This is a draft paper. Please contact one of the authors before citing or referencing

Contrast, Inference and Scientific Realism

MARK DAY and GEORGE BOTTERILL

Sheffield, September 2006

I: ARGUING FOR SCIENTIFIC REALISM

Scientific realism is motivated by desires to make claims about both the ontological aspirations of science and its epistemic achievements. On the ontological side, the realist maintains that at least some scientific theories do try to inform us about the constituents of matter and the entities and forces which generate natural processes. Science is not just a systematic attempt to "save the phenomena". Although there have been many opponents of ontological realism in the history of philosophy of science, who have believed, for one reason or another, that it is desirable to rein in the ontological commitments of scientific theorizing, there are good reasons for preferring ontological realism. One might, after all, say that it is more an anthropological observation than a philosophical claim that we hope our theories can describe the structures and forces in the world and capture the laws which govern nature.

However, it clearly is not enough for scientific realism to urge the ontological commitments of some scientific theories. For presumably there were Vikings who were ontologically committed to Odin, Thor, and Valhalla. Science is more deserving of credence than mythology and other non-scientific belief-systems. A robust scientific realism will want to *endorse* the ontological commitments of modern science. But getting this endorsement right can seem like a problem of some delicacy, with the realist running the risk of appearing naively optimistic about the correctness of current scientific theories. We think that a correct appreciation of the role of contrastive explanation in science and its connection to causal explanatory inference can help scientific realism get the right mix of ontological and epistemic commitment.

In this paper we aim to explore the connections between a certain form of explanatory inference and scientific realism. The attempt is prompted by dissatisfaction with more familiar ways of defending scientific realism. In general, two different kinds of argument have been deployed, an *argument from experimental interaction* and an *argument from theoretical success*. However, there are reasons to doubt whether either of these kinds of argument can provide convincing support for a robust form of scientific realism.

Taking the argument from experimental interaction first and Hacking's clarion call 'If you can spray them they're real!'[1], it is easy enough to see why this line of reasoning cannot be entirely independent of assumptions about the correctness of theoretical interpretations. It is no doubt salutary for philosophers of science to be reminded that fascination with theories should not blind them to the significance of experiments. And it's certain that what we can interact with really exists. But this only prompts the further question as to how we can be sure about what interactions are occurring in the course of an experiment. Different experimental procedures may put this issue in a rather different perspective, as we intend to demonstrate. Yet in the final analysis claims about interaction seem dependent upon theoretical interpretations, and so dependent upon the adequacy of our theories. After all, Priestley believed that he was interacting with phlogiston. So advocates of the argument from interaction need to explain why Priestley's isolation of 'dephlogisticated air' does not show that he was actually right to believe in such a substance as phlogiston. It looks, therefore, unlikely that an *argument from interaction* for scientific realism can be developed in a way which is independent of considerations of the correctness of theoretical interpretations.

For this reason, some form of *argument from theoretical success* has remained the most popular way of trying to support the epistemic aspirations of scientific realism. Arguments from theoretical success appeal to the predictive and technological successes of scientific theories and contend that the best way of explaining such success is to take the theories to be true and the entities and forces they appear to refer to as actual existents. Much of the discussion about scientific realism has been dominated by the 'No Miracles' or Inference to the Best Explanation argument, regularly attributed in the first instance to Hilary Putnam[2] and famously contested by Larry Laudan[3]. In our view it would be unfortunate if debates over scientific realism should become fixated on discussion of the 'success' of theories and what can be read off such success in terms of ontological adequacy, for at least two reasons.

One of these reasons is that there is something problematic about the idea that there might be distinctively philosophical arguments for the truth of scientific theories that are different from — or could possibly be any improvement on — the actual considerations used in the course of scientific investigation itself. In other words, if scientific realism is arguing that scientific theories are true, it must of course mean *some* scientific theories (the preferred and successful ones) rather than other contenders. But in that case what could make the arguments of scientific realism any different from the reasoning deployed in science that leads the practitioners themselves to prefer some theories to others in the first place?

