OBSERVATION AND INDUCTION[1]

Theodore J. EVERETT

ABSTRACT: This article offers a simple technical resolution to the problem of induction, which is to say that general facts are not always inferred from observations of particular facts, but are themselves sometimes defeasibly observed. The article suggests a holistic account of observation that allows for general statements in empirical theories to be interpreted as observation reports, in place of the common but arguably obsolete idea that observations are exclusively particular. Predictions and other particular statements about unobservable facts can then appear as deductive consequences of such general observation statements, rather than inductive consequences of other particular statements. This semantic shift resolves the problem by eliminating induction as a basic form of inference, and folding the justification of general beliefs into the more basic problem of perception.

KEYWORDS: observation, induction, problem of induction, Karl Popper, hypothetico-deductive, foundationalism

In this article, I offer a simple technical resolution to the problem of induction, which is to say that general facts are not always inferred from observations of particular facts, but are themselves sometimes defeasibly observed. I suggest a holistic account of observation that allows for general statements in empirical theories to be interpreted as observation reports, in place of the common but arguably obsolete idea that observations are exclusively particular. Predictions and other particular statements about unobservable facts can then appear as deductive consequences of such general observation statements, rather than inductive consequences of other particular statements. This semantic shift resolves the problem by eliminating induction as a basic form of inference, and folding the justification of general beliefs into the more basic problem of perception.

In the first section of the paper, I analyze the problem of induction in terms of five jointly inconsistent propositions, of which the weakest is the statement that

[1] I would like to thank Earl Conee, Richard Fumerton, Alan Sidelle, Elliott Sober, audiences at SUNY-Geneseo and the Creighton Club, and several anonymous reviewers for helpful comments on earlier versions of this article.

all observations are particular rather than general. In the second section, I complain about the standard particularistic theory of observations, which depends on a cluster of assumptions that are commonly taken for granted, but that deserve little support in the light of recent progress in philosophy. In the third section, I give a brief sketch of a possible holistic account of observations, and show how it might work as a positive solution to the problem. I suggest that a main weakness in the classical hypothetico-deductive model of scientific reasoning can be removed if at least some hypotheses can be seen as defeasible observations of general facts.

Let me be clear about what I think I can establish. My primary concern is to point out that there is a possible new approach to the problem of induction in terms of general observations – an approach that ought to be considered, but is somehow missing from the standard treatments of the issue. My secondary concern is to argue that there really are such general observations. I do not want the value of this essay to depend entirely on that idea's being independently more plausible than other theories about observation. I am not certain that it is. But if it has any plausibility at all, and if it really gives us a way to resolve the problem of induction, then it will be worth some future effort to work the idea out in detail.

I. The problem of induction

An inductive inference is often defined as one in which the conclusion does not follow necessarily from the premises – so it is not deductively valid – but in which the premises seem to render the conclusion more likely.[2] This is sometimes seen as a matter of the conclusion's somehow adding to the content of the premises. As Brian Skyrms puts it, "If an argument is inductively strong, its conclusion makes factual claims that go beyond the factual information given in the premises."[3] Wesley Salmon calls anything like this an 'ampliative' inference.[4] (E1) and (E2) below are simple examples of these ampliative inferences.

[2] I will concentrate on one standard type of definition of induction, convenient for my purposes. I believe that what I say can be extended to apply to other common formulations, but will not attempt to do so here. James Cargile provides a discussion of various definitions in "The Problem of Induction," *Philosophy* 73 (1988): 247-275.

[3] Brian Skyrms, *Choice and Chance* (Belmont: Wadsworth, 1986), 8.

[4] Wesley C. Salmon, *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press, 1967).

(E1) This raven is black.
<u>That raven is black.</u>
All ravens are black.

(E2) <u>All ravens observed so far are black.</u>
All ravens are black.

A third common form of inductive argument moves from what is known or observed to particular unknown cases, for example:

(E3) <u>All ravens observed so far are black.</u>
The next raven observed will be black.

This third form may be seen as deductive extension of form (E2), since if we take our observations to imply some general fact, then we can also take them to imply whatever is entailed by that fact. It might also be seen by some as having independent standing as a form of inductive argument. In any case, I will concentrate on forms (E1) and (E2) in what follows. These examples best fit Karl Popper's largely syntactic understanding of induction:

> It is usual to call an inference 'inductive' if it passes from *singular statements* (sometimes also called 'particular' statements), such as accounts of the results of observations or experiments, to *universal statements*, such as hypotheses or theories.[5]

The conclusions of (E1) and (E2) do not follow necessarily from their premises, evidently because the conclusions say more than the premises, in that they talk about all ravens, not just those mentioned in the premises. The problem of induction is, then, often understood to be the problem of justifying non-deductive inferences like these.[6] As Hume was the first to point out, since such

---

[5] Karl Popper, *The Logic of Scientific Discovery* (New York: Basic Books, 1959), 27.

[6] This is controversial. There are many who would like to believe in some kind of ampliative inference, but who also think that the little forms listed are worthless in themselves. We know that the sun will rise tomorrow, not simply because we have a series of past risings of the sun; there must be something else involved, that distinguishes the law-like regularities from the merely accidental ones. A recent strategy attempts to replace enumerative induction with abduction or "inference to the best explanation" (see Hilary Putnam, "The Meaning of Meaning," in his *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), 215-271). I am inclined to agree with Richard Fumerton, in "Induction and Reasoning to the Best Explanation," *Philosophy of Science* 47 (1980): 589-600, that this form of reasoning is

inferences cannot be justified deductively, and cannot be justified inductively either (on pain of circularity), it appears that they cannot be justified at all.

