Demonstrative Induction and the Skeleton of Inference

P.D. Magnus

October 9, 2008

This is the final author's draft of a paper forthcoming in *International Studies in the Philosophy of Science*. Earlier versions were posted on the author's website `fecundity.com` under the title 'Eliminating Induction', first in December 2005.

P.D. Magnus is at the Department of Philosophy, University at Albany.

Abstract

It has been common wisdom for centuries that scientific inference cannot be deductive; if it is inference at all, it must be a distinctive kind of inductive inference. According to demonstrative theories of induction, however, important scientific inferences are not inductive in the sense of requiring ampliative inference rules at all. Rather, they are deductive inferences with sufficiently strong premises. General considerations about inferences suffice to show that there is no difference in justification between an inference construed demonstratively or ampliatively. The inductive risk may be shouldered by premises or rules, but it cannot be shirked. Demonstrative theories of induction might, nevertheless, better describe scientific practice. And there may be good methodological reasons for constructing our inferences one way rather than the other. By exploring the limits of these possible advantages, I argue that scientific inference is neither of essence deductive nor of essence inductive.

1. Introduction

According to some accounts, scientific inference often needs nothing more than deductive rules. Statements of the phenomena, along with constraints on permissible theories, are sufficient to deductively entail a single correct theory. Inferences of this kind are called *demonstrative inductions*, and I use the phrase *demonstrative theory of induction* to indicate an account according to which demonstrative inductions are centrally important in science. Demonstrative theories of induction are offered by *inter alia* Bird (2006), Dorling (1973), Kitcher (1993, ch. 7), Massimi (2004), and Norton (1993, 1994, 1995).

In what follows, I begin by describing demonstrative theories (§1). I then offer general considerations about inference to show that there is no logical reason to prefer demonstrative theories over other ways of reconstructing scientific inference (§§2–3). Although this places serious limits on demonstrative theories, they might still succeed as descriptive or heuristic accounts (§4). In the penultimate section, I extend these lessons from demonstrative theories to other recent accounts of scientific inference (§5).

2. Demonstrative Theories of Induction

It is typical to say that scientific theory choice involves forming conclusions that would not be licenced by good old deductive reasoning. Indeed, this is often taken to be the starting place for any philosophical account of confirmation. Clark Glymour writes, "The aim of confirmation theory is to provide a true account of the principles that guide scientific argument insofar as that argument is not, and does not purport to be, of a deductive kind" (1980, 63). Some philosophers have disagreed that there are any such principles to be discovered. Following Hume, perhaps it is strictly a psychological process. Following Popper, perhaps there is no such thing as confirmation. Regardless, all of these approaches agree that confirmation of scientific theories would require more than deductive principles. Induction and deduction must be different things.

This point is closely connected to the so-called underdetermination of theory by data. If phenomena are insufficient to deductively entail a single scientific theory, then there are indefinitely many incompatible theories that are consistent with the phenomena. It is common to say, as Laudan (1990) and Lipton (2004, ch. 1) do, that answering underdetermination requires acknowledging that science involves non-deductive, ampliative inference rules.

The common wisdom is correct at least in this respect: We cannot start with singular observation reports and derive scientific theories as deductive consequences. Supposing that observation reports are the only permissible premises, it follows that scientific inference cannot be deductive. Yet there are accounts of scientific inference that permit richer premises and so make scientific inference into a deductive affair. The terminology is somewhat tangled: Deductive inferences are variously called demonstrative inductions, deductions from phenomena, eliminative inductions, and Holmesian inference.

These are not 'induction' in the sense of ampliative inference. They are deductive in the full-blooded sense that the conclusion is a valid consequence of premises that scientists (fallibly) know to be true. For example: Taking the invariance of the speed of light as an empirical fact and the theoretical constraint that the laws of physics must the same in all reference inertial reference frames, it is possible to derive the Lorentz transformations.[1]

In the Newtonian tradition, demonstrative inductions are called *deductions from the phenomena*. As Worrall (2000) notes, the premises of these deductions include not only the phenomena (appropriately described) but also theoretical constraints.