It may be that there are versions of the argument from theoretical success which can allay this worry. But there is in addition one very general difficulty with any form of argument from theoretical success, and that is the issue of underdetermination. It has long been a paradigmatic tenet within philosophy of science that no finite body of empirical data, however large, can fix any one single theory as the theory which alone explains, accommodates, or predicts all that data. This is maintained on the grounds that given any finite body of data there will always be many other theories (perhaps even infinitely many theories) equally compatible with that body of empirical evidence and all equally implying (with suitable initial conditions) any assertible data-statements. And all these theories, therefore, are apparently equally capable of explaining the data, unless we impose strong quality controls on what we are prepared to allow as an explanation.

Once this principle of the underdetermination of theory by empirical data is accepted, it leads directly to a powerful argument against any attempt to justify epistemological realism by using an argument from theoretical success. A theory that we have hit on may be successful in explaining and predicting all the data we will ever be able to obtain. But if many, many other theories could equally do this, then such predictive and explanatory success does little to show that our theories are true. It is true that there may be other differences between empirically adequate theories (theories which are capable of accounting for all available data) in such respects as simplicity, ontological parsimony, mathematical elegance, and general aesthetic appeal. But nobody has been able to provide any convincing argument why such non-empirical theoretical virtues should be reliable indicators of truth.

So if we do accept the principle of the underdetermination of theory by evidence, it appears that no robust version of epistemological scientific realism can be rationally supported. The predictive and explanatory success of our theories shows only that they are among the set of empirically adequate theories, not that they are the uniquely true ones. And of course this is precisely the argument that van Fraassen has advanced in favour of preferring his *constructive empiricism* (the view that the aim of science is to arrive at empirically adequate theories and that we can have empirical grounds only for accepting them as such) to a more ambitious scientific realism which claims that we can be justified in believing scientific theories to be true (or approximately true).[4]

How should we react to the fact that the thesis of underdetermination is an obstacle to the epistemological aspirations of realism? An obvious rejoinder, to anyone at all familiar with the history of scientific theorising, is that the thesis of underdetermination is not only unsupported by the history of science, but even appears to be a philosophical doctrine of peculiar perversity. For it is often remarked that if one examines any domain of science in which there are competing theories, we are not likely to be troubled by an embarrassment of empirically adequate riches. On the contrary, the normal situation is that not a single one of the theories on offer is capable of accounting for all the data in the domain.[5] Feyerabend maintained that theories are always born in an ocean of anomalies[6], and continuing to sail in such choppy waters is a common enough life-history for scientific theories. However, it seems premature to reject the underdetermination thesis on the grounds that in most, if not all, scientific domains we cannot point to any actual example of an empirically adequate theory. For there is a worry that this situation might only be testimony to the poverty of our imagination. How do we know that there are not vast numbers of empirically adequate theories which to date have not occurred to any human scientists? To assume that current science must already have arrived at the best theories would appear to be begging the question in favour of epistemological realism.

Underdetermination is at the heart of van Fraassen's provocative critique of Inference to the Best Explanation (IBE). He argues that it cannot be rational to use IBE as a rule for belief-fixation because it 'only selects the best among the historically given hypotheses. ... So our selection may well be the best of a bad lot.'[7] While there may be good arguments against van Fraassen's attack on IBE, it is clear that it would hardly be sufficient to challenge the underdetermination thesis on grounds of fit with actual scientific theorizing. One might therefore conclude that scientific realists urgently need to defend IBE against van Fraassen's critique. For it might seem that unless we can claim that acceptance of some scientific theories is warranted on the grounds of IBE, we have no way of justifying the epistemological aspect of scientific realism. However, we think a better approach is to start by establishing a more carefully articulated account of how we can reason from the phenomena. For it must be admitted that at the present stage of research we do not have a very well developed understanding of how IBE operates: a point Fodor has been fond of repeating, reiterating the claim that we are very much in the dark about abductive belief-fixation[8]. Moreover, we should be wary about the notion of 'theory' as it is used in both arguments from underdetermination, and in defence of IBE. Without care, such use of the term 'theory' is liable to warp the debate in a way not at all helpful to the realist position, as we shall now try to explain.