Why should we care about the problem of induction? The answer is that we seem so heavily to depend on such inferences, in science and in ordinary life. That is, we accept as justified many beliefs that can be viewed as the conclusions of inductive inferences, and we further believe that such beliefs originate in inductive inferences. If no such inferences are rationally justified, it looks like we ought to give up much of what we now believe.

Why do we think that what might be called "inductive conclusions," such as that all ravens are black, require inductive arguments? Perhaps because we are empiricists, in at least the broad sense that we believe (or would like to believe) that there are two and only two basic ingredients in human knowledge: observation and proper reasoning (where by proper I mean valid, or else rationally justified in some other way). It may be that we can figure out some things, such as truths of mathematics, *a priori*, through valid reasoning alone. But our knowledge of such things as ravens is not like that; it must be based on observation as well. Unfortunately for general beliefs, it seems that all we can observe at any one time is this or that raven (or, at most, some small number of ravens) and their properties. The general statement that all ravens are black is not deducible from any available set of reports of observations about particular ravens, though those are all that we have to go on. This is why we have a problem, and why it looks as if we need to find some way of justifying ampliative arguments. But I want to reconsider the implicit claim that the general facts in question are always unobservational. I want to suggest that we come to believe them in essentially the same way that we believe particular facts, and with the same kind of justification.

The distinction that I will employ between general and particular statements, facts, or observations is not identical to Popper's, and needs a more definite characterization. There are three types of statements that we usually find listed as the premises in inductive arguments. Some are singular claims of the form "this A is B" or "the C A is B," such as "this raven is black" or "the twelfth observed raven is black." Others are existential claims of the form "Some A's are B," "A least two A's are B," and the like. And still others are universal statements of the form "all C A's are B," such as "all of the ravens in such-and-such a sample are black," or "all

---

effectively reducible to induction. If I am wrong, and abduction must be seen as a distinct form of ampliative inference, it nevertheless stands in the same need of justification as induction. What I say in this paper may be applied as well to the resulting "problem of abduction" as to the traditional problem of induction.

observed ravens are black." It appears that none of the statements usually used as inductive premises have the simple form "all A's are B."[7] This seems a contingent, language-dependent feature of ordinary observation reports. We could always introduce a term like 'obsraves' to denote the class of ravens that have been observed, and then produce the simple universal statement "all obsraves are black." We could also artificially produce a statement like "all ravens are unobserved-or-black." But given the way that we normally speak, it appears that the usual inductive premises *about A's* are effectively particular, in the sense that none of them affirms anything straightforwardly about the entire class of A's, but only about some members, or about a certain subclass.

I will call any contingent statement that is effectively particular in normal language in the way that I have described a *p-statement*. I will call any statement that takes the form of a simple universal affirmative sentence a *u-statement*. In what follows, I will call the facts (if they exist) to which p-statements and u-statements correspond p-facts and u-facts. I will call the objects (if any) to which the subject terms of those statements refer p-objects and u-objects. And I will call observations (if they occur) of p-facts and u-facts p-observations and u-observations. My point is just to focus on the kinds of statements that are involved in alleged inductive inferences, as distinct from the epistemic roles that these statements are supposed to play.

Now I can summarize my understanding of the problem of induction as a set of five jointly inconsistent statements:

(S1) Our knowledge (or justified belief) has the form of a set of observation reports and their consequences closed under proper inference.[8]

(S2) All observation-reports are p-statements.

(S3) All proper inferences are deductive.

(S4) It is impossible to deduce a u-statement from any set of p-statements.

(S5) We have knowledge (or justified belief) of the truth of some u-statements.

---

[7] An exception would be "All of my fingers are unbroken," or something of the sort, where one knows that the entire relevant class is present to the observation.

[8] The class of analytic propositions should be included as well, if these are considered to be substantive objects of knowledge.

Any reasonable approach to the problem of induction must falsify at least one of these five statements. To reject (S5) would be to embrace skepticism with respect to the whole class of universal statements. This is a possible view, of course, but not what we should call a solution to the problem.

Statement (S4) is hard to deny. I cannot prove that it is true, for the obvious reason that the classes of u- and p-statements are only partly defined. But it is demonstrably true for the standard cases that I have in mind – for example, no proposition of the form "all A's are B" can be deduced from any set of propositions of the forms "this A is B" and "all C A's are B."[9]

In most standard presentations of the problem, such as Salmon's, it is simply presupposed that something like statement (S3) must be rejected if the problem is to admit of a solution. There have been many attempts to prove that one or another non-deductive inference pattern is proper. None of these efforts has gained very wide acceptance. Popper and other deductivists affirm (S3) and treat inductive inference as an illusion, arguing that science works essentially through the falsification of some tested hypotheses. But this leaves the positive justification of surviving hypotheses problematic.

(S1) is intended as a concise statement of the central claim of empiricism. While it is surely subject to objections and qualifications, few traditional philosophers of science would deny it wholesale or in spirit. This does not entail that (S1) is true, of course. My point is rather that induction is primarily a problem for broad-sense empiricists in the first place.

There is room in this analysis for another approach to the problem: Deny statement (S2) above. Assert in its place that ordinary u-statements like "All ravens are black" can sometimes be accepted as reports of observations, or as deductive consequences of more general u-statements that are reports of observations. This approach could give us a quick, snappy solution to the problem of induction, if it did not seem so obviously to be false. I want to say that it is actually true, despite appearances – or, at least, that it can be treated as true for purposes of philosophical analysis. In what follows, then, I will do what I can to make the idea of non-particular observations less implausible. To that end, I will try to undermine the common assumptions that support (S2), and to replace them with a quick sketch of an alternative theory of observation. The result will sympathize with Popper's rejection of induction as a fundamental form of reasoning, but offer the idea of general observations as a positive means of justifying "inductive conclusions."