So-called *eliminative induction* is a variety of demonstrative induction. This is the sort of inference that Sherlock Holmes recommends to Watson: "[W]hen all other contingencies fail, whatever remains, however improbable, must be the truth" (Doyle 1905, 926). Bird (2006, 11) advocates eliminative induction under the moniker *Holmesian inference*. Mystery stories are often structured to facilitate this sort of inference. There is a list of suspects, one of whom must have committed the crime. If all but one suspect can be exonerated, then the remaining suspect must be guilty. This inference is, properly speaking, deductive. Holmes knows that some set T contains all the possible accounts of the crime. He determines that each $t \in T$ is false except for one. By disjunctive syllogism, he concludes that the remaining element must be true. Holmesian inference cannot be used to justify the premises that the members of T exhaust the possible explanations or that particular members of T are false, on pain of regress. The premises will typically be known by other means. Where the premises are known, scientific inference may be deductive.

In what follows, I will use the phrase 'demonstrative theory of induction' to indicate accounts according to which demonstrative inductions are centrally important in science. Borrowing a term from Norton (2003), I'll call the extra premises supplied by demonstrative inductions 'material postulates.'

3. Everybody Needs Rules

In this section, I begin with two observations about inference systems and then apply them to the issue at hand.

> *Observation 1:* If two proof systems licence the same inferences, then they are equivalent.

The business of a proof system is to licence inferences. Given true premises, a deductive system only leads to true conclusions. If two proof systems lead from the same premises to the same conclusions, then choosing between them is a practical rather than a logical matter.

This is most obvious in cases where it is possible to construct a translation manual that converts valid proofs in one system into valid proofs in the other. They agree on which inferences are legitimate and which are not, so they are merely different formulations of the same logic. Of course, there may be practical reasons to prefer one logic over the other. Perhaps proofs in one system are longer or less intuitive than proofs the another. When economy and simplicity vary between two equivalent formal systems, we work in whichever most readily does what we want.

The logical force of a proof system is constituted by the inferences it licences or censures. For any deductive system, this amounts to a set containing all of the premise-conclusion sequences that the system recognizes as valid. Proof systems that identify the same arguments as valid may differ on the number of steps required to get from some specified premises to a specified conclusion, on what intermediate conclusions must be derived, on what axioms (if any) must be invoked— but those are differences in implementation, not in logical force. A system of inductive inference does not identify arguments as 'valid', but is still similarly associated with a set of premise-conclusion sequences it licences. If two systems licence the same set of arguments, they have the same logical force.[2]

> *Observation 2:* Any proof system must involve rules of inference, but there is no specific requirement as to what those rules must be.

A deductive proof system cannot yield any conclusions if it is constructed entirely of axioms. This is one lesson of Lewis Carroll's *What the Tortoise Said to Achilles* (1895). To put it in modern symbolism, the Tortoise grants to Achilles that P and $(P \supset Q)$ are true. Yet how are we to move from these premises to the conclusion, Q? The Tortoise insists that we should only accept the inference if we accept another axiom: $[P\&(P \supset Q)] \supset Q$. And so on.

The Tortoise insists that the proof is only complete when the rule of inference itself is written down as a line of the proof; indeed, he insists that each specific instance of the schema be written down. The regress results because— in each abortive attempt to apply the rule— a further instance of the general rule is required to move from the latest conditional to the conclusion. There is a distinction between *axioms* (which one is always warranted in writing down) and *rules of inference* (which one may apply without writing down).[3] We must have some rules of inference. However, we have a choice as to what weight will be carried by axioms and what weight will be carried by rules of inference.

In a natural deduction system, there are no axioms but a wealth of rules: an introduction and elimination rule for each logical operator. This approach has definite virtues; it illustrates the meaning of the operators and highlights the unified structure of the deductive system. In other systems, there are lists of axioms and only a few rules of inference. Sometimes only a rule like modus ponens is provided. Of course, there is nothing unique about modus ponens— it is just a traditional and perspicuous choice. Such a rule-poor but axiom-rich proof system is provably equivalent to a rule-rich natural deduction proof system; i.e., it is possible to construct a translation manual to turn any proof in one system into a proof in the other. The rules in the natural deduction system can do all the work done by axioms in the axiomatic system. So, by Observation 1, they are different expressions of the same logic. The choice between them could only be motivated by practical, extra-logical factors such as pedagogical clarity or formal elegance.

