Material Theory of Induction and Scientific Realism

Juha T. Saatsi*

26th February 2007

Abstract

John Norton has advanced a general view of induction—'Material Theory of Induction'—that renders ampliative reasoning in a deep sense local. This paper is a sympathetic appraisal of this view, applying it to the scientific realism debate. It argues that the scientific realist should turn to such local construal of ampliative reasoning in her attempt to justify beliefs about unobservables. More generally, the distinction that Norton draws between 'material' and 'formal' theories of induction is helpful in contrasting the intuitions behind various realist arguments, and in assessing their strengths and weaknesses. As far as justificatory challenges of induction are concerned, it is in this context that the Material Theory of Induction pays most dividends.

1 Introduction

Many have defended scientific realism against inductive scepticism. Science progresses by making various inferences about the unobservable world, and the realists aims to justify (some of) the resultant beliefs. The cognate arguments are multifarious and the debate is rambling; there is a real zoo of different positions. It has been said that there are as many realist arguments as there are realists.

This paper attempts two tasks. In the first place it compares and evaluates a broad spectrum of realist arguments. This is done by adopting a *meta-level* perspective, by discerning the basic intuitions behind different positions, and the types of arguments these intuitions lead to. We can ask, in particular, whether different realist arguments naturally couple with different general views on induction. Some recent work on induction is highly pertinent here. John Norton (2004) has characterised induction as *local*, rather than global; as *material*, rather than formal. Norton's meta-level analysis of induction can be employed to shed light on the whole realism debate. This is because the various realist positions can also be viewed as global or local, depending on the generality and form of their arguments. Hence a useful parallel can be drawn between the spectrum of realist arguments and the spectrum of general views on induction, regarding their strengths and weaknesses, and underlying intuitions.

Secondly, the paper attempts is to answer the following question. Given Norton's local understanding of induction—which I support, with some qualifications—how

*ADVERTISE HERE! Just kindly tell me what's wrong with this paper, and you're in... Seriously though: this is a rough draft. Please don't quote it.

should realism be defended against the selective inductive sceptic? Norton is optimistic about the reverberations of his theory with respect to Hume's general problem of induction. Despite my reservations in this regard, I am optimistic regarding the more narrow problem posed by selective inductive scepticism. I will argue that the local understanding of induction can to some extent reform the realism debate.

The paper proceeds as follows. The initial third of the paper gives an account of Norton's 'material' theory of induction, summarising his statement of the position in §2, and interpreting it further in §3. Then I will look at the zoo of realist positions in §4, first in the abstract and then in terms of concrete examples. Section §5 draws a parallel with Norton's analysis and evaluates its implications to the more global arguments for realism. The local *experimental* realist arguments are considered in §6, with the conclusions that the gist of these arguments suggests how the whole realism debate should be reformed.

2 Induction localised

Much of philosophy of science has focused on devising theories of induction. Since the ancient times the Holy Grail has been to find a single (or a small number of) principle(s) of induction that would lead inductively from true premises to a true conclusion, with some level of reliability. To this end philosophers have constructed elaborate theories of inductive generalisation, hypothetical induction, and probabilistic induction. (Norton, 2005) Norton calls such constructions *formal* theories of induction: they attempt to provide a formal schema to distinguish good, or licit, inductive inferences independently of case-dependent detail.

The term 'inductive logic' is most appropriate for theories of probabilistic induction. Much like deductive logic, these theories abstract away the content of particular propositions, and we are left with a formal structure of probabilistic notions that ought to apply universally to rational belief revision, say, in the face of evidence. Whilst theories of induction *not* framed in probabilistic terms—theories of inference to the best explanation, for example—are clearly not formal in *this* sense, these constructions are nevertheless formal in the sense of providing a universal *schema* which similarly abstracts away the content of particular explanations. Hence, we end up with general recipes like this:

> IBE: Given the evidence E, we inductively infer whatever would be the best explanation of E (if it were true), given our background knowledge B.

The advocates of inference to the best explanation argue that we can equally apply this schema to explain how the Big Bang theory is supported by the data, and how I make my circadian inference that a piece of bread in my hand nourishes rather than poisons me. What makes both of these very different inductions licit is the fact that our evidence is best explained by the conclusions inferred; both inductions are licit qua *instances* of IBE.

Norton provides a simple yet credible argument against the ideal of trying to achieve a satisfactory formal theory of induction: two millennia of effort hasn't resulted in a well functioning theory, so perhaps it is high time to recognise the Holy Grail as unachievable. And not only have all our efforts failed, but they have failed *for the same reason*. For any schema to function it always needs to be supplemented with local case-dependent detail. That is, there is an ever-present tension between universality and function. In Norton's diagnosis of this tension it is the local 'material postulates'

that ultimately do the work in licensing an induction, not the form per se. So, Norton advocates a *material theory of induction* instead of any formal one. For him all inductions are ultimately "underwritten by local material facts."

To understand what this means, consider the classic case of enumerative inductive generalisation, for example. There are several good inductive arguments of this form, of course.

> Sample A of lead melts at 327.5 ˚C
> Sample B of lead melts at 327.5 ˚C
> Sample C of lead melts at 327.5 ˚C
> ———
> Any lead sample melts at 327.5 ˚C.