IBE is usually employed to support some theory. But how big is a theory? There really is no answer to this question. Explanatory systems and hypotheses of quite different degrees of scope and generality in terms of their domains of application all get to be called theories. Compare, for example, Darwin's theory of evolution through natural selection, with the theory that the dinosaurs became extinct because of a meteoric impact. Clearly, the term 'theory' is a highly elastic one and can be stretched to cover explanatory accounts of all sorts of sizes. By focusing on theories of wider scope, there is a tendency to draw a too rigid distinction between the source of the inference and what it is intended to support. Emphasis on the distinction between 'the context of discovery' and 'the context of justification' also encourages us to think in terms of a complete divide between explanatory theories and data, with an absolute division between the theory to be supported and the data to be explained. This is potentially damaging to the realist position, because in the case of any very significant or fundamental theory it can suggest the picture of a vast data set in need of one single explanatory hypothesis. Such a conception readily generates problems over how we can ever be at all confident that one particular theory is uniquely eligible as 'the best explanation'.

One thing which is badly misleading about this picture is that, in setting the data on one side and some fairly big theory on the other, it makes it seem as if all the data may be equally important and that the data set is *unstructured*: it is just a conjunction of a vast number of otherwise isolated, empirical facts. The data set is just so many observations or readings, whose significance awaits their theoretical explanation. This picture of the unstructured data set is liable to be reinforced by any habit of falling back on unrealistic examples for the purposes of philosophical reasoning. One of the starkest illustrations of this is given by any examples in which the data set is imagined to be merely a sequence of 1s and 0s, with the candidate hypotheses being different equations giving long-term patterns in the sequence.

In real science the data set always has a complex causal structure. We are not taking readings of something generated we know not how. There must be some causal knowledge built into the data, or we would not know what we were doing in making observations or what we were seeking to explain. Let's take an extremely simple example to make this point more vivid: the inference to the conclusion that it has rained in the night because things are wet in the morning. This has frequently been used in philosophical exposition as an accessible and everyday illustration of IBE. But describing the data which that IBE draws on in a casual way ('it's wet out there', 'there are puddles in the street', 'there's water on the ground') can make it seem that what warrants the inference to the conclusion that *it rained during the night* is that other hypotheses are inferior in terms of various theoretical virtues independent of particular evidence: that they are not so simple, so fecund, or are just downright far-fetched (e.g., *a film crew has been trying to make it look as if it was raining*). But before resorting to those considerations, look at the data more carefully. There is water on the street: maybe the city council has a new policy on overnight cleaning? But they wouldn't have hosed my neighbour's lawn. The grass is wet in these gardens. Maybe keen gardeners left their sprinklers on? But they wouldn't do anything that would leave those puddles on the pavements. Even this window is wet on the outside. So maybe late-night window-cleaners have been at work? But no. Anybody who has ever looked at such things would be able to distinguish a pattern of trails and wind-blown droplets from the after-effects of cleaning. More detailed description might well be required. But our point is that the details are there, in any domain, and are all potentially targets for explanation. The 'data set' is a real, richly structured, phenomenon, naturally described using causal terminology.[9]

For this reason we think that, rather than concentrating on IBE and whether it can be defended against van Fraassen's critique, it is a better strategy for the scientific realist to develop an account of a form of causal-explanatory inference which can be deployed wherever there is some potential explanatory target. This sort of 'micro-inference' can thus be used at a sub-theoretical level to introduce causal structure into the body of data which some more grand and general theory seeks to integrate. This strategy may yield a double harvest. If causal-explanatory inference is at work in developing the data to be accounted for by some theory, as well as any larger scale explanatory inference (*IBE proper*, perhaps) from the data in support of that more general theory, we may begin to understand why the philosophical underdetermination thesis appears at odds with the actual theoretical situation in so many scientific domains. So it is possible that by thinking in terms of causal-explanatory inference we may gain a new perspective on what, from the point of view of the realist, seems to be the problem of unlimited underdetermination.