An inductive inference system will licence more inferences than a deductive system. Just like a deductive system, however, an inductive system must have some rules of inference. A demonstrative theory of induction describes a system of inference that includes all the rules and axioms of some deductive proof system plus a rich collection of topic-specific, contingent, local premises. Following Norton (2003), call these *material postulates*. They play the rôle of further axioms, sentences which may be introduced in derivations so as to yield conclusions.

Just as the work done by axioms may be done by rules, the work done by material postulates may be done by ampliative rules. Such ampliative rules are often called *material inference principles*. Formally, there is no reason why deductive work must be done by new axioms rather than by new rules. So too, there is no reason why inductive work must be done by material postulates rather than by material inferences principles. The two approaches are inferentially equivalent, so a demonstrative theory of induction is a representational choice rather than something forced on us by the nature of scientific inference.

4. Your Enthymeme is My Risky Inference

In this section, I begin with two more observations about inference systems. They provide a different path to the conclusion of the previous section.

> *Observation 3:* Any putatively ampliative argument can be reconstrued as a deductive argument.

Any deductively invalid argument may be construed as an enthymeme, a valid argument with suppressed premises. Trivially, one can add as a premise the conditional that has the conjunction of the argument's premises as its antecedent and the argument's conclusion as its consequent. If we think that the premises of an ampliative argument justify belief in its conclusion, then we should accept at least a cautiously stated premise to that effect. It is typically possible to construct more plausible, less trivial formulations of such a premise, although there is no general algorithm for doing so.

This observation was familiar even before the heyday of Newtonian deductions from phenomena. In the *Port-Royal Logic*, Arnauld and Nicole claim that enthymemes are more common than explicit deductive arguments. The suppression of premises in discourse, they say, "flatters the vanity of one's listeners by leaving something to their intelligence and, by abbreviating speech, it makes it stronger and livelier" ([1683] 1996, 176).

> *Observation 4:* Any deductive argument with less-than-certain premises can be reconstrued as an ampliative argument in which the uncertainty is explicitly acknowledged.

Scientific facts may be described in a more or less cautious way, with what Pinch (1985) calls greater or lesser degrees of *externality*. 'The needle points to 11' is a low externality report. 'The speed of light is invariant'— treated as "an empirical fact of experience" (Dorling 1995, 98)— is a high externality report. Higher externality reports, when questioned, can be treated as inferences from lower externality reports. Those inferences will typically be ampliative.[4]

Demonstrative inductions take high externality reports and theoretical constraints as premises. Formulated in this way, they are deductive. By making the argument for the high externality premises explicit, the arguments can be reconstrued as ampliative inferences from lower externality premises.[5]

Holding the low externality reports fixed, scientific inference may be treated as involving ampliative inference up to a point and then deductive inference. The same inference may be treated as involving ampliative inferences all the way, provided we accept ampliative rules of the right form. It may even be treated as a hodgepodge of ampliative and deductive rules applied higgledy-piggledy, given different ampliative rules. Each of these reconstructions requires *some* ampliative rule or rules. It may be that some set of ampliative rules will be more intuitive or more elegant, but the competitors can be cooked up so as to licence the same inferences.

By Observation 1, it follows that there is no logical difference between the different treatments. Although any scientific inference that can be construed as deductive may also be construed as non-deductive, this difference does not immediately ground a disagreement about the logic of science. We only disagree about the logic of science if we disagree about which conclusions are justified when.

5. The Aim of Demonstrative Theories

What do these arguments mean for demonstrative theories of induction? Their fate depends on what they are meant to show.

5.1. Theories of Induction as Descriptions

Demonstrative theories might be *descriptive* accounts. When I am constructing a deductive proof, I know whether I am working in an axiomatic system or a natural deduction system. In much the same way, even though a proof system that exploits demonstrative inductions and one that does not can be inferentially equivalent, scientists use some system or other when drawing an inference. Just as an observer can look at the chalkboard and determine the sort of proof system I am using, we can look at scientific practice to see what system scientists have employed.[6] The historical record may be mixed, ambiguous, or incomplete, but it is legitimate to ask: Did scientists actually use demonstrative induction in this or that case? Advocates of demonstrative theories have shown that, at least in some cases, the answer is 'yes.' (See esp. Bonk (1997), Dorling (1973), Kitcher (1993, ch. 7), and Massimi (2004).)

