed, especially at Harvard College Observatory, to identify

variable stars. Of the 12,500 variables discovered at Harvard between 1886 and 1956—out of the 14,700 variables listed in the General Catalogue of Variable Stars—80% had been found by women (Hoffleit 1993, 8 and 12). One of the principal reasons for employing women was that their minds were thought to be uncontaminated by any scientific theory, enabling them to classify variables impartially. Regarding Annie Cannon, who discovered several hundred variable stars as well as devising the OBAFG Harvard system of spectral classification (which has remained standard), a contemporary (Payne-Gaposchkin 1941, 63) wrote: "Miss Cannon was not given to theorising; it is probable that she never published a controversial word or a speculative thought. That was the strength of her scientific work—her classification was dispassionate and unbiassed." (Galison (1997) has noted similar attitudes toward women employed as photograph scanners in high-energy physics laboratories.) While their observations may indeed have been neutral among various theories regarding the mechanisms by which a star's luminosity may vary, I have argued that their observations were not wholly uncontaminated by theory. They carried on as they did, from the canonical cases, only because they knew that their categories were supposed to constitute genuine stellar species.

[19] It might be objected that in observing various stars to be Cepheids, astronomers were not being guided by their understanding of Cepheids as constituting a natural kind of variable. Rather, they went on from the prototypes in the manner they did because they understood the Cepheids simply as whatever variables fall between the prototypes and the nearest gap in the stellar population density; they held that by definition, a star is a Cepheid if and only if there is no gap between its light curve and delta Cephei's but there are gaps between its light curve and Algol's, beta Lyrae's, Mira's, and so forth. However, although this interpretation would account for the range of cases that astronomers designated "Cepheid," this view interprets a star's Cepheid character as depending not just on its own light curve, but also on the distribution of other stars' light curves. In other words, a star's Cepheid character would not be among its intrinsic properties; had there been additional stars filling in the gaps, for instance, there would automatically (on this view of Cepheid-hood) have been no Cepheids (compare note 16), although the intrinsic properties of various actual Cepheids would have been unchanged. Astronomers certainly treated a star's Cepheid character as one of its intrinsic properties, alongside the star's period of variation and intrinsic luminosity "as determined by [its] mass, density, and surface brightness" (Leavitt 1912, 3), but unlike its apparent brightness on Earth. For astronomers to have gone on as they did in applying "Cepheid" to new cases, while meaning anything like what they did, astronomers must have regarded the Cepheids as a species of star.

[20] There is a good deal more to be said about the belief that there are various sorts of Cepheid laws, prescribing various features having something to do with the Cepheids' luminosity, and about the relation between having this belief and regarding examined Cepheids as relevant confirmation-wise to every unexamined Cepheid. The traditional view is that candidate laws of nature can be confirmed by their instances whereas candidate accidental generalizations cannot. For example, since we believe "All emeralds are green" to be a natural law, if true, we take the discovery of a given emerald to be green as confirming, of any unexamined emerald, that it is green. But, since we believe "All of the families on my block have two children" to be an accident, if true, we do not regard my family's having two children as confirming the same of the Jones family next door. I have shown (1999a, 2000a) that although this traditional view is incorrect as I have just elaborated it, there is a kernal of truth to it: a claim that we believe to be accidental, if true, cannot be confirmed "inductively" in a special sense. Moreover, in regarding emeralds to be a natural kind (of a certain sort), we believe certain "inductive strategies" to be the best, and therefore we adopt them. These strategies have us take (certain sorts of) features of examined emeralds as relevant confirmation-wise to (certain sorts of) features of any unexamined emerald.

[21] For a recent discussion, see Putnam 1999. Bayesians should find direct realism quite congenial. Indeed, Jeffrey's rule codifies one aspect of the "theory-ladenness" (Hanson 1958, 19) of our observations: our priors influence the new opinions that our perceptual experiences directly bestow upon us, and hence influence *what we have experienced things as*. (See my 2000b.)

[22] My thanks to a referee for pressing this objection.