[9] I am ignoring the possibility that C is a vacuous property like 'self-identical.'

II. The common theory of observation

Why does it seem so obvious that all observations are particular? The claim that only p-facts can be observed is not essential to broad-sense empiricism. It stems, rather, from a certain theory about observation. This theory has its roots in common sense, to be sure, and has appeared in philosophical writings since Aristotle's *Posterior Analytics*. But its largely unchallenged status in epistemology may stem more from convenience and simplicity than from any claim to universal truth. It is, in fact, a theory of observation that most present-day philosophers will cheerfully reject when it causes problems in other contexts.

According to the common theory, the philosophically best cases of observation are quite local and brief, such as an individual person's seeing that a certain object in his presence has a certain color. These quick, individual observations find their most natural expression in the form of p-statements. All other cases will be seen as proper observations only to the extent that they approximate these paradigms. This view of observation accords well enough with pre-philosophical intuitions. It is obvious that we can't see everything at once, and we can surely see things better when they are nearby and reasonably small. But for this idea to function as a philosophical theory of observation, not just a rule of thumb, requires further metaphysical, semantic, and epistemological assumptions.

There are three most important such assumptions, and all three have been losing force within philosophy over the past several decades. The first assumption is that, since observational beliefs are epistemically foundational, they should be absolutely certain, or at least as close as possible. The second is that knowledge and justified belief ought to be seen as existing primarily or exclusively in individual minds. The third is that discrete individual objects and their properties are fundamental to the metaphysical and semantic structure of the world. All of these common assumptions were important to the positivists' original project of rationally reconstructing scientific knowledge within something like a classical first-order logical language. Absent the requirements of that project, however, the claim that only particular, immediate facts are observable can be at least reopened for discussion among broad-sense empiricists. Let me reconsider the three background assumptions of the standard theory, then, one at a time.

It used to be held that observations, or at least a certain foundational class of them, must yield absolutely certain knowledge. But few philosophers think this way anymore, and it was never very plausible to apply that criterion to ordinary reports of observations, as distinct from artificial statements about sense-data. For example, if I think I see that a particular raven is black, I can be wrong in a number of ways. It could turn out to be a big crow, not a raven. It could be navy

blue, not black. It could be black on the side facing me, but pink on the other side. I could even be dreaming or hallucinating the whole experience. If we are to speak about ordinary objects rather than immediate sense-data, we can say at best that observing (or seeming to observe) a particular fact gives us good, *prima facie* reason to believe in that fact, but nothing more. As we now say, observational beliefs are *defeasible*. With additional observations and reports from other people (in case there's something wrong with our own eyes, for example), we might get closer to certainty, though we will never get all the way. But if there is no special need for certainty, if all we require of observation is that it give us *prima facie* justification, then there is less reason to restrict the scope of observation to local facts and objects. If I can report, defeasibly, the observation that a certain Roman driver ran his motorcycle into a certain pedestrian, why can I not report defeasibly the observation that Romans in general are reckless drivers? Neither is certain on its face; both would require further investigation to pronounce as definitely true. And many American tourists do claim to observe the general fact that Romans are reckless drivers, calling it an observation in the ordinary sense of the word, just as they claim to observe this or that particular collision or near miss. It is not clear that there is any philosophically essential difference here.

Traditional empiricists have also worried about skepticism with respect to memory. If we believe in foundational observations, we can only get around the problem of memory by requiring that those observations be discrete and very brief events – too brief for memory to play an internal role in the process. Bertrand Russell's remark to the effect that sense-data last "about two seconds" is sometimes seen an amusing example of philosophical bullet-biting. But why does this straightforward statement strike us as funny? I think it is because everybody knows that observations are the sort of thing that can be individuated only arbitrarily. As we speak about them outside of philosophy, observations are often highly indeterminate in duration and scope. Two seconds may actually be an approximate lower bound of sorts: it is about the length of time it takes per sentence to make a series of oral reports at top speed, like a play-by-play announcer at a football game. But this is hardly significant for epistemology. Nor is it relevant that it takes something like a tenth of a second for a person to notice any particular change in his surroundings, since those intervals are not discrete, but plainly overlap each other in a more-or-less continuous way. And unless we wanted to maintain that perception was infallible, while memory was not, there would be no good reason to be concerned about such lower bounds in the first place.

As we usually speak, events and processes of all durations can be observed, and those observations reported. A person can say that he has seen the sun set, seen a new bridge go up, seen an army lose a war, and so on. Why should any of these things be ruled out as proper observation reports? If even the rise and fall of the Third Reich can be considered as one big event, comprising lots of particular and general facts, we should be able to describe William Shirer's lengthy book of that title as the report of one big observation: this very big thing happened, he watched it happen, and the book is his report.

A second background assumption to the standard theory of observation, hence to the problem of induction, is what is called methodological individualism, or sometimes, rather pejoratively, methodological solipsism. This has also been widely rejected in recent decades. It has one source in traditional concerns about the problem of other minds. If, as above, we are determined to base our beliefs on a foundation of certainty, and if the existence of other people's minds is impossible to establish, then we can hardly grant the observations of others equal status with our own. This results in the restricted view that each person's knowledge must be based solely on the observations that he is able to make for himself.