5.2. Theories of Induction as Advice

Demonstrative theories might be *practical* recommendations; cf. Dorling (1995) and Norton (1995). Just as different deductive systems may be more intuitive or more elegant, demonstrative inductions may have practical advantages over other types of inference. In mundane cases, like those constructed in mystery stories, eliminative induction is a useful inference strategy.

Even Worrall (2000), who has mixed feelings about demonstrative inductions, admits that they can be heuristically powerful. He has suggested a principled way in which to regiment the use of ampliative and deductive rules (Worrall 2002). On the ampliative side, he advocates a heuristic 'use novelty' account. If a phenomenon was not taken into consideration when constructing a theory and the theory predicts the phenomenon, then observation of the phenomenon supports the theory. Kepler's laws of planetary motion were not included as a presupposition of Newton's theory of gravitation, but they can be derived from it along with some background assumptions. Even though Kepler's laws were already known, they provided some support for Newton's theory. (Regarding the heuristic account, see e.g. Worrall 1989.) On the deductive side, there are demonstrative inductions. The precise value of the gravitational constant, for example, is not a shocking prediction of Newton's theory of gravitation but something to be derived from careful observation.

The difference, Worrall suggests, is that the deductive work can only begin once the ampliative work is done:

> The judgment of support in the "deduction from the phenomena" cases is ineliminably *conditional*, an ineliminably *intra*-paradigm

> or *intra*-research programme judgement. *Given* that a general framework, or research programme, is already accepted, then the data give... not just *some* support for the scientific theory, but conclusive support. (2002, 203, emphasis in original)

This suggestion makes sense of some cases. For example, Bonk (1997) describes a nineteenth-century deduction that gravitation must be an inverse square law and a 1920s deduction of the quantization of energy. In these cases, the theoretical constraints effectively specified a theory with free parameters; the demonstrative inductions provided a precise way of fixing the parameters given the data, without further assumptions. Demonstrative induction served as "a convenient computational technique for fine tuning a given theoretical model, with little confirmational impact" (Bonk 1997, 68). The demonstrations did not justify believing a theory, but rather facilitated inferences *within* a theory. It is important to note that Worrall's suggestion is not merely meant to show *that* demonstrative inductions are sometimes but not always useful in science. Rather it is meant to show exactly *when* they are useful and when they are not. Two reasons suffice to show that Worrall's criterion of demarcation is unsatisfactory.

First, the heuristic account of confirmation tells us when a theory is incrementally confirmed but has no detachment rule. It does not tell us when sufficient confirmation has accumulated to actually accept a theory. Nor does it provide any way of measuring the degree of confirmation or the extent to which we can rely on the theory. As such, there is no rule that tells us when to adopt a paradigm or research programme and, as such, when to shift from ampliative to deductive inference. A given conclusion might be the product of demonstrative induction for scientists who have already made the leap to a new paradigm but the product of ampliative inference for scientists who have not. The demonstrative induction is "ineliminably conditional" and Worrall's criteria do not tell us when to accept the antecedent of that conditional. As a descriptive matter, one might claim that scientists do begin to accept demonstrative inductions when they come to believe (rather than merely entertain) the premises. Regardless, if demonstrative theories are to serve as practical recommendations rather than as mere descriptions, then it is insufficient it is insufficient to note that scientists might and sometimes do offer demonstrative inductions— one must say when and how it is appropriate to do so.

Second, separating confirmation of a paradigm from intra-paradigm inference requires individuating paradigms. This is a notoriously difficult problem. As Kitcher argues, "the game of finding paradigms... in the actual course of events becomes highly arbitrary and often unprofitable" (1993, 89). The problem is little easier if we replace 'paradigm' with 'research programme' or 'theory.' Any inference within a broad theory may be redescribed as the adoption of a narrower theory against a background of other theories. For this reason, underdetermination arguments are often pitched as arguments about total sciences rather than arguments about specific theories; cf. Magnus (2005).[7]

Without some criterion for when to employ demonstrative induction and when to employ ampliative rules, demonstrative theories as practical recommendations are importantly incomplete. Many case studies show that scientists have profitably used demonstrative inductions in the past, and so one might infer that scientists will profitably do so in the future. The resulting demonstrative theory is rather weak, however. We infer by an ampliative rule that demonstrative inductions are useful, but we can only say where they have been useful in retrospect.