This is (presumably) a licit argument. But this simple form of enumerative induction is, of course, illicit when the target of generalisation is a kind which is less homogeneous with respect to the property in question. Hence, one cannot thus infer the melting point of any sample of plastic, say. What makes the inductive argument above licit is not its form, argues Norton, but the fact that it is an enumerative induction *about lead* (or about an element). This fact is typically left implicit, but it can be included explicitly as a premise—'lead (as an element) is uniform in this respect'—rendering the argument (in this case) deductive, an instance of *demonstrative* induction. Norton calls such often unwritten local premises 'material postulates'.

Similarly, consider the following abductive argument, about the melting point of mysterious (temporarily) unclassifiable material, suddenly encountered around the globe:

> Sample A of the mystery material melts at 100.2 ˚C
> Sample B of the mystery material melts at 100.2 ˚C
> Sample C of the mystery material melts at 100.2 ˚C
> Sample D of the mystery material melts at 100.2 ˚C
> (BE) The best explanation for the regularity is that the mystery material has a constant melting point at 100.2 ˚C.
> ———————
> Therefore, any sample of the mystery material melts at 100.2 ˚C.

This is (presumably) a reasonably good argument. What makes it so? We should notice that there are a couple of local parameters that need to be fixed (to get the premise (BE)) in order to apply the schema IBE. First of all, we need to say what *counts as an explanation* in this context. Secondly, we need to say how different explanations are to be compared so as to make (BE) true. Hence, what makes the above argument licit is not the fact that it is *an instance of* the abstract schema of inference to the best explanation. Rather, Norton argues, it is the fact that in this local context our judgements regarding the two parameters are such that they ensure this schema really functions; that is, our judgements reflect the relevant ways the world is (the *facts*, the *local material postulates*).

Like the general recipe of enumerative induction, the abstract abductive schema IBE also furnishes many an illicit inference. In some contexts our explanatory judgements mislead us. We are prone to look for a causal explanation even when there isn't one, and any good conspiracy theory exemplifies how the schema gets misapplied by us. A well functioning formal theory of induction would tell us when the schema is applicable, and how the two parameters are fixed. Attempts to ameliorate the schema

in its universal form are doomed, however: our judgements about what counts as an explanation, and what counts as the best explanation, always get fixed locally, reflecting the underlying local material facts and the domain under investigation. If we want to capture what makes a particular abductive argument licit, we always need to include a local material postulate. This effectively states that in the particular context such-and-such explanatory virtues are also inductive virtues, reflecting the facts prevailing in the relevant domain.

More generally, these lessons about enumerative induction and inference to the best explanation arguably apply across the board. Consider, for example, the *Hypothetico-Deductive model* of confirmation with its notorious indiscriminateness: logic does not rule out arbitrary conjunctions being equally confirmed, and arbitrary disjunctions being equally confirming. For this scheme to function at all, this underdetermination needs to be tamed. According to Norton, all the proposed ways to achieve this turn on tracking local facts. For example, augmenting the HD-model by introducing considerations of *simplicity* cannot be done in general, universal terms, for 'our decisions as to what is simple or simpler depend essentially upon the facts or laws that we believe to prevail.' (2004: 656) In dealing with some cyclic phenomena we find sine and cosine functions nice and simple, instead of attempting to fit a linear curve, say.

Norton extends his thesis to Bayesianism as well: 'Bayesianism is vacuous until we ascribe some meaning to the probabilities central to it. Until then, they are just mathematical parameters. Each way of ascribing meaning brings factual presumptions with it.' (2004: 661) It is undeniable that the mathematical parameters need to be interpreted for the Bayesian probability calculus to model inductive reasoning, but it seems to me that this interpretation can be provided in rather universal, abstract terms. Hence Bayesianism does not seem local and case-dependent in quite the same way as, say, inference to the best explanation. Nevertheless, due to the complete openness of the prior probabilities Bayesianism all by itself is a rather weak theory, and perhaps it is best viewed as complementing, rather than competing with other theories of induction such as inference to the best explanation. (Lipton, 2004) Furthermore, given the focal point of scientific realism more needs to be said about the inductive method of science, since it doesn't seem that realism can be save by Bayesian considerations alone.[1]

3 What does Norton's theory achieve?

I have summarised Norton's explication of his slogans '*local*, rather than global', and '*material*, rather than formal'. But it is not easy to put a finger on what his theory exactly achieves. I will now try to provide some further interpretation.

It is important to distinguish here between *descriptive* and *justificatory* challenges. It is one thing to try to describe the way(s) we inductively reason—whether in everyday or the most advanced scientific life—and whole another thing to justify our ways of reasoning as profitable (truth-tracking, or empirical-adequacy-increasing, say). (Lipton, 2004) Norton's thesis about the locality of induction is first and foremost a descriptive one. His theory aims to locate a distinction between good inductions and bad inductions, without making any further claim whether are actually in a position to *know* which are which. It locates the distinction between licit and illicit not in the form (or any universally describable feature) of an inductive argument, but in its *content*. The lesson is that philosophers have been trying to find informative generality where there

[1] Douven (2005) aptly criticises a simple Bayesian argument for realism, and goes on to develop a sophisticated alternative. I don't think this works...

simply isn't any, and the Holy Grail has resulted from modelling inductive reasoning too closely on deductive reasoning (with its truly universal formal patterns). This philosophical theory about the distinction between good and bad inductive arguments does not in itself amount to knowledge that any particular inductive argument is good, since we may not know that the 'local material facts' really are *facts*, for any particular induction. So justification is a further question.