But again, it is not clear that we ought to impose this limitation on the range of observable facts. In ordinary life, we often take reports of others' observations (for example, those of our parents or doctors) as perfectly good grounds for our own beliefs.[10] Moreover, we frequently make reports of shared observations, speaking in the first person plural. (For example, the previous sentence.) Observation reports are given by teams of researchers, by businesses and government agencies through their public relations offices, and by all sorts of other groups.[11] Consider also Hilary Putnam's discussions about metals and trees.[12] Most of us know many things about aluminum, he says, for example that it's cheap and shiny, without being able to distinguish the stuff from molybdenum, or any number of other metals, face-to-face. This implies that our even knowing what we are talking about, in some cases, relies on the existence of distant experts who could make the meanings of our statements more precise. In general, it is

---

[10] I have argued elsewhere that such deference is rationally required of us in a very broad range of cases (Theodore J. Everett, "The Rationality of Science and the Rationality of Faith," *Journal of Philosophy* 98 (2001): 19-42), and that it is through such rational acceptance of the statements of others that we come to know that other minds exist (Theodore J. Everett, "Other Voices, Other Minds," *Australasian Journal of Philosophy* 78 (2000): 213-222).

[11] John Hardwig gives an example of a scientific paper with 99 co-authors, in "Epistemic Dependence," *Journal of Philosophy* 82 (1985): 335-349.

[12] Putnam, "The Meaning of Meaning," 225-227.

increasingly clear that much of human knowledge is distributed socially, rather than duplicated inside each of our heads. If this is right, then there should be no harm in our accepting at least some groups as capable of making at least some observations. The larger are the groups of people who can act together as observers, the bigger and more broadly scattered are the facts and objects we should take as minimally observable.

Suppose I want to say that central planning in agriculture always reduces output. I might describe this as an inductive conclusion of my own, based mainly on written sources, most of which are based on other testimony, books, reports, and scholarly analysis. But there is no reason that this general statement could not be classified as an observation that people have made collectively, rather than an inductive conclusion that I have drawn individually. Statements about well-known facts are often phrased this way in literature, to indicate points that are taken for granted by the writers and their readers. Thus, "…we have seen that it is the Holy Spirit who brings about the wonderful communion of believers in Jesus Christ,"[13] and "…we have seen that no religion stands on the basis of things known… so must it ever be at once a source of error and contention,"[14] and innumerable similar statements.

A third obsolete assumption that supports the traditional theory of observation is logical or metaphysical atomism. The broad idea is that there is one basic level of objects or properties in the world, and that everything else is analyzable in terms of these simplest items. For the early logical positivists, this was a matter of fitting the world to the structure of first-order logic and set theory. Since the collapse of the positivist project in the mid-20th century, almost nobody now thinks that classical logic is adequate to mirror the structure of the world or to analyze scientific discourse. For those who saw the world as fundamentally a set of what I am calling p-objects or p-facts, a particularistic theory of observation was only natural: if there are not really any u-facts or u-objects to begin with, if such things are only logical constructs, then there is nothing special for a u-observation to report. But most of us now hold a less restricted view of the relation between particular and general things. Some find it better, for example, to view the relation of individuals to kinds (e.g. to species in

[13] Pope Benedict XVI, in a speech at the World Youth Day Vigil, held in Australia in 2008.
[14] Francis Wright, "Morals," in *Course of Popular Lectures* (BiblioLife, 2009), 108.

biology) as more like the intrinsic relation between parts and wholes than like the formal relation between members and sets.[15]

There may also be other reasonable choices for the form of an observation report than particular and universal statements as they are classically understood. For example, an improved, non-atomistic semantics might be able to provide an adequate analysis of generic statements.[16] Why shouldn't we say that we have observed the fact that "ravens are black," where the word "ravens" can be understood as picking out the species, in the way that the phrase "this raven" picks out the individual? "All ravens are black" might then be seen as fundamentally similar to "All of this raven is black." Each refers to a certain piece of the world, and says that the entire piece is black.

It could be objected that a *causal* theory of perception favors particularism, in that only a small number of ravens can ever figure causally in any act of observation. But it is not clear that this is true. If the part-whole idea is to be taken seriously, it may be correct to say that whenever particular ravens are involved in an event, ravens in general are also involved, just as an observation or some other event involving one room in my house necessarily involves my whole house too. Moreover, the objection presupposes an atomistic view of the entire causal situation: particular light bouncing off of particular ravens into particular eyes. But there are causal facts at macroscopic levels, too. Unless we are still trying to work within something like positivist limits, nothing prevents us from talking about light *in general* bouncing off of ravens *in general* into the eyes of people *in general*.

We often do use generic statements, rather than u-statements, to report our observations of general facts, and we do so for practical reasons. We are all concerned that our statements reflect, if not full certainty, at least a reasonably high degree of confidence in what we report. Unless we are deliberately engaged in philosophical or scientific theorizing, it is ordinarily safer simply not to report our u-observations as such – that is, not to "generalize" unduly, even if what we are observing is a universal fact. One alternative is just to report those p-observations that we are making at the same time, as we do in scientific lab reports, since these are in practice less likely to be defeated later on. The other is

---

[15] See, for example, David Hull, *Philosophy of Biological Science* (Englewood Cliffs: Prentice-Hall, 1974), 48f, and David Sloan Wilson and Elliott Sober, "Reviving the Superorganism," *Journal of Theoretical Biology* 136 (1989): 337-356.