5.3. Theories of Induction as Justifications

Demonstrative theories might claim that demonstrative inductions provide greater *justification* for scientific theories. For example, authors such as Norton (1993) and Massimi (2004) have claimed that demonstrative inductions resolve the problem of underdetermination; scientists are able to draw deductive conclusions from phenomena (described at a high level of externality) and background constraints on theory. Massimi suggests that "a demystification of the problem of underdetermination" results from acknowledging that "the theory that comes to be chosen is derived from the salient phenomena with the help of the salient theoretical constraints" (2004, 266). As anyone must admit, the justificatory burden is passed off to the high-externality descriptions and theoretical commitments. If these in turn are justified by another demonstrative induction, then more premises are required and a regress ensues.[8]

When scientists do accept the reports of phenomena and theoretical constraints that serve as premises of a demonstrative induction, reason demands that they accept the conclusion. As a psychological matter, this may explain why scientists do not worry that their theory choice is underdetermined. Demonstrative induction "makes the phenomenon of underdetermination, so beloved by philosophers, invisible to working scientists" (Worrall 2002, 200). The demonstrative induction also provides a kind of justification for the theory choice: *Given* the beliefs that scientists have, they ought to accept the conclusion.

Norton argues that this makes the regress acceptable. He writes, "It merely describes the routine inductive explorations in science. Facts are inductively grounded in other facts; and those in yet other facts; and so on. . . . What remains an open question is exactly how the resulting chains. . . will terminate and whether the terminations are troublesome" (2003, 668). The regress stops with the premises that we actually accept, and that is OK.

Appealing to background theories in this way suffices to answer underdetermination worries; cf. Kitcher (1993) and Magnus (2005). To count as a defence of demonstrative theories, however, it needs to be shown not only that accepted background theories may resolve underdetermination— but also that it is impossible for ampliative rules to fill that rôle. It seems *prima facie* as if appealing to beliefs we find unproblematic is neither more nor less secure than appealing to inferences we find unproblematic. Norton (2003, 667–668) insists that this is not so, but does not give much of an argument.

In favor of demonstrative theories, one might argue in this way: Scientists are better at evaluating whether to believe explicit *claims* than they are at evaluating whether to commit themselves to *rules* of inference. If that were true, it would make demonstrative inductions more reliable than the parallel inferences carried out with ampliative rules. This is a practical matter, but it is also a matter of whether certain practices are apt to lead to truth or not— and so it is a matter of justification. I have expressed this hypothetically because I do not know whether there is this difference in reliability between scientists' assessments of premises and their assessments of rules. Norton and others have an intuition that there is this difference, but something more than an intuition is required— because, of course, their opposition does not share their intuitions. Whether there is a psychological difference that makes us better at assessing general premises than ampliative rules is a contingent matter. Champions of demonstrative induction need more than the mere suggestion or intimation of a difference. Moreover, if such a difference in reliability is to ground a general defence of demonstrative inductions, then the difference must be robust. It must obtain not only for some particular group of scientists, but for most or all scientists. It must endure across changes in methodology. I do not know of any evidence that there is such a difference.

Alternately, detractors of demonstrative induction might argue in this way: A demonstrative induction typically requires both high-externality observation reports and premises that constrain the form of admissible theories. These latter constraints are non-empirical, and so (the argument goes) expose us to *more* risk than an ampliative argument from the same observations to the same conclusion. Of course, the material postulates that figure in demonstrative inductions are contingent and so involve some risk. However, ampliative rules will only be reliable in some subset of possible worlds—simply put, that is what makes them ampliative and not deductive.[9] The rules, too, involve risk. Where we are led to the same conclusions, we accept the same risk regardless of whether we are led there by material postulates or inductive rules. If the detractor means to say that we have some greater insight into which rules to trust than which premises, then again we are in the realm of psychological speculation. An appeal to intuition will get them no further than it gets champions of demonstrative induction.