Nevertheless, Norton suggests that his theory does have interesting justificatory repercussions regarding Hume's problem of induction.[2] Hume's description of induction focused on enumerative generalisation, and thus his argument against the possibility of justification of induction naturally turned on the idea that enumerative generalisation hangs on the assumption of uniformity of nature. Although it is nowadays clear that enumerative induction is woefully inadequate as a description of our variegated ways of inductive reasoning, many think that Hume stated his argument in a general enough form for it to apply to any form of non-deductive reasoning. Norton disagrees. For the way that Hume's problem is typically presented relies explicitly on global and formal understanding of induction. Consider attempting an inductive justification of induction. We've got our first-order inductions about the world, and we've got a meta-induction about the past success of these first-order inductions. Such constructions, Hume's argument goes, are blatantly circular because both arguments are of the same *form*: 'more of the same'. If we are trying to thus establish the reliability of this formal schema of enumerative induction (irrespective of what the schema is applied to), we irrefutably end up running in circles.[3] But according to the material theory of induction, no induction is licit purely by virtue of its form anyway. So the classic circularity challenge is based on a misconstruction of the whole justificatory challenge.

In the material theory of induction, by contrast, a good induction is grounded on the facts correctly described by the material postulate. So justifying a particular induction is a matter of justifying the relevant material postulate. This material postulate cannot be just taken as given, and justifying a particular material postulate requires another induction. But this is a *different induction,* grounded on *different facts* described by *different material postulates*. No circularity ensues, and arguably our best actual inductions are background-dependent and local in exactly this way:

> 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. As we trace back the justifications of justifications of inductions, we are simply engaged in the repeated exercise of displaying the reasons for why we believe this or that fact within our sciences. (2004: 668)

But isn't there an obvious regress here? Norton is optimistic in this regard.

> What remains an open question is exactly how the resulting chains (or, more likely, branching trees) will terminate and whether the terminations are troublesome. As long as that remains unclear, these considerations have failed to establish a serious problem in the material theory analogous to Hume's problem. (2004: 668)

[2] I will briefly mention this line of thought, and my reservations about it, but I do not wish to get too embroiled in this debate. (See Okasha, 2006, and references therein) For whatever its outcome is, my appropriation of Norton's descriptive thesis to the realism debate is equally valid.

[3] A logical possibility is to try to justify one schema by applying a meta-induction of a different schema, to be justified by applying a meta-meta-induction of a yet different schema, and so on. An obvious regress ensues, and it is not clear what all these different schemas really are.

The problem is different from Hume's problem, for sure, but it seems equally damaging to me. It does not seem reasonable that the justification of any induction depends on justification of further facts in this way. Consider the paradigm induction of the sun rising tomorrow. For pre-scientific human beings it was a basic regularity of the world. We now justify this regularity by appealing to different, science-discovered facts as material postulates. These facts about gravitation and dynamics are more general than the facts about the sun and the Earth per se, but they are still local by virtue of not being a priori universal postulates about the worldly uniformity, but local postulates concerning the dynamical-gravitational aspects of the world (and not the greenness of the emeralds, say). But how is the science-informed starting point—a more general basic regularity—any less problematic *qua* basic regularity? Both regularities are inferred from a finite set of experiences.[4] And although there is a sense in which we have scientifically justified what was taken to be a primitive regularity beforehand—and hence the scheme '*describes* the routine inductive explorations in science'—the philosophical challenge of justifying induction concerns the respective starting points. I do not see any indication that that the regress isn't going to be infinite and vicious. For whilst we can give reasons for this or that regularity by appealing to different, broader regularities, I don't see any reason to assume that those reasons will converge to our experiences. What reason do we have to believe that our best science will still work tomorrow?

In sum, I take Norton's *descriptive* thesis about the locality of inductions, of their inherent background dependency, to be significant. Regarding justification, it does show that one needs to be more careful how the all-out justificatory challenge is posed. A typical two line statement of *The* Problem of Induction is not inline with the fact that we do not have a universal formal schema (or a set of schemas) to capture the difference between licit and illicit inductions. However, an equally difficult problem of justification may remain.

Finally, a clarificatory point regarding the status of global and formal theories of induction that do not take into account Norton's thesis. Although the locality of induction needs to be acknowledged, it does not by any means render the descriptive work done at higher-levels of abstraction wholly redundant. For example, we can gain significant insight in induction by modelling it in Bayesian or abductive terms, as long as we keep in mind that these descriptions are not the whole story, but gained by abstracting away some content that is local and an essential part of what makes the induction licit. But all in all, Norton's emphasis on locality is certainly not misplaced.

* * *

Whether my reservations regarding Norton's dissolution of Hume's problem are warranted or not, we can use Norton's insight vis-à-vis a more limited problem of induction. First of all, the distinction between formal and material can be used to evaluate and throw light on the plethora of realist arguments. Secondly, the local understanding of induction can to some extent reform the realism debate by more explicitly spelling out the intuitions and motivations behind the various (local) experimental realist arguments.