But whether we get that new perspective or not, we can establish a connection between causal-explanatory inference and realist epistemological commitments. In order to see that connection, we only need a reminder of the commonplace that explanation and inference are inversely related. If one can explain E in terms of C, then one can also infer C from E.[10] Moreover, as we shall go on to argue in general and exemplify from cases of scientific research, contrastive causal explanation can be used actively in order to establish what differences have been causally efficacious, and in consequence what differences can be inferred in similar cases. So contrastive explanation helps us find out about the causal factors differentially involved in the production of the phenomena which are taken as our explanatory targets. This seems to us to be a very important form of non-demonstrative inference, and therefore deserves to be acknowledged and made memorable through being named. We will call it 'Differential Inference'[11].

And how does all that help to support the epistemological aspect of scientific realism? Our basic thought is the simple idea that our ontological commitments are founded in the first instance upon those things which we must recognise as existent because we interact with them. But what else must we acknowledge? *Everything which is in causal commerce with the items of our base ontology*. When you suffer from influenza, you know that the physiological changes you are undergoing as part of the illness, and the discomfort they cause you, are real, if anything is real. If it is a virus which is causing you to suffer from that illness, then that virus must be equally real. As Descartes insisted long ago, there must be at least as much reality in the cause as in the effect.[12] Notice also, by way of confirmation of this simple view, that in any case in which such causal involvement is absent or problematic an ontological commitment is liable to seem fantastic and gratuitous. Thus, for example, it is plainly the lack of any possible causal link with items in our base ontology which makes Platonism such an awkward position in the philosophy of mathematics. So if we can describe a form of inference which is undeniably causal (as we take Differential Inference, the inference from effect to cause which is the inverse of contrastive explanation, to be), then we are also describing a form of inference which has the capacity to build a network of processes, objects and forces which are all to be acknowledged as co-existing, because they are in causal interaction with each other.

It is true that this may still leave scope for theoretical disagreement and further investigation concerning the best theories of how those explanatory entities and forces operate (as we have pointed out with regard to Hacking's argument). So it is possible that we may have to revise our current theories in some respects. But we can have good reason to think that we will not need to junk the ontological commitments of our current theories. Whatever the details, there can no longer be any doubt that such things as electrons and viruses exist. Claims of this sort can be buttressed by the closer examination of differential causal inference undertaken in the remainder of this paper. This examination will be shown to support a sufficiently robust scientific realism permitting a claim which might otherwise appear to be a fudge: that our current theories 'are true — or at least approximately true'.

II: CONTRASTS: QUESTIONS AND HYPOTHESES

Scientific realism is best supported by focusing on the *process* of scientific inference: the interrelated activities of questioning, hypothesising, explaining, causally inferring. This should be contrasted with the more usual attempt to read off realism from finished scientific *products* (whole theories and their successful application).

Inquiry begins with questioning. The paradigmatic form of scientific questioning is contrastive: 'why does this happen, rather than that?' Why does the sun appear yellow, rather than some other colour? Why does this neighbourhood, rather than that one, have such a high rate of unemployment? Why is there infection to be found in many of the mothers in ward A, rather than ward B? Why do most animals have two eyes, rather than three or more? Why is there life on Earth, rather than Mars? Why is it that adult humans succeed in locating objects on the basis of landmark information after disorientation, whereas animals such as rats, as well as human children, fail? Why are the noble gases, as opposed to other similar elements, so unreactive? So regarded, the resulting explanation will have three elements: explanans, explanatory target, and contrast or contrasts ('foil' sometimes being employed in place of 'contrast'). Detailed accounts of contrastive explanation were first provided by van Fraassen[13] and Garfinkel[14], each of whom stressed the use of contrast in understanding explanatory relevance.