Champions and detractors alike face a serious challenge.

6. Further Consequences

The points above regarding the relation between deductive rules and ampliative rules hold *mutatis mutandis* between different sets of ampliative rules. I want to briefly consider consequences of this for two other accounts of induction.

First, consider Norton's (2003) distinction between *material* and *formal* theories of induction. A formal theory characterizes the correctness of an inductive inference in terms of general, ampliative rules. A material theory of induction treats induction as an inference relying on *material postulates*, explicit premises about the kind of system in question.

He argues that there are no general conditions for the correctness of inductive inference. For example: It is legitimate to infer the melting point of pure gold from melting one sample of pure gold, but it is not legitimate to infer the melting point of beach sand from melting one sample of beach sand. Although the inferences are formally identical, one is about an element and the other about a variable mixture. Thus, Norton concludes, induction requires paying attention to facts and not merely to form: The correct account of induction must be a material theory.

A material induction might not be strictly a demonstrative induction. For example, inferring the melting point of gold might rely on a specific material postulate like "Chemical elements are *generally* uniform in their physical properties" (Norton 2003, 651, emphasis in original). Elements being *generally* uniform is not quite the same as their being *all and always* uniform. The weasel-quantifier 'generally' means that an inference using this postulate will not be deductive. (I would like to thank Norton for drawing my attention to this point.) Nevertheless, the points made above about demonstrative inductions may be profitably be applied here.

Norton's contrast between formal theories (general, ampliative) and material theories (local, deductive) conflates two separate questions. (1) Can all inductions be corralled under one tent or must they be savvy to features of the world, such as whether something is an element or a mixture? (2) Should induction involve distinctive ampliative rules or not? Norton implicitly suggests that taking the first horn of one of these questions means taking the first horn of both, and similarly that taking the second horn of one of them means taking the second horn of both.

Like a demonstrative induction, a material induction works because of an inductive premise that supports a generalization (a material postulate) rather than by an inductive rule (a material inference principle). As we have seen, however, whether or not an inference involves ampliative rules is a matter of how we reconstruct it. Norton's argument against general theories of induction addresses question 1; his conclusions about question 2 are a non sequitur.

Perhaps Norton would readily admit as much, since nothing said here poses problems for his answer to question 1. He may be right that material postulates or principles need to be characterized in a domain-specific way. The point here is only that this context sensitivity might be built into content-sensitive rules or concrete premises with equivalent inductive risk.

Second, consider Lipton's (2004) strategy of exploiting underdetermination: Begin with some inferences we accept, the rules of deduction, and some acknowledged inductive rules. If the rules are insufficient to justify the inference, then we must be employing some further inductive rules. This allows us to discover where unrecognized ampliative rules are at work. Lipton (2004, 5–7) attributes this strategy to Chomsky and Kuhn, and he himself uses it himself in attempting to suss out inference to the best explanation.

This is unproblematic if we have all the premises of the inference explicitly formulated. Yet, as we have seen, the inferential gap may be filled by suppressed premises as easily as by further rules. Since Lipton allows for previously acknowledged ampliative rules, the new premises need not be so strong as to make the argument deductively valid. One might strengthen the explicit premises so that the acknowledged rules are sufficient to licence the inference, removing the impetus for adding further rules.[10]

Where there is sufficient evidence as to how scientists actually carried out an inference, we can discover what their premises were and what rules they employed; see §5.1., above. The point against Lipton's method is the converse of this. Where there is insufficient evidence as to how scientists actually carried out an inference, we cannot discover whether their inference relied on hidden premises or inexplicit rules.

7. Conclusion

For any set of conclusions, there is no difference (in terms of risk) between committing myself to a risky rule that warrants all and only those conclusions and committing myself to risky premises that deductively entail all and only those conclusions. The equivalence between demonstrative and non-demonstrative theories of induction means that they expose us to the same degree of risk. There may be legitimate practical reasons when constructing or reconstructing scientific inferences to construe them as deductive or as ampliative. Nevertheless, the scientific inferences could in principle be construed the other way and so are neither essentially deductive nor essentially ampliative.