[4]Norton rejects the 'simple argument that that such brute facts are always singular and that no collection of singular facts can license a universal', but I don't understand his argument.

4 The Realist Zoo

The realist arguments considered here aim defend realism against selective inductive scepticism, not against Hume's all-embracing scepticism.[5] The realist does not have to possess an argument against Hume, if we modestly begin with the premise that we have got some substantive inductive knowledge of unencountered observable affairs (the sun rising tomorrow, Jupiter having moons etc.). Assuming this much, the challenge is to argue that we should also take seriously some ampliative inferences to the unobservable.

The first task is to understand and sort out the various intuitions and motivations behind the different realist responses to this challenge. There's a real zoo of arguments out there. Here are a couple of initial observations.

1. Mirroring the form of the sceptical challenge, the basic form behind each and every realist argument is this:

> We are happy with *such-and-such* inductions $\{I_O\}$ about the observable.
> *Such-and-such* inductions about the unobservable $\{I_U\}$ are *relevantly similar* to $\{I_O\}$.
> ———————-
> We should be happy with $\{I_U\}$.

The various arguments then differ with respect to how inductive inferences are construed and classified, to spell out what the two classes of arguments $\{I_O\}$ and $\{I_U\}$ are, and in what sense they are 'relevantly similar'.

2. Different realists are driven by different understandings of the inductive method of science. Inference to the best explanation is often the central point of contention. Some say it is a unifying feature of all scientific reasoning, and realism turns on arguing that explanatory virtue is a truth-tracking virtue. The most optimistic line of thought appeals to IBE in a rule-circular fashion at the meta-level. Others deny the optimistic meta-level application of IBE, and argue for realism directly on the grounds that scientific inferences are explanation-driven. Yet others argue for realism without appealing to explanation at all, preferring to leave it open whether any (extant) descriptive scheme captures the inductive method.

3. The aim and scope of the realist arguments differ. Some wish to produce a single overarching argument that does it all, once and for all. Others are happy to argue for realism about this or that in a more piecemeal fashion. We can talk about *global* and *local* realist arguments, depending on their scope.

We should try to analyse the whole gamut of different arguments. Can we say something interesting about the way the realists differ in their attitudes towards induction? Can we analyse the pros and cons of these different leanings, by abstracting away from the details of particular arguments? We can bring Norton's novel understanding of induction as a fundamentally local business to bear on this task. But let's first take a bit closer look at the spectrum of realist arguments, first in the abstract and then in reference to some specific arguments.

[5] A kind of selective scepticism is typically attributed to van Fraassen.

4.1 The realist spectrum

Realist arguments can be usefully located on a *spectrum* spanning from global to local. The word 'spectrum' indicates that locality/globality is a matter of degree. We can also speak of global and local *tendencies* to signal one's realist orientation.

What characterises the global tendency of a realist argument is the attempt to justify some inductive inferences by reference to some rather general attribute unifying all these inferences. The local arguments, by contrast, take there to be more justificatory analysis to be done on case-by-case basis. Locality comes in degrees. A set of inferences can be unified by virtue of some single characteristic/form that acts as the vehicle of justification for each instance of inductive inference featuring that characteristic/form. Corresponding to the level of generality at which such characteristic/form is described—how encompassing the set of such inferences is—we have more local and less local realist strategies. This abstract preliminary distinction between global and local is best clarified via concrete examples of actual realist positions.

Advocate	**Argument**	**Tendency**
Boyd / Psillos	No-Miracles Argument	100% Global
Lipton	1st-order explanationism	Rather global
Kitcher	Galilean Strategy	Rather global
Hacking / Achinstein	Experimental Realism	Rather local
?	Arguments of Science	100% Local

Where would Norton fall on this spectrum, given his material theory of induction? Since according to Norton the form of an inference *alone* never makes it licit, it would seem that it is never going to be enough for the realist to focus on mere form alone to unify some inferences across the observable-unobservable boundary. So would his preferred argument go at all beyond the first-order arguments of science? Do these scientific reasons in themselves constitute a philosophical argument at all? I'll come back to these questions towards the end. Next I will detail some examples furnishing the table above, before drawing a parallel between this analysis and Norton's meta-level analysis of induction.

4.2 Examples

100% Global. The standard explanationist arguments for realism, originating from Putnam, and finessed by Boyd and Psillos amongst others, are *fully global*. Psillos (1999) is a notable recent author in this lineage. He argues (roughly speaking) that the scientific method is based on inference to the best explanation, and that by a meta-level use of IBE we can (in the externalist epistemological framework) justify the scientific use of IBE as truth-tracking. Hence, we've got an attempt to justify realism by reference to an extremely general attribute unifying all scientific inferences to unobservables. Namely, they are all of the same form: IBE.[6]

There are a couple of noteworthy ideas underlying Psillos's global explanationism. First of all, Psillos (following Harman 1965, and Josephson 1996, 2000) takes IBE to be a fundamental, primitive foundational form of inductive inference (see Psillos 2002,