[16] For a collection of recent efforts, see Gregory N. Carlson and Francis Jeffry Pelletier, *The Generic Book* (Chicago: University of Chicago Press, 1995).

to use the generic form of statement instead, which hedges on the possibility of some indeterminate number of exceptions (not necessarily a minority) to the universal claim. These statements are vague, obviously, but not inherently more vague than ordinary singular statements. In both cases, the subject term picks out some object in the world (say, ravens in general, or some particular raven or group of ravens), and the predicate is used to say something about it. In neither case is it strictly entailed that all, or even most, parts or instances of the subject have the property predicated of the subject as a whole. What is entailed is only that *enough* of the subject has the predicated property. The appropriate sufficiency conditions are not implicit in the statements themselves.[17]

Universal statements like "all ravens are black" are more precise. Such statements correspond to the world in the same way as do those that could be called *universalized* singular statements, such as "all *of* this raven is black." The subjects are again things like an individual raven or ravens in general, but the word "all" has the function of applying the predicate to exactly all, not merely enough, relevant parts or instances of the subject. We can imagine ordinary singular and generic statements as opposite ends of a spectrum, with subject-

[17] E. J. Lowe has made a partly similar, but to my mind needlessly subtle, suggestion, in "What is the 'Problem of Induction'?," *Philosophy* 62 (1987): 325-340. Lowe claims that the class of what I am calling inductive conclusions should not be formulated as u-statements in the first place, but rather as generic statements, which express *laws*, as he understands them, rather than universal generalities. Lowe does not quite say that these law-like facts about biological species and other kinds are themselves observable, but rather claims that observations of their "normal" instances are strong *prima facie* evidence of their truth. This is an attractive view, but hard to evaluate because the concept of a law is so elusive. For one thing, Lowe notes that in order for his laws to count as useful knowledge, we must be able to draw predictions from them in a justified way. But how, for example, can we draw "this is black" from the premise "this is a raven" and the generic formulation "ravens are black?" Not deductively, as Lowe concedes. He relies instead on the principle that most members of a kind must be normal members, so that we can make this sort of inference, in effect, probabilistically. Lowe sees the principle as analytic – it is "incoherent", he says, to suppose it false (Lowe, "What is the 'Problem," 336). But one can easily imagine cases where most of the actual instances of some type are abnormal. For example, some new plague or political development could bring it about that the majority of Canadians have no teeth, without falsifying the claim that a normal (as distinct from average) Canadian does have teeth. More recently, both Howard Sankey ("Induction and Natural Kinds," *Principia* 1 (1997): 239-254) and Brian Ellis ("An Essentialist Perspective on the Problem of Induction," *Principia* 2 (1998): 103-124) have approached the problem of induction along the same broad lines as Lowe, through consideration of the essential properties of natural kinds. Both stop short of claiming that the relevant facts are observable.

predicate statements about mass-type objects (which are often thought of as "scattered particulars") in the middle. At one end of a parallel spectrum would be ordinary u-statements, and at the other end of that spectrum would be universalized singular ones. I do not know what it would take to prove that these connections are as real and as gradual as I suggest. But perhaps these features can be *observed* in the following matrix of statements:

	simple (s/p)	universal
general	Ravens are black.	All ravens are black.
	Apples are red.	All apples are red.
	Peas are green.	All peas are green.
	Pease is green.[18]	All pease is green.
	Corn is yellow.	All corn is yellow.
	Snow is white.	All snow is white.
	The snow is white.	All (of) the snow is white.
	The sky is blue.	All (of) the sky is blue.
	The moon is silvery.	All of the moon is silvery.
particular	This raven is black.	All of this raven is black.
		This raven is all black.

The statements in each column are similar in form. The subjects get less 'classy' and more 'massy,' then less 'massy' and more individual as we move down the page. My claim is that these differences are not very important from an epistemological point of view, unless we are already committed to an atomistic analysis.

Atomism skews the sample for the problem of induction. It forces us to take the most particular singular statements as paradigmatic observation-reports, and to wonder how we get from them to the least particular universal statements. It is more reasonable to take all subject-predicate statements (including generics) to be

[18] 'Pease' is an archaic mass noun for peas, as in "pease porridge hot, pease porridge cold, pease porridge in the pot, nine days old."

equally possible reports of observation, and then to ask how they all relate to the *corresponding* universal statements. Residual problems about confirmation should be the same in principle for the most particular cases as for the most general.

III. An alternative theory of observation

Here is the main idea for an alternative, holistic theory of observation. Think of the world not as a set of pre-cut facts, but as a single, variegated but undivided object. Think of experience not as a series of pre-cut, sentence-like events, but as a more-or-less continuous flow that needs interpretation to be represented propositionally. Think of single experiences as non-random chunks of this whole flow of experience, unified under a broad range of possible criteria. Think of observations as articulate representations of experiences, expressed as statements. On this view, an observation could be large or small, brief or enduring, individual or social. When someone says "I see that your dog is wearing trousers," this expresses a particular observation that fits the standard subject-predicate model, made by an individual more or less momentarily. When someone says "We see that solar activity influences climate," this expresses a general observation, made not individually but socially, and very extensive in time and space. Both are legitimate sorts of observations, because the world has larger and smaller parts, and our experience has larger and smaller parts to match.

There are no *a priori* limits on what sort of empirical theory might best represent our total experience. Therefore, any amalgam of individual or collective experience could theoretically count as an observation, and any statement could count as an observation statement. Ultimately, our decisions as to what to count depend of how our total experience is best systematically articulated into a theory about the whole world. Proximately, though, we do need to rely on rules of thumb regarding what to count provisionally as observation and observability. What I am doing here, then, is debating the restrictive rules of thumb currently in use, and suggesting a more open approach as helpful to philosophical analysis, if not to practical science. I say that we have insufficient reason to insist that one syntactically-defined subset of beliefs is based on observation alone, and the rest only on inference. I think that no belief should be seen as either purely observational or purely inferential. All are functions of a total process that takes in information from the world at various levels of generality, framing hypotheses from these observations, deducing consequences, testing, taking in more observations, and gradually forming an articulate and stable model of the whole, complex system.