Massimi claims that "the advantage of [demonstrative induction] over [ampliative] inductive methodologies consists in confining the 'inductive risk' to the phenomenal premises… without spreading it throughout the skeleton of the… inference" (2004, 263). It is sound advice to lift with your legs and not with your back. It might similarly be sound advice to lift with the muscle of premises rather than the skeleton of inference, but the weight is the same in either case.

Acknowledgements

An earlier version of this paper was presented to the Society for Exact Philosophy meeting in La Jolla, California and to the philosophy department at Southern Methodist University. I want to thank members of the audience there for helpful comments. Thanks also to John Norton (for commenting on a prior draft) and Bradley Armour-Garb (who engaged induction as a lunchtime conversation topic on many occasions.)

References

Arnauld, A. and Nicole, P. ([1683] 1996). *Logic or the Art of Thinking*. Cambridge: Cambridge University Press. The Port-Royal Logic. Translated and edited by Jill Vance Buroker.

Bain, J. (1999). Weinberg on QFT: Demonstrative induction and underdetermination. *Synthese*, 117:1–30.

Bird, A. (2006). Abductive knowledge and Holmesian inference. In Gendler, T. S. and Hawthorne, J., editors, *Oxford Studies in Epistemology*, volume 1, 1–31. Oxford: Oxford University Press.

Bonk, T. (1997). Newtonian gravity, quantum discontinuity and the determination of theory by evidence. *Synthese*, 112:53–73.

Carroll, L. (1895). What the Tortoise said to Achilles. *Mind*, 4(14):278–280.

Dorling, J. (1973). Demonstrative induction: Its significant role in the history of physics. *Philosophy of Science*, 40(3):360–372.

Dorling, J. (1995). Einstein's method of discovery was Newtonian deduction from the phenomena. In Leplin 1995, 91–111.

Doyle, A. C. (1905). The adventure of the Bruce-Partington plans. In *The Complete Sherlock Holmes*, volume 2, 913–931. Garden City, NY: Doubleday& Company.

Glymour, C. (1980). *Theory and Evidence*. Princeton, NJ: Princeton University Press.

Kelley, K. (1996). *The Logic of Reliable Enquiry*. Oxford: Oxford University Press.

Kitcher, P. (1993). *The Advancement of Science*. Oxford: Oxford University Press.

Laudan, L. (1990). Demystifying underdetermination. In Savage, C. W., editor, *Scientific Theories*, 267–97. Minneapolis, MN: University of Minnesota Press.

Leplin, J., editor (1995). *The Creation of Ideas in Physics*. Dordrecht: Kluwer.

Lipton, P. (2004). *Inference to the best explantion*. London: Routledge. second edition.

Magnus, P. (2005). Background theories and total science. *Philosophy of Science*, 75(2):1064–1075.

Massimi, M. (2004). What demonstrative induction can do against the threat of underdetermination: Bohr, Heisenberg, and Pauli on spectroscopic anomalies (1921–24). *Synthese*, 140:243–277.

Norton, J. D. (1993). The determination of theory by evidence: The case for quantum discontinuity, 1900-1915. *Synthese*, 97(1):1–31.

Norton, J. D. (1994). Science and certainty. *Synthese*, 99(1):3–22.

Norton, J. D. (1995). Eliminative induction as a method of discovery: How Einstein discovered general relativity. In Leplin 1995, 29–69.

Norton, J. D. (2003). A material theory of induction. *Philosophy of Science*, 70(4):647–670.

Pinch, T. (1985). Towards an analysis of scientific observation: The externality and evidential significance of observational reports in physics. *Social Studies of Science*, 15:3–36.

Schell, L. and Magnus, P. (2007). Is there an elephant in the room? addressing rival approaches to the interpretation of growth perturbations and small size. *American Journal of Human Biology*, 19(5):606–614.

Stroud, B. (1979). Inference, belief, and understanding. *Mind*, 88(350):179–196.

van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Clarendon Press.

Wittgenstein, L. (1953). *Philosophical Investigations*. New York, NY: Macmillan. Translated by G.E.M. Anscombe.