[6]Admittedly there are many subtleties to Psillos argument, regarding the rule-as-opposed-to-premise-circularity of the meta-level justification, for example, and the fact that Psillos allows for different degrees of confirmation: not all explanatory inferences are epistemologically on a par as far as their confirmatory strength goes. But the basic form of the justificatory argument is this, and it is global.

for example). Secondly, Psillos explicitly appeals to the similarity in the *form* of our reasoning about the observable and the unobservable matters, displaying clearly his underlying intuitions:

> Theoretical beliefs in science are formed by means of abductive reasoning. But so are most of our every-day commonsense beliefs. Realists have exploited this fact in order to argue that if one has no reason to doubt commonsense abductive reasoning, then one should have no reason to doubt abduction in science. The *pattern of reasoning*, as well as justification, are the same in both cases. (Psillos, 1999: 211, my emphasis)

And on these grounds Psillos accuses the selective sceptic (van Fraassen) of adopting a selective attitude against inferences about the unobservable:

> Clearly, van Fraassen sustains a selective attitude towards IBE. The latter is a means of going beyond the realms of what has been actually observed and forming warranted beliefs about unobserved things and processes. Yet IBE is not a means of forming warranted beliefs about the realm of unobservable things or processes. (Psillos, 1996a: 34)

Rather Global. Other realists are wary of a meta-level application of IBE, and also of the idea that there is pertinent justificatory unity to *all* scientific inferences that can be viewed as abductive in form. Yet some of these realists wish to tap into the pivotal explanatory dimension of science, and appeal to explanatory virtues in a less global way. These realists argue that the gap between ampliative inferences to observables, on the one hand, and to unobservables, on the other hand, is bridged by virtue of the fact that the respective inferences are *not only* of the same general form (IBEs), *but also* of the same more specific 'inferential kind'.

Lipton (2004) presents such an argument. He develops an overarching *descriptive* account of confirmation and induction in terms of inference to the best explanation.[7] Regarding the justificatory challenge of realism, he puts forward a very general argument to unify and justify a significant class of abductive inferences of science. After repudiating the No-Miracles Argument, Lipton considers a less global, first-order explanationist strategy:

> Can explanationism defend realism instead by appeal to the *structure* of those first-order inferences? ... The *structure of causal inferences* is the same, whether the cause is observable or not. ... So there is a prima facie case for saying that all these inferences should be construed in the same way: granting the truth-tropism of inferences to observable causes, we ought also all to be realists about inferences to unobservable causes, since the *inferences have the same form* in both cases. (2004: 199–200, my italics)

Although Lipton avoids a meta-level global abductive inference about science, he still provides a very general template for the justification of scientific inferences. For him any scientific first-order instance of causal abduction is (probably) approximately true by virtue of being '*structurally similar*' to everyday ampliative reasoning about

[7] Unlike Psillos, Lipton is not a totalitarian 'IBE fundamentalist', claiming that all inductive inferences are best construed as abductive. Rather, for him inference to the best explanation simply plays a significant role in understanding inductive reasoning.

the observable. *Pace* Psillos, Lipton takes it to be incumbent on the realist to provide a more specific description of the kind of abductive reasoning that allows us to generalise from the everyday theorising to the scientific. Hence Lipton stresses the *causal-contrastive* mode of IBE. But is that enough said? I will return to this below.

How local is this species of justificatory argument? It depends on how tightly the relevant 'inferential kind' is delineated. Just appealing to causal explanations (spelled out as contrastive explanations – cf. Lipton 2004, ch. 3) yields a rather global argument. The notion of contrastive causal explanation is a broad one, even at the level of observable matters.

Rather local. At the local end of the spectrum we have got a set of arguments whose advocates are collectively known as *experimental* realists.[8] Experimental realists do not (have to) advocate any level of explanationism, not even as a significant descriptive thesis about the scientific method. Traditional questions about general characterisation of induction are simply irrelevant to their realist arguments, for these arguments for realism about some theoretical posit do not rely on the idea that a scientific inference to this posit has a particular form. Rather, these local arguments rely on domain, or case-specific considerations, typically closely following the reasons that scientists themselves supply for their beliefs about something unobservable. For example, regarding the classic paradigm of an unobservable entity, the atom, it has been popular to examine Perrin's original reasoning to the existence of atoms on the basis of Brownian motion. (Achinstein, 2002; Miller, 1987; Salmon, 1984)

Although the individual realist arguments differ considerably in detail and rhetoric, in my interpretation of them there is a common pattern to be seen. All the "experimental" arguments—I will question the aptness of the title below—seem to rely on some local basic assumptions about the uniformity of the world, crossing the observable-unobservable boundary. The unwritten premise is that such uniformity assumptions are simply as innocent as the assumption required for us to induce, say, the sun rising tomorrow. The uniformity assumptions pertain to local matters, and the corresponding epistemic warrant is thus localised.

Achinstein (2002) is the latest (and the clearest) representative of this line of thought. He analyses Jean Perrin's reasoning to the existence of atoms as *causal-eliminative*, also giving a more general account of the conditions on which this kind of reasoning is justified. Achinstein presents Perrin's reasoning as follows (2002: 474)

> 1. Given what is known, the possible causes of effect E (for example, Brownian motion) are C, C_1,..., C_n (for example, the motion of molecules, external vibrations, heat convection currents).
> 2. C_1,..., C_n do not cause E (since E continues when these factors are absent or altered).
>
> So probably
>
> 3. C causes E.