Contrastive questioning serves to focus inquiry on specific features of the target phenomenon. Such questioning focuses upon those features by which target and contrast differ. What requires explanation is some difference between target and contrast. What explains that difference is some difference between the histories of target and contrast. This valuable insight can be credited to Peter Lipton's Difference Condition for contrastive explanations[15]:

> "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[16] in the history of not-Q."[17]

To take one of the previously cited examples, in seeking to explain why there is life on Earth rather than on Mars it will not help to cite the fact that the Earth was and is a solid planet, for there is a corresponding fact in the causal history of Mars. Rather, we should focus on actual differences in history – hypothesising, perhaps, that the relevant feature is the differing planetary distances from the Sun. Note that the Difference Condition correctly requires that a good contrastive explanation cite a *causal difference* between target and contrast. This does not, however, obviate the use of contrastive questioning in inferring causes, since at this stage of the enquiry all that need be supposed is that the hypothesis meets the necessary, but not sufficient, criterion that the hypothesised explanans is one which refers to a *difference* between target history and contrast history.

Are all explanation-seeking questions contrastive? One can appreciate the importance of contrastive questioning to scientific inquiry without answering in the affirmative. So we do not here need to argue that most, if not all, explanations are contrastive.[18] However, two points should be made which widen the applicable scope of contrastive terminology. First, it is important to recognise that we may properly regard a question and its answer – the resultant explanation – as contrastive, even where no contrast is explicitly stated. Contrasts may be implied by such conventions as verbal stress, or written italics (compare the effect of asking why most *animals* have two eyes, why most animals have *two* eyes, why most animals have two *eyes*). Contrast is also understood (often in conjunction with these sorts of conventions) in virtue of the pragmatic circumstances in which the question is asked. Explanations are, in this sense, dependent upon the interests of the questioner.

Second, it may be thought that in a good many cases, the only possible contrast is simply ‘rather than not’. But if that is so, it may seem that such a contrast does no real work, adding no content to the plain non-contrastive question. This possible response should be resisted, however, for the reason that the negation referred to by adding ‘rather than not’ would need, in any given context, to be understood in a limited, or determinate, way. ‘Why did the bridge collapse?’ might best be regarded as being contrastive only in the sense that one might add ‘rather than not collapsing’. But not every case of the bridge not collapsing is to be regarded as a potential contrast: in a pragmatically typical situation, no consideration should be given to the possibility of the bridge being vaporised, melting, or simply disappearing. The bridge’s not collapsing is potentially realisable in an indefinitely large number of ways. But only a small number of these possibilities will need to be excluded by a satisfactory explanation of why the bridge collapsed.

Contrastive explanation is, therefore, a useful and widely applicable tool. Contrastive questioning directs attention towards differences between target and contrast, and thereby serves to suggest explanatory hypotheses. It is a truism that not any contrastive question will be equally appropriate as a means of generating hypotheses. What, then, are the general features that suitable contrastive questions possess? Ideally, P and not-Q should introduce two separate though similar systems, of which a single differential aspect is selected as that to be explained. In more detail, we can point to three necessary requirements. First, the two systems referred to should be similar; the differences between them should be few (‘why is there life on Earth, rather than in a black hole?’, would be an unsuitable question for this reason). Second, the difference to be explained should raise only a single aspect of the target case. The oddity of the possible question ‘why is there complex life on Earth, rather than single-celled organisms on Mars?’ can be attributed to failure to meet this second requirement. For two aspects are referred to: the potential location of life, and the potential complexity of that life.

Third, in order to be a source of contrastive hypothesis, P and not-Q should be a contrast which introduces separate systems. (This requirement, unlike the first two, is not a demand upon contrastive questioning in general, but only upon the use of contrastive questioning in the process of differential causal inference.) Earth and Mars are indeed separate systems, and hence have separate causal histories. For this reason, actual differences can be found, and hence hypotheses suggested. These actual differences are either directly observed, where the historical difference is also a present difference as is the case with the radius of planetary orbits, or inferred by means of other knowledge in the case of a historical difference no longer present. However, to ask why Earth developed life after one billion rather than one million years is not so amenable, since this question does not refer to two separate systems whose historical differences can be compared.