Worrall, J. (1989). Fresnel, Poisson, and the white spot: The role of successful predictions in the acceptance of scientific theories. In *The Uses of Experiment*, 135–157. Cambridge: Cambridge University Press.

Worrall, J. (2000). The scope, limits, and distinctiveness of the method of 'deduction from the phenomena': Some lessons from Newton's 'demonstrations' in optics. *The British Journal for the Philosophy of Science*, 51:45–80.

Worrall, J. (2002). New evidence for old. In Gärdenfors, P., Woleński, J., and Kijania-Placek, K., editors, *In the scope of logic, methodology and philosophy of science*, volume I, 191–209. Dordrecht: Kluwer.

Notes

[1]Of course, scientific inferences are not generally offered as deductive. Einstein did not present the derivation as a demonstrative induction (Dorling 1995, §2). This is no mark against demonstrative theories, as accounts of how scientific inference *should* be understood and not of how it *was* understood by scientists; "inductions are *reconfigured* as deductive inferences with suppressed premises" (Norton 2003, 652, my emphasis). In the final section, I return to the issue of whether demonstrative theories are meant to be descriptive or normative.

[2]One might worry that the observation equivocates between two kinds of equivalence. Logical systems that share the same set of valid inferences are semantically equivalent. Proof systems that are intertranslatable are proof-theoretically equivalent. For sound and complete proof systems, however, this distinction does not matter. Most examples of demonstrative inductions are not fully formalized, but there is no reason to think that they could not be formalized in a language with a sound and complete proof theory. As such, we can ignore this complication in what follows.

[3]Wittgenstein makes a similar point when he observes that there must be a way of following a rule that is not an explicit interpretation of the rule (1953, §201); the Tortoise's trick is to insist that the rule must be interpreted explicitly for each instance before it can be followed. Stroud (1979, 195) also notes the affinity between Carroll and Wittgenstein.

[4]The distinction between high and low externality is similar to the distinction between observational and theoretical language. Although the latter might be more familiar to philosophers, it has the unfortunate effect of suggesting a binary distinction rather than a matter of degree. Even allowing that a report might be 'more or less observational', the distinction suggests a continuum with pure observation at one end and pure theory at the other and so creates unnecessary difficulties. For example, van Fraassen (1980) argues that the distinction between theoretical and observational language is a category mistake. No such charge can be leveled at the distinction between high and low externality reports.

[5]The high externality claim might be derived from background theory; e.g., the constancy of the speed of light may be derived from Maxwell's equations. But then the claim is a theoretical constraint rather than an observation report.

[6]Motivated by Quinean worries about meaning and translation, one might argue that there is no fact of the matter as to which of two equivalent proof systems an observed agent is using— and so that demonstrative theories of induction necessarily fail as description accounts. I have nothing helpful to say about the nature of translation here, but it suffices to note that such a Quinean argument would provide an independent route to the conclusion I argue for below. That line of argument would offer no help to champions of demonstrative induction.

[7]Of course, it is sometimes possible to fruitfully distinguish between theories and research programmes; e.g. Schell and Magnus (2007). For Worrall's criterion to be general, however, it must always or almost always be possible.

[8]Bain (1999) provides an illustrative example by reconstructing an argument by Steven Weinberg. Briefly, the argument is that "local quantum field theory is the only way to reconcile the principles of quantum mechanics with special relativity and Cluster Decomposition" (Bain 1999, 6). Bain thinks that this answers a specific kind of underdetermination worry by showing that there is no other theory that has all of these virtues (Bain 1999, 12). Yet underdetermination arguments typically turn on 'empirical equivalence', not equivalence with respect to strongly stated theoretical desiderata. An anti-realist who does not believe the principles of quantum mechanics will not believe their deductive consequences either. Weinberg's argument thus presumes a great deal of physical theory that could only be established by (partly) ampliative means. The ampliative support for some of the premises may be pushed back, but the contingency and inductive risk do not go away.

[9]A systematic discussion of what ampliative rules must assume about the world is given in Kelley (1996).

[10]The resolution of this issue will depend ultimately on what counts as evidence. It is no surprise that Bird (2006) begins his discussion by accepting a specific account of evidence.