Observing the microscopic particles dancing around, continually accelerating and decelerating, indicates the existence of internal forces responsible for such behaviour,

[8] Some of the *entity* realist arguments (e.g. Hacking, 1982.) are also naturally interpreted as belonging to this category.

assuming that no plausible external cause can be found. And the meticulous experiments performed by Guoy did indeed allow Perrin to eliminate the plausible external candidate causes C_1,..., C_n. The various experiments performed by himself and others then allow Perrin to claim quantitative evidence for his initial conclusion and for the numerical value of Avogadro's constant.

This reasoning raises two obvious anti-realist worries. First of all, there is the possibility that the hypothesis of internal molecular forces singled out by the eliminative reasoning is merely the best of a bad lot. How do we know that all the possible alternative causes of the phenomenon have been cited and eliminated by experiments? Achinstein's response is to insist that the realm of possibility here is restricted by our background knowledge.

> The claim that the possible causes cited probably include the actual one can be defended by appeal to the fact that the phenomenon in question is of a certain type that, experience has shown, in other cases is caused by one or the other of the causes cited. (2002: 478)

But this immediately raises the second anti-realist worry: how can we justify inferences to the unobservable on the basis of the observable, on the basis of what 'our experience has shown'? For example, in Perrin's argument we need to justify the inductive generalisation from 'All observed accelerating bodies in contact with other bodies exert forces on them' to 'All accelerating bodies, including molecules (if any exist), in contact with other bodies exert forces on them' (ibid., 481). And empiricists like van Fraassen, of course, take such inductive inferences to the unobservable to be unjustified and unjustifiable.

Achinstein responds to this second worry as follows, bringing the locality of his argument to the fore. The idea is that the realist can provide a positive empirical reason for taking observability *not* to be a biasing condition for an inductive generalisation from a sample.

> One can vary conditions or properties in virtue of which something is observable (or unobservable). For example, items can be observable (or unobservable) in virtue of their size, their distance from us in space or time, their duration, their interactions (or lack of them) with other items, and so on. ... If we vary the conditions in virtue of which bodies are observable and find no differences in whether bodies have mass, and if we have no contrary empirical information, then we have offered an empirical argument to support the claim that the fact that all observed bodies are observable does not bias the observed sample with respect to the property of having mass. (ibid, 484–485)

Hence, the anti-realist's selective scepticism should feel some tension here. In particular, the kind of variation in the conditions and properties that the realist here appeals to arguably also count for the legitimacy of ampliative inferences about unobserved observables. So whence the difference? After all, the logical possibility of *observability* being a biasing condition is on a par with the logical possibility of *having been observed* being a biasing condition.

The locality of Achinstein's argument resides in the fact that it concerns only a very particular uniformity of the world: the relevant mass-related properties (e.g. conservation of momentum for massive bodies) are independent of the properties in virtue of which bodies are observable, or otherwise. Achinstein's realist analysis takes explicitly

into account those local matters of fact which underwrite Perrin's inductive argument for the existence of atoms. And his response against the selective sceptic turns on a kind of *Tu Quoque*: arguably our inductive inferences about the unobserved (but not unobservable) massive bodies are underwritten by material postulates that are epistemologically no different from the ones used by the local realist.

* * *

Many arguments in this last class are notoriously imprecise and rhetorical (but also intuitively pulling, e.g. Hacking, 1982), and a certain amount of interpretation is required. The interpretative gloss presented here, emphasising the locality of the arguments, gains impetus from the following observation. Many have reacted to these arguments by objecting that they are *just as abductive in form* as the standard 'theory-realist' argument based on the success of science ('the miracles argument'). (e.g. Resnik, 1994; Psillos, 1999) These reactions are unsurprising: if one is looking for a global and formal justification of induction, the experimental realist arguments seem to rely on an abductive form. However, the intuition driving the local realist is that the justificatory work is done at the level of local material assumptions. The fact that a realist argument turning on a material assumption can be naturally construed as an inference to the best explanation is wholly irrelevant, given the very different view on what makes an induction licit.

5 Parallels with Norton's analysis of induction

Having sketched some prominent arguments on the realist spectrum, we can now discern interesting parallels with Norton's analysis of induction. Given the foregoing (interpretation-laden) outline of the various positions, this is straightforward.

The (more) global arguments are driven by (more) emphasis on (more) formal similarities, whilst the domain/case-specific details are downplayed. The seeming advantage of these arguments is that one gets more with less: a justification of a significant class of scientific inferences by their shared form, without having to pay much attention on what these inferences are about. The downside is an increased epistemic risk which, I will argue, is unacceptable. By contrast, the (more) local arguments have (more) emphasis on domain-specific factual matters. Local realists are happy to admit that justification of knowledge of the unobservable world is a business that always hangs on local matters of fact driving the ampliative inferences. Although this is not how the "experimental" arguments are typically portrayed, I will claim that this is the best way to cash out the underlying intuition. (Section §6).