[23] It is not entirely unproblematic, however. One might well ask how the observer arrives at her knowledge of her past. If a bit of this knowledge, as a recollection, is to have noninferential justification, then by Sellars's lights, the observer must be able to support it from her track record of recollecting things, an argument that in turn depends on further premises about the past. I shall not investigate here how Sellars purports to avoid such potential difficulties—except to mention that, according to Sellars, "one could not have observational knowledge of *any* fact unless one knew [empirically] many *other* things as well" (1963, 168) and no bit of the foundational stratum of one's empirical knowledge is epistemically prior to any other bit.

[24] For more general remarks on having to trust yourself epistemically, see my (1999b).

[25] So on Sellars's view, making observations is bound up with using induction; induction does not enter the picture only after the data have already been collected. An inductive sceptic would be unable to observe anything. Although I think that Sellars is correct here in his reply to the inductive sceptic, it is not clear that much science can arise merely from the sorts of inductions that Sellars says even the inductive sceptic must countenance (on pain of not being able to make observations at all). Hence the burden of this paper.

[26] That a Cepheid is expected to remain a Cepheid for a long time (see note 16) makes it easier for you and me to check what you have seen, and hence for you to build up a track record of Cepheid reports that were revealed to be accurate. If you and I cannot look at the thing on which you have already reported, or if that thing may well in the meantime have changed from the condition that you reported it as being in, then it will be more difficult for you and me to check what you have reported, in the absence of theoretical background that would enable currently unobservable conditions to be inferred from what is currently observable. Here is a further respect in which Cepheid identifications are especially amenable to being made as observation reports despite a theoretically barren context.

[27] At the close of section III, I said that astronomers cannot grant that various stars are actually Cepheids, and then ask what basis there is for regarding the Cepheids as a real class of variable. But astronomers *can* step back from their own practice and ask why they are entitled to regard the *so-called* "Cepheids" as a natural kind.

[28] A fan of Goodman's new riddle of induction would note that no matter what a star's future magnitude (say) might be, there will be a characteristic of its light that remains unchanged, since a suitable grue-type predicate can be manufactured. For example, a star of magnitude 5 before the year 2005 and magnitude 7 after 2005 has a constant schmagnitude 5*7. Goodman's general point, of course, is that any behavior could be made to count as nature's going on in the same way; the principle of the uniformity of nature is empty without some further specification of what counts as a uniformity. But if (as I have argued) astronomers possessed background beliefs about like patterns of starlight having like causes, then presuming this background belief to be genuinely contentful rather than utterly empty, Goodman's point entails that astronomers must have had some background beliefs limiting what kinds of behavior could count as the star's going on in the same way. (These beliefs would have been revised as more variables were discovered.) Goodman's new riddle of induction arises only if some purely formal, syntactic relation is alleged to determine which evidence would inductively confirm a hypothesis.

References

Austin, J. L. (1962) *Sense and Sensibilia* (Oxford: Clarendon).

Brueckner, Anthony. (2001) "BonJour's A Priori Justification of Induction," *Pacific Philosophical Quarterly* 82: 1–10.

Burnham, Robert. (1978) *Burnham's Celestial Handbook* (New York: Dover).

Campbell, Leon and Jacchia, Luigi. (1946) *The Story of Variable Stars* (Philadelphia: Blakiston).

Campbell, W. W. (1901) "The Motion of zeta Geminorum in the Line of Sight," *Astrophysical Journal* 13: 90–7.

Clerke, Agnes. (1903) *Problems in Astrophysics* (London: Adam and Charles Black).

Curtis, Heber D. (1921) "The Scale of the Universe," *Bulletin of the National Research Council* 2 (part 3): 194–217.

Fine, Terrence. (1973) *Theories of Probability* (New York: Academic).

Galison, Peter. (1997) *Image and Logic* (Chicago: University of Chicago Press).

Hacking, Ian. (2001) *An Introduction to Probability and Inductive Logic* (Cambridge: Cambridge University Press).

Hanson, Norwood Russell. (1958) *Patterns of Discovery* (Cambridge: Cambridge University Press).

Hetherington, Norris. (1996) *Hubble's Cosmology* (Tucson: Pachart).

Hoffleit, Dorrit. (1993) *Women in the History of Variable Star Astronomy* (Cambridge: AAVSO).