Even in the latter case, the contrastive question might be useful in focusing attention upon particular features of the target case. Substitute contrasts are intuitively employed: perhaps an actual case which is suitably similar to the contrast (an Earth-like planet on which life developed after one million years), or, more likely, simply an imaginary case. The effect would be to substitute separate systems for a single system with distinct fact and foil outcomes. However, there are important advantages to comparing actual separate systems for differences. Such comparison is highly tolerant of ignorance: one need not know much about the origins of life in order to locate potentially relevant differences. Comparing actual cases also has the advantage of being a source of surprising and previously unimagined differences. To switch examples: when patient X does not contract a disease whereas patient Y does, both seemingly equally exposed to infection, we need only seek some difference between X and Y which might account for X’s immunity. That is something we can investigate before we have anything much at all in the way of knowledge of what the agent of the disease is and how it interacts with human physiology. At that stage of investigation most of those causal factors and processes can be presumed to be *common* to the causal histories of X’s immunity and Y’s illness. We need only search for a difference.

It is worth exploring this requirement of ‘separate systems’ by comparison with a distinction introduced by Lipton, that between compatible and incompatible contrasts. In opposition to previous writers (including van Fraassen and Garfinkel), Lipton maintained that contrasts need not be incompatible[19]. ‘Why Earth rather than Mars?’ refers to compatible facts, while ‘Why after one billion rather than one million years?’ refers to incompatible facts. The former pair could have both been realised (though in fact they weren’t) whereas the latter pair could not. Generation of hypotheses from contrastive questioning is particularly clear in the case of compatible contrasts, as has been illustrated by the differing hypothetical roles of the contrasts Earth/Mars and $10^9/10^6$ years. Yet ‘separate/dependent systems’ does not mark out quite the same distinction as ‘compatible/ incompatible contrasts’.

Consider, for example, the incompatible contrastive question 'why did Peter win the Nobel Peace Prize for 2005, rather than Quentin?' The two 'systems' in this case are the life and work of Peter, and that of Quentin. These two systems can be contrasted, and differences discovered: perhaps Peter had written better selling books, or had a more influential political role. The histories referred to by an incompatible contrastive question are entangled in a way that those referred to by a compatible contrastive question are not. The two advantages that Peter possessed are not only causes of Peter gaining the Prize, but also – in virtue of the fact that there is only one Prize on offer – causes of Quentin not gaining the Prize. Nonetheless, with respect to the Difference Condition, Peter's advantages are potentially relevant differences insofar as there are no *corresponding* features in Quentin's history: corresponding features being Quentin's selling a similar amount of books, or having an equally influential political role. To put the point in another way: although these causes of P are also causes of not-Q, that does not obviate the possibility that these causes answer the question 'Why P rather than Q?'

In any case, the suggestion of hypotheses by contrastive questioning is a *heuristical* feature, and as such depends on simply recognising differences, no matter how the target and contrast turn out to be related. For judgements of compatibility are both fallible and admit of degrees; and neither perfect certainty nor perfect compatibility are required in order for hypotheses to be generated[20].

To summarise: a demand for contrastive explanation is not only a demand that prior causal knowledge be deployed in answer to that question, but is also the first step to the inference to new causal knowledge. There are distinct advantages to this sort of hypothesis generation, which explain why it is so prevalent in both common sense and more developed scientific thinking. Contrasts which exhibit independent causal histories are of particular importance as a source of causal knowledge because they are so tolerant of ignorance. For there may be many causal factors and complex processes involved in the production of both fact and foil. But so long as they are *common* to the causal histories of fact and foil, by a corollary of the Difference Condition, we can take it that those factors and processes cannot by themselves account for the difference in outcome.

III: INFERENCE AND MANIPULATION

So a proposed difference between causal histories of fact and foil is a candidate explanans for a contrastive explanandum. But it may not be the right explanans because there will be other differences, perhaps amongst these differences which are easily confounded with the candidate because they accompany it not only in this instance, but quite regularly. Finding a difference is a necessary condition for successful contrastive explanation, but we need to ascertain whether that difference really does account for the difference in outcome. And how are we to do that?