Given this parallel, we can provide a meta-level argument against the global tendency in realist arguments. The global realist arguments suffer from what could be termed the *Description–Justification Gap*. Too much emphasis is paid on formal descriptive unity, without realising that descriptive unity can be cheap, and does not amount to justificatory unity. The former can be achieved at the level of abstract induction schemas, but the latter requires more. Since a licit induction is always underwritten by a material postulate, we need to make sure to give some justification at that level as well. But this is exactly what is missing in the realist arguments which attempt to cross the gap between the observable and the unobservable by comparing the respective inductive inferences vis-à-vis their *form*, or *structure*. Worse still, the ambitious meta-level use of inference to the best explanation appeals to a descriptive unity on a

much wider scale, now spanning from scientific to *philosophical* explanations.[9] Given how open the required characterisation of IBE is, the two parameters (cf. section §2) determining what counts as an explanation, and what counts as a good explanation, are left wide open. But the mere form of an inference cannot carry the justificatory burden.

It is insufficient for the realist to *not have* any reason to think that a particular inference form is unreliable in some unobservable domain, given its reliability in some observable domain. More than that is required: a *positive* reason to think that the respective inferences are in same epistemological boat. To have a reason to suspect an inductive inference in a particular theoretical context requires a reason to suspect that the material postulate underwriting that inference does not correspond to the worldly facts. For example, we have such grounds to suspect the inductive generalisation from 'All Turkish adult males I've encountered are bearded' to 'All adult male Turks are bearded'. We know enough of *human beings* to know that nationality simply isn't a strong enough unifying factor in this respect. But *not* having such negative grounds for suspecting an induction does not amount to having positive grounds for it, either. For example, we may not have any *particular* reason to suspect that scientists' evaluation of the explanatory virtues in quantum physics is any less reliable than farmers' evaluation of the explanatory virtues required to catch a flock-harassing beast. But we may not have any positive reason to think that the explanatory virtues are on a par as inductive virtues either. And surely the abductive form of the respective inferences isn't enough on its own, given the huge difference in the domains and the kinds of inferences made.[10]

Norton describes a tension between functioning of a descriptive inference schema, and its universality. The global justificatory arguments are under pressure to go more local, too, to rule out illegitimate use of cheap descriptive generalisations. For this reason (I think) Lipton focuses more narrowly on *causal-contrastive* explanations. But this is still an attempt to justify a rather wide class of explanation-driven inferences by a single argument, corresponding to the rather open notion of causal explanation (understood in contrastive terms). Although the unifying characteristic here is not purely formal—a causal explanation obviously needs to reflect a causal fact about the world—it does not seem that this alone captures what makes each instance of causal-contrastive abductive inference licit. There is still much contextual variability in how the best explanation is chosen.[11] One could respond by further narrowing down the class of abductive inferences by fixing more case-dependent variables. This clearly amounts to more local realism. But how local should we go?

6 The realism debate reformed?

The justificatory challenge for the realist, as described in section §4, is to argue for a *Unity of Inductions* that makes selective scepticism unnatural and unappealing. The

[9] Explaining the success of the scientific method by its truth-tracking ability is a philosophical explanation, albeit a naturalistic one.

[10] See also Magnus's critique of Kitcher's Galilean strategy (2003). Kitcher's strategy is quite global, and subject to corresponding difficulties. The main difference between Kitcher and the rather global arguments considered above is that Kitcher does not operate at the level of forms of inductive inferences, but rather finds the relevant unity at the level of their success-conditions.

[11] According to Lipton 'for the causal explanations of events, explanatory contrasts select causes by means of the Difference condition: *To explain why P rather than Q, we must cite a causal difference between P and not-Q, consisting of a cause of P and the absence of a corresponding event in the case of not-Q*' (2004: 42) This is clearly a rather open characterisation of what is required of these explanations, and much hangs on case-dependent detail. For example, the notion of 'corresponding event' is highly contextual, and get fixed by the situation at hand.

literature contains a great variety of ways to argue for such unity, and I have argued that the intuitions behind the different arguments correspond to more local and more global understanding of induction. This offers a useful way to order the sprawling debate, for comparison and evaluation of the alternatives.

Assuming that Norton's insight about the locality of induction is correct, what are its repercussions on the realism debate? For one thing, there seems to be a serious problem with the global arguments. This is due to the gap between achieving a descriptive unity and achieving a justificatory unity: we haven't been given any positive reason to think that the kind of descriptive unity that the global arguments capitalise on amounts to a relevant justificatory unity. This pushes the realist towards the local argumentative strategies. I have argued that the "experimental" realist arguments can be viewed as (rather) local arguments. But this raises further questions. What is the best way to construe these local arguments in *general* terms? Exactly how local are they? I'll finish the paper with some tentative remarks on these issues.

The material theory of induction acknowledges that any licit inductive argument has both a form, and an underlying material postulate. The global realist arguments try to argue for the Unity of Inductions at the level of the shared form, whilst the local realist arguments are construed as depending on an analysis of the relevant material postulates. Hence, the justificatory work is done by comparing the material postulates corresponding to inductions-to-observables, on the one hand, and inductions-to-unobservables, on the other. But this general way of putting it makes it clear that these arguments have nothing to do with "experiments", or "entities" *per se*. Rather, they have to do with a local-as-opposed-to-global, material-as-opposed-to-formal, comparison of the respective inductions.