Even in the case of an individual observer having a very local experience over a short time, there is no essential particularity in the experience itself. There is no difference in the initial set of sensations between those representing the blackness of ravens in general, for example, and those representing the blackness of this raven in particular. The psychological content of an ordinary observation is not very much like a sentence, after all. From the subjective or internal point of view, we begin with an experience, i.e. *some* experience, and that experience may bring some sentence or sentences to mind. We may or may not articulate that experience with such sentences, but the experience itself is something else. In reality, our observational life is much more like a *flow* of initially inarticulate sensations than it is like a series of sentences being fed in through the senses like input to a computer. Nothing prevents our expressing some of that flow of experience in general terms. We may come into a certain stream of impressions that is both ravenly, as it were, and black. We may then articulate these impressions in an appropriately vague particular form ("this raven is black") or generic form ("ravens are black"), or both. But then to universalize these simple, subject-predicate reports requires something else, a decision that sufficient evidence exists to count the object in question as consistent in all of its parts. We may need to examine more of this raven to conclude that all of it is black, or to examine more of the species raven to conclude that all of them are black. How complete these further tests must be depends on the level of certainty that we require for the resulting universalized beliefs.[19]

Moreover, when we think of observations taking place over longer periods of time (such as a detective's observing that a staked-out gangster always visits a certain nightclub at about one in the morning), all the less does it seem like importing a sentence through the eyes, and all the more like the selection or creation of a sentence to articulate some feature of an otherwise unseparated mass

[19] It is also possible to construct or interpret empirical theories without including definite judgments as to the truth of any particular or universal statement. Instead, we can associate each statement with a probability, and let those probabilities rise and fall according to new evidence, but never reaching either 0 or 1. Bayesians consider a certain formulation of this idea, using Bayes's Theorem in the probability calculus to govern changes in subjective probabilities, definitive of empirical rationality. Wesley Salmon makes the case for this view in "Rationality and Objectivity in Science *or* Tom Kuhn Meets Tom Bayes," in *Philosophy of Science: The Central Issues,* eds. Martin Curd and J. A. Cover (New York: Norton, 1998), 551-593. Clark Glymour argues against it in "Why I am not a Bayesian," in his *Theory and Evidence* (Princeton: Princeton University Press, 1980), and also in *Philosophy of Science: The Central Issues*, eds. Martin Curd and J. A. Cover (New York: Norton, 1998), 594-606.

of impressions. And the more so still, when we consider that some observations might be scattered over many persons, as with a group of veterinarians and ranchers who collectively perceive an outbreak of mad cow disease in their vicinity. To put it sweepingly: there is a whole subjective world, in complex, causal contact with the whole objective world. This contact produces (or possibly constitutes) a mass of evidence. This evidence is then cut up in various ways for various purposes, with appropriate degrees of generality, from one baby seeing one red ball, to a team of scientists observing the long-term effects of a drug on tumors, to humanity as a whole discovering that cooked meat is easier to chew.

I am relying, plainly, on a certain broad faculty of choice, which is involved in our deciding how to aggregate or individuate ourselves as the subjects of our observations, how to aggregate or individuate the objects of our observations, and how to articulate the content of the resulting evidential mass. But we cannot just say whatever we want; there are important constraints that must be placed on any plausible theory of observation. It must be possible, for one thing, to distinguish good observations from bad ones. It must also be possible to distinguish what is observable in principle from what is not. And both theoretical distinctions must accord reasonably well with common intuitions.

First, then, a theory of general observations must leave room for mistakes. It must be possible to distinguish a real general observation (i.e., a correct observation of an actual general fact) from an apparent observation of a general fact that does not exist. For example, if we can observe the fact that all ravens are black (which I have been taking to be true) while directly confronting only some of those ravens, then why do we not properly observe that all swans are white (which is false), when confronting a similar number of white swans? Such mistaken general observations will have to be understood in the same way that we understand mistaken particular observations. I may see a blue car from a distance and perceive that it is blue, in which case I have observed that fact correctly, but I may also see a green car as a blue one, in which case I have made a mistake. In many cases, I may not be able to tell the difference without further research. As I said above, even a single raven in my hand may appear to me to be black, but turn out to be navy blue, or to be pink in those parts I am not directly looking at. We would still say that if it is black, then I am seeing that it is black, not inferring that it is black. This is true, even though my ability to see the raven as a whole relies on the truth of my assumption that the partial surface that I directly see is fairly representative of the entire raven.

Next, the new theory must also preserve something of the intuitive distinction between observable and unobservable objects, facts, etc. This can be

done, I think, along the same lines. What is observable in the new view will be any object or fact, particular or not, in an appropriate relation to the observer, individual or not. Presumably, this will include such general objects as the species raven (which is observed along with its instances, like every other natural kind), and such u-facts as that all ravens are black, as well as such p-facts as that this or that observed raven is black. But it will necessarily exclude those specific facts and objects which are entirely unobserved, such as the species Martian, or the fact that this or that unobserved raven is black, or that all ravens after the year 2500 are black, or that all Martians carry swords.