In general there can be no single recipe or procedure for resolving this issue. Causal-explanatory inference in general, and Differential Inference in particular, is a knowledge-rich and informal business in which we are entitled to draw on anything we know, any connection with general theory, or any further experiment which can be devised. While it isn't true that 'anything goes' in causal-explanatory inference, there is no general way of framing and delimiting what body of knowledge will be relevant, and no *a priori* limitation on the scope for experimental ingenuity. But what we can do here is characterise some fundamental investigatory procedures which help to supply from experience, in terms of both examination and trial, the causal information needed for this sort of inference to be used. We shall argue that contrasts are not only useful, indeed ubiquitous, in the generation of hypotheses, but also are of central importance to the testing and justification of those hypotheses.

Observed or suspected differences between fact and contrast constitute a shortlist (perhaps, initially, a very short list of just one difference, but perhaps more), to be tested in a way to be described. Yet, it is likely that no matter how well chosen, there will be too many differences between the two histories to allow all to be actively tested. In comparing the histories of two neighbourhoods in order to inquire as to the causes of the differing unemployment rates, a multitude of more-or-less relevant differences will be apparent. Historical differences might include not only differing levels of educational attainment, numbers of business start-ups, and levels of business rates; but also differing distances from the nearest glacier, differing kerb widths, and whether more streets run East-West than North-South. Clearly, it would be an understatement to claim that a great many of these differences are highly unlikely to be relevant. They are unlikely to even enter the consciousness of the investigator. We might term this implicit stage of the inquiry 'relevance filtering': a necessary closed-mindedness on the part of the inquirer.

It is instructive to compare 'relevance filtering' to what Mill had to say about causal inquiry, and in particular his Method of Difference (a method which, we are happy to acknowledge, bears much similarity to the description of scientific activity using Differential Inference outlined in this paper). Mill required of the 'most perfect' method of causal investigation that we find two cases which differ in just a single antecedent factor.[21] Yet it has often been remarked that such a requirement can rarely, if ever, be met. Mill himself was forced to recognise the difficulty of meeting the requirement of a unique difference in his belief that the social sciences must be entirely applied sciences, given that if there were one respect in which two communities differed they would inevitably differ in many other respects too[22].

Filtering assumptions can be general (concerning what sort of mechanisms could possibly account for the difference) or specific. And these assumptions will, as suggested, usually not be explicitly articulated. There is a danger that the differences we are prone to notice may be less causally significant than we imagine, and the real causal agents may not be particularly salient to us. This difficulty is clearly related to the 'problem' of underdetermination. For, just as van Fraassen's argument that IBE does not allow the inference of scientific realism because the explanation may be 'the best of a bad lot', it could be that the overlooked differences are the causally efficacious, explanatory, differences. However, the crucial point is that assumptions of non-efficacy are defeasible. Where each member of the short list of potentially relevant differences is shown to be unsuccessful (in a way to be described below), then the investigator is committed to re-examining the previously ignored differences. By appreciating the process of scientific inference as extended, not as a result of comparing 'ideal' (but unattainable) contrasts, Mill's 'problem' can be seen to be problematic neither in theory nor, usually, in experimental practice. The trickiest cases will involve *confounds*, 'silent partners' with no causal efficacy of their own. We will conclude our account of causal inference by further consideration of this important issue.

The use of contrastive reasoning in experimental testing is most noticeable in the paradigmatic case of contrasts having separate causal histories, accessible to direct intervention. The experimental manipulation can in this case be modelled using the notion of 'Difference Closure'. In short, either the target or contrast is manipulated according to the feature hypothesised, with the aim of closing the difference. Either the target feature is suppressed, and hence made similar to the contrast (not-Q), or the feature is produced in virtue of manipulating the contrast such that it is made similar to the target. If the difference is closed, then the hypothesis is proved correct, the cause inferred. More generally, the explanation of the contrastive difference is sufficient to the extent that the difference is experimentally closed. The requirement of difference closure is more demanding than simply that the target be altered in some way, for the target might be altered and yet not become like the contrast. In such a case, we are justified in claiming to have manipulated some relevant causal feature, though not to have answered the contrastive question successfully.