Hoskin, Michael. (1976) "The 'Great Debate': What Really Happened," *Journal of the History of Astronomy* 7: 169–82.

———. (1982) *Stellar Astronomy: Historical Studies* (Chalfont St. Giles: Science History Publications).

Howson, Colin. (2000) *Hume's Problem: Induction and the Justification of Belief* (Oxford: Clarendon).

Howson, Colin and Urbach, Peter. (1989) *Scientific Reasoning: The Bayesian Approach* (LaSalle: Open Court).

Jevons, W. Stanley. (1905) *The Principles of Science* (London: Macmillan).

Kornblith, Hilary. (1993) *Inductive Inference and its Natural Ground* (Cambridge: MIT).

Kuhn, Thomas. (2000) "Afterwords," in *The Road Since Structure* (Chicago: University of Chicago Press), pp. 224–52.

Lange, Marc. (1999a) "Why are the Laws of Nature So Important to Science?", *Philosophy and Phenomenological Research* 59: 625–52.

———. (1999b) "Calibration and the Epistemological Role of Bayesian Conditionalization," *The Journal of Philosophy* 96: 294–324.

———. (2000a) *Natural Laws in Scientific Practice* (New York: Oxford University Press).

———. (2000b) "Is Jeffrey Conditionalization Defective By Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experience," *Synthese* 123: 393–403.

———. (2000c) "Salience, Supervenience, and Layer Cakes in Sellars's Scientific Realism, McDowell's Moral Realism, and the Philosophy of Mind," *Philosophical Studies* 101: 213–51.

Leavitt, Henrietta. (1908) "1777 Variables in the Magellanic Clouds," *Harvard College Observatory Annals* 60: 87–108.

———. (1912) "Periods of 25 Variable Stars in the Small Magellanic Cloud," *Harvard College Observatory Circular* 179: 1–3.

Loud, F. H. (1907) "A Suggestion Toward the Explanation of Short Period Variability," *Astrophysical Journal* 26: 369–74.

McDowell, John. (1979) "Virtue and Reason," *The Monist* 62: 331–50.

———. (1981) "Non-Cognitivism and Rule-Following," in S. Holtzman and C. Leich (eds.), *Wittgenstein: To Follow a Rule* (London: Routledge), pp. 141–62.

Newcomb, Simon. (1900) *Elements of Astronomy* (New York: American Book).

———. (1902) *The Stars* (London: John Murray).

Okasha, Samir. (2001) "What Did Hume Really Show About Induction?", *The Philosophical Quarterly* 51: 307–27.

Pais, Abraham. (1986) *Inward Bound* (Oxford: Clarendon).

Payne-Gaposchkin, Cecilia H. (1941) "Miss Cannon and Stellar Spectroscopy," *The Telescope* 8: 62–3.

Putnam, Hilary. (1999) *The Threefold Cord* (New York: Columbia University Press).

Roberts, Alexander. (1895) "Close Binary Systems and their Relation to Short Period Variation," *Astrophysical Journal* 2: 283–92.
Salmon, Wesley. (1967) *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press).
Sankey, Howard. (1997) "Induction and Natural Kinds," *Principia* 1: 239–54.
Sellars, Wilfrid. (1963) "Empiricism and the Philosophy of Mind," in *Science, Perception, and Reality* (London: Routledge), pp. 127–96.
Shapley, Harlow. (1918) "On the Determination of the Distances of Globular Clusters," *Astrophysical Journal* 48: 81–116.
Smith, Robert. (1982) *The Expanding Universe* (Cambridge: Cambridge University Press).
Sober, Elliott. (1995) *Core Questions in Philosophy*, 2nd ed. (Englewood Cliffs: Prentice-Hall).
Van Cleve, James. (1984) "Reliability, Justification, and the Problem of Induction," in Peter French, Theodore Uehling, and Howard Wettstein (eds.), *Midwest Studies in Philosophy* IX (Minnesota: University of Minnesota Press), pp. 555–67.
Williams, Bernard. (1985) *Ethics and the Limits of Philosophy* (Cambridge: Harvard University Press).