This may seem to generate a bit of a paradox, in that I am classifying some general facts as observable while some of their deductive consequences are not. If we have observed that *all* ravens are black, how can it be sensibly said that we have not observed that *each* raven (including all of the specifically unobserved ones) is black? But I think that we are already familiar with such relations between facts about wholes and facts about parts. From observed events concerning visible bodies in chemistry, for example, we can infer many properties of their constituent atoms, which cannot be seen as individuals. It might, of course, be protested that the relevant micro-facts are indeed observable, though indirectly, precisely through their effects on larger bodies. But I could happily adopt the same formulation, and claim that inductive predictions are, after all, just another fallible form of indirect perception. We perceive, albeit dimly, that all ravens are black, and infer or indirectly seem to see (why should it matter which we say?) that each 'part' of all ravens, i. e. each individual raven, is black as well.

Consider this brief discussion:

> Amy: How is the pizza at Mario's?
>
> Bob: Pretty good. I've eaten there twice.

There are two ways to analyze Bob's epistemic situation here. One analysis is to say that Bob has tasted certain particular *slices* of pizza at Mario's on a certain two occasions, enjoyed them, and is now reporting an inductive inference to the effect that most of the millions of other slices of pizza at Mario's are equally good. The other analysis is to say that Bob has on two occasions tasted a certain general thing, namely *the pizza at Mario's*, found it pretty good, and is now reporting this directly as an observation. On the first analysis, Bob makes a thorough observation of a few entire small things (give or take some crumbs), about which he is able to judge with a high degree of certainty: those slices were pretty good. As to the pizza at Mario's generally, that should be seen as the set of all such slices, of which

Bob has only tasted a tiny sample. Therefore, he is able to make only a fairly weak induction from his few samples to an enormous class, though such inductions are supported by other inductive beliefs about the usual consistency of restaurant food.[20] On the second analysis, Bob has no greater total certainty about the general quality of Mario's pizza, since his observation of the stuff is slight and could easily be defeated by further experience. But he does, at least, have epistemic contact with the stuff as stuff, not just with members of a set. As a practical matter, it makes no difference which analysis we choose – although I think, as I have said above, that there is no good reason always to favor the first. What makes a difference here is that the first analysis leaves us with the problem of induction, while the second one does not. There is still the problem of grounding beliefs in sufficiently good evidence, and there is still the background problem of perception: how do we know that *any* observation is reliable? But there is no problem of induction where there is no induction.

But, is there really no induction here at all, or am I sneaking it in somehow? You must suspect that I am sneaking it in somehow. Based on my observations of some things, I am claiming to derive beliefs about other things that I have definitely not observed, for example ravens in the year 2050. What else can there be to connect the observed facts with the unobserved facts, other than some form of induction?

Here is my answer. There is indeed an inference from observed facts to unobserved facts, but it is a deductive, not an inductive inference. I observe the universal fact that *all* ravens are black, if it is a fact, when I observe the general fact that *ravens* are black, which I do at the same time that I observe the particular fact that *some* ravens are black. My belief that future ravens will be black is logically entailed by my belief that all ravens are black. It is not observed directly, but it does not have to be. There is no general law, after all, that the deductive consequences of our observational beliefs must be observed themselves, or even observable. Suppose a car goes by, and I observe that it is blue. I already know that all cars have registration forms, and that the color of each car is listed on its form. Therefore, I come to believe, based on my observation of this car, that the word 'blue' appears on its registration form, though I will have no opportunity to see the

[20] At another restaurant, someone makes a little joke:

Carla: How is your filet mignon?

Dexter: I don't know yet. I've only eaten half of it.

This is a joke because we do commonly take our direct knowledge of parts and surfaces of most small things to count implicitly as knowledge of the whole things.

form itself. Now, it may be that this deductive inference yields a false conclusion, of the sort that everyone agrees induction sometimes produces. If I had made a faulty observation of the car that went by, and it was really green instead of blue, then it would not say 'blue' on the car's registration form, so my deduced belief would be a false one. Similarly, if I falsely observe that all swans are white, based on my observations of swans in America, and deduce that swans in Australia are all white as well, then I am simply wrong. But what is wrong is not a faulty inference – my deduction was perfectly valid – just a misleading observation.

This proposed solution can be seen as providing an element that has always been missing from the classical hypothetico-deductive approach to scientific reasoning. On the hypothetico-deductive model, there is no such thing as an inductive argument *per se*. What happens instead, freely translated, is that scientifically interesting u-statements are initially written down only in pencil – that is, as mere hypotheses, not to be believed (because there is no initial reason to believe them), but just to be considered. Once they are on the list, we test them by deducing predictive p-statements from them, and then observing whether or not the predictions come true. In a standard version like Carl Hempel's, a hypothesis is held to be more believable the more it is confirmed by true predictions.[21] In Popper's deductivist alternative, the hypothesis is never confirmed, but merely 'corroborated' by surviving attempts to find predictions that turn out to be false.[22] Now, these procedures (one or both) strike most of us as a better description of actual scientific reasoning than simple inductive arguments. It does seem right to say that u-statements acquire greater credibility as they pass successfully through more comprehensive and more rigorous tests. But, as Salmon and others have pointed out, neither variant of the hypothetico-deductive approach provides a real solution to the problem of induction, because each fails to show how testing actually justifies belief in a hypothesis.[23] No account is given as to why one hypothesis should be initially considered rather than another, and it is not made

[21] Carl Hempel, *Philosophy of Natural Science* (Englewood Cliffs: Prentice-Hall, 1966).

[22] Popper insists in *The Logic of Scientific Discovery* that he is not attempting to justify either induction or the hypothetico-deductive model, as he understands these terms. Instead, he wants his approach to be seen as entirely deductive.