But the material postulates underwriting inductions to the observable and to the unobservable, respectively, are still going to be different, of course, so there is no question of *identifying* the postulates required by the realist with those required by the selective sceptic. The best one can do is still a judgement of the *naturality*, or otherwise, of drawing the line of epistemic incredulity at some point. This is how realists have always argued, admitting that there is always ample *logical* room for inductive scepticism, selective or not. This is just the nature of induction *qua* non-deduction.

But of the various ways of arguing against the unnatural scepticism of the anti-realist, the local approach, I maintain, is the best. This follow from Norton's insight. For if the local material facts are what make an induction licit, then the realist takes an unnecessary epistemic risk by appealing to descriptive unity. By appealing to the form of an inference, instead of its material postulates, raises the possibility that an inductive inference is taken to be licit when there is no relevant material fact to underwrite it. Of course, the stronger the appeal to descriptive unity, the higher the epistemic risk. But the absolute minimum—corresponding to the strongest realist arguments—is achieved by focusing on material postulates themselves.

But how local are these arguments then? Do they go beyond the first-order scientific reasoning at all? Sure they do. The local realist arguments are bona fide philosophical arguments. Scientists latch onto the correct material postulates by the methods of science which may or may not make the material postulates transparent. If a scientist appeals to a theory T because it is the simplest and the most unifying, it a task for the philosopher to make explicit how these contextual judgements of simplicity and unification reflect the local material facts, given the scientific background knowledge of the domain in question. Only once material postulates have been made transparent can we compare them with the sorts of local assumptions that are needed to underwrite inductions to the observable. Hence, although the specific arguments of local realism

hang on case-dependent detail, its master plan can be described in general terms. This presents a new challenge for the philosophers of science, reforming the realism debate. The recurring question is: can we argue for realism about this, or that, in terms of local material postulates?

To conclude by paying homage to Norton's insight, it can be pointed out that at times his writing comes very close to what I have said above. For example, Norton's discussion of the control of inductive risk is suggestive:

> As long as the inductive risk resides within the schema, we must assess it through a highly problematic judgment of the overall reliability of the relevant schema. We have little chance of coming to a clear judgment let alone determining how to reduce the risk. However once the risk is relocated in a material postulate in some local domain, our assessment of the inductive risk will depend in large measure on our confidence in the material postulate. If the inductive risk is great, we now also have a program for reducing it. We should seek more evidence relevant to the material postulate and perhaps even modify the material postulate in the light of the evidence. The result will be a more secure induction. (2004: 665)

This guidance is given in the context of how *science* uses induction. My claim is that the realist philosopher arguing for realism should follow suit. And this is exactly the spirit of the local realist arguments: rely more on detailed scientific practice, on the first-order reasons that convinces a scientist, knowledgeable of those 'local material' details, of the existence of something.

References

Achinstein, Peter (2002). Is There a Valid Experimental Argument for Scientific Realism? *Journal of Philosophy*, **99**(9), 470–495.

Douven, Igor (2005). Evidence, Explanation and the Empirical Status of Scientific Realism. *Erkenntnis*, **63**, 253–291.

Hacking, Ian (1981). Do We See Through a Microscope? *Pacific Philosophical Quarterly*, **62**, 305–322. Reprinted in Images of Science. (Eds) P.M. Churchland and C.A. Hooker.

Hacking, Ian (1982). Experimentation and Scientific Realism. In *Philosophy of Science: The Central Issues* (ed. M. Curd and J. Cover), pp. 1153–1168. W.V. Norton and Company, 1998, New York.

Lipton, Peter (2004). *Inference to the best explanation* (2nd edn). Routledge, London.

Magnus, P. D. (2003). Success, Truth and the Galilean Strategy. *British Journal for the Philosophy of Science*, **54**(3), 465–474.

Norton, John D. (2000). How Do We Know About Electrons. In *After Popper, Kuhn and Feyerabend; Recent Issues in Theories of Scientific Method* (ed. R. Nola and H. Sankey), pp. 67–97. Kluwer, Dordrecht.

Norton, John D. (2003). A Material Theory of Induction. *Philosophy of Science*, **70**, 647–670.

Norton, John D. (2005). A Little Survey of Induction. In *Scientific Evidence: Philosophical and Historical Perspectives* (ed. P. Achinstein), pp. 9–34. Johns Hopkins University Press.

Okasha, Samir (2005). Does Hume's argument against induction rest on a quantifier-shift fallacy? *Proceedings of the Aristotelian Society*, **105**(2), 253–271.

Psillos, Stathis (1996). On van Fraassen's Critique of Abductive Reasoning. *Philosophical Quarterly*, **46**, 31–47.

Psillos, Stathis (1999). *Scientific Realism: How science tracks the truth*. Routledge, London.

Psillos, Stathis (2002). Simply the Best: A Case for Abduction. In *Computational Logic: From Logic Programming Into the Future* (ed. A. Kakas and F. Sadri), Berlin, pp. 605–625. LNAI 2408: Springer-Verlag.

Reiner, Richard and Pierson, Robert (1995). Hacking's Experimental Realism: An Untenable Middle Ground. *Philosophy of Science*, **62**, 60–69.

Resnik, David B. (1994). Hacking's Experimental Realism. *Canadian Journal of Philosophy*, **24**, 395–412.