[23] As Salmon points out in *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press, 1967), 25-26, if corroboration is supposed to give us any *reason to believe* the general hypothesis in question, based ultimately only on particular results of observations, then this amounts to an ampliative (hence non-deductive) element in Popper's theory, whatever he chooses to call it.

clear why confirmation or corroboration makes the hypothesis in question more likely to be true than its surviving competitors.

On the view that I am suggesting, however, our initial choice of one hypothesis over another can be accounted for, since some general statements will appropriately articulate our general observations, and some will not. An account can also be given of why both confirmation and non-falsification tend to add epistemic weight to these hypotheses. If we take the u-statement in question initially as the tentative report of an imperfect observation, then what are usually considered to be separate observations of confirming or non-falsifying instances can be seen instead as extensions and clarifications of the *same* observation. It would be a matter of making sure that our initial observation is a good one – in the same way that someone who thought he had seen an individual black raven might catch the bird and study it carefully, in order to add ink to his initial penciled-in report.[24]

As long as there are some observationally acquired u-statements available from which appropriate theoretical hypotheses could be deduced, there is no need to hold that all types of general fact can be observed directly. It is in principle only necessary that there be *one* sufficiently general u-statement, the truth of which can be affirmed provisionally through observation – perhaps even something like "inductive inferences are generally reliable." Kant tried to show that some such principle of nature's uniformity is knowable *a priori*, though Hume's arguments against that possibility seems to have proven more persuasive over time. In any case, once we had such a universal hypothesis penciled-in through observation, more specific u-statements could be deduced from it, and jotted down as likely to be true. The two-stage argument would go something like this:

(U1) Induction is reliable, i.e. if all observed A's are B, then probably, all A's are B (observed).

(U2) Therefore, if all observed ravens are black, then probably, all ravens are black (deduced from (U1)).

(U3) All observed ravens are black (observed).

(U4) Probably, all ravens are black (deduced from (U2) and (U3)).

[24] This is why we take some scientific experiments to yield general knowledge on the first try, and view repetitions as providing reassurance to our initial results, rather than new, logically separate facts. For example, it required only one carefully observed solar eclipse (in 1919) for physicists to perceive that light bends around massive objects.

In this way, the idea of inductive inference is ultimately vindicated by means of observation. But it is not vindicated as a fundamental form of reasoning – only as a certain conditional formula that has been observed to work well in general. The high-level principle of uniformity would not have to be observed in an immediate way, either. We could start with a few lower-level observations, to the effect that all ravens are black, all rats have tails, and the like. We could then submit some of these basic statements to the usual sorts of testing. If successful, the whole resulting situation could be said to be contained in an observation of the fact that this observational-deductive method usually works. Thereafter, we could with greater and greater confidence deduce unobserved hypotheses from the initially-weakly-observed general principle, and then through usually-successful testing add credence to both. This kind of 'bootstrap' procedure would require only that there be enough initial observational input at some level for the whole process to get going.[25]

IV. Conclusion

In this paper, I have argued that the problem of induction, as it is usually conceived, presupposes the impossibility of our observing general facts. This is

[25] The idea of general observations might also help a bit with Nelson Goodman's variant problem of induction. Goodman asks, in *Fact, Fiction, and Forecast*, 4th edition (Cambridge: Harvard University Press, 1983), how we can rationally choose to generalize on the basis of our ordinary concepts, like the colors blue and green, rather than such odd but clearly describable properties as 'grue,' which he defines as either green if first examined before a certain time t, or blue if first examined afterwards. Any prediction which 'projects' the property green before time t will be justified by precisely the same evidence, he says, as the corresponding prediction which projects the property grue – but clearly these are different, incompatible predictions, and arbitrarily many such equally incompatible, but equally well-evidenced, predictions could be generated just as easily. What good reason is there to prefer one to the other? One possible response is to replace the question of good reason with a question as to what one actually sees. If it is a fact that all emeralds are green, then this is one of the facts that one observes, when one observes a number of green emeralds in the absence of any observations to the contrary. But since it is not an actual fact that emeralds are grue, it is not, *a fortiori*, an observable fact. I could believe, of course, that I was seeing something grue, not green, when I was looking at a pile of emeralds. And this would certainly be a mistake, like looking at a solid green stone and somehow believing it was blue on the reverse side. But there may be some advantage to analyzing this mistake as a faulty observation, as distinct from an irrational inference. Grue, as defined, would not then be seen as an unprojectable property in principle, but rather as an invisible (or indistinguishable) one in practice. This is only to suggest an angle on Goodman's problem, of course, not to pretend to have solved it.

why we seem to need inductive inferences to justify our general beliefs. But such inferences are hard to specify and seemingly impossible to justify in their own right; hence, the problem. I have attempted to undermine the common view of observation as always particular in scope, by arguing that the foundationalism, atomism, and individualism on which it seems to depend are all rightfully obsolete. I have suggested an alternative, holistic account of observation as a replacement, according to which general statements are indeed observable, albeit typically with low initial certainty. And I have tried to show how these defeasible general observations would neatly fit into the standard hypothetico-deductive model of scientific reasoning, by providing hypotheses, previously viewed as unempirical, with some measure of *prima facie* justification.

Somehow, I doubt that every reader has been totally convinced by these remarks to abandon the traditional idea of observation as exclusively particular, and to accept my sketch of a holistic account as adequate to the resolution of the problem of induction. But perhaps some readers are convinced to this extent: that the problem of induction depends on a certain theory of observation, that this theory is questionable, that a different theory can provide at least a superficial answer to the problem, and that there is an approach here worth exploring further.