Of course, this sort of experimental intervention is not always possible. We can hardly make the Earth like Mars in some respect to see if we thereby succeed in suppressing life altogether. And while it would be interesting to see if we could stimulate the evolution of life on Mars by making it in some way more similar to Earth (a sort of terra-forming by seeding), any such intervention is well beyond our capabilities. So in order to use Difference Closure we need not only separate causal histories, but separate causal histories which are available for manipulation. An additional, and notable, class of cases which are excluded from this sort of investigation are historical contrasts: the coming of the French Revolution, or of universal suffrage, or for that matter the winning of a Nobel Prize. For in those cases the causal histories, being unique in their realisation, are no longer available for intervention. But we can often apply Difference Closure both to general compatible contrasts and to compatible contrasts in enduring systems in which different outcomes are still being generated.

Indeed, we can also apply Difference Closure to ***in***compatible contrasts, provided the causal systems are still available to experimental intervention. As noted above, we consider the case of compatible contrasts to be a particularly favourable case for the discovery of causal explanations, because distinct causal histories can be surveyed and this survey is a source of hypotheses as to what the causally significant difference is. But Difference Closure is also possible in cases in which a contrast is incompatible because we are only dealing with a single causal system. Compatibility of contrast psychologically facilitates hypothesis-formation. But it is not essential to the testing of a causal-explanatory hypothesis. This is well illustrated by the example of Semmelweis' investigation into child-bed fever[23]. As a matter of fact, there were two maternity wards in the Vienna Hospital, with significantly different rates of mortality from child-bed fever. That obviously helped Semmelweis to formulate hypotheses concerning what difference might account for the higher level of mortality. But even with a single maternity ward one might seek to explain why mortality from that infection was so high, and seek to reduce it by variation in treatment.

Where the difference is not closed by manipulating the hypothesised factor, the natural next step is to return to another candidate on the shortlist. Where there are no such remaining candidates, the assumptions made at the filtering stage must be re-examined. Where the hypothesis is shown to be successful, it is to be expected that new questions are raised, and thus that the process of inquiry loops back to the start (though, of course, with differing erotetic content). When Semmelweis justified his hypothesis that cleanliness was the answer to the difference between mortality rates between the two wards of the Viennese maternity hospital, more precise questions became possible: what specific cleaning procedures would be most effective? What, precisely, was it about previous procedures that had caused the infections?

Difference Closure is the basis of some standard experimental paradigms. One widely-used example of this is the so-called "Common Garden" experiment. The rationale of a common-garden experiment may be formulated in this way, in botanical terms. Suppose a kind of plant is found in two different environments, exhibiting variation in some trait between those two environments (perhaps tall-growing plants being found in the valley-bottom, while shorter specimens are found at higher altitudes). We might want to resolve the question whether this variation is due to environmental factors, or whether it is due to some intrinsic difference in the plant populations found in these environments. By moving plants (or, rather seeds or cuttings taken from these plants) and growing them in a common garden we can eliminate the differences in environmental factors. If variation in the plants grown in the common garden persists at the same level as in the wild, then it cannot have been due to environmental factors. Although we have characterised the experimental rationale in botanical terms, this is a paradigm which can be applied equally well to animal populations. It has even been employed in a sociological context in which human immigrant populations with different cultural backgrounds have moved to the same host country.[24]

We have emphasised the point that differential inference based on contrastive explanation is very tolerant of ignorance about how most of a causal system is operating. It seems appropriate, therefore, to consider an example taken from Cognitive Science, given that we have to admit there is little we know about how cognitive processing operates in the human mind. In this domain a good example of an experimental procedure using difference closure is the dual-task paradigm in cognitive psychology, employed by Spelke and her associates to investigate the hypothesis that natural language is used to integrate information processed by distinct modular systems.[25] In this example the contrastive explanandum is: *Why can human adults succeed in locating objects on the basis of topographical and landmark information after disorientation, while children and rats are not able to do so?* The application of Difference Closure was heavily influenced in this case both by a general commitment to the modularity of the mind and a quite specific hypothesis that natural language is a medium in which information taken from different modules can be combined. Rats have no natural language, and children at an age at which their linguistic abilities for describing location are limited perform no better than rats on the task of locating hidden objects distinguished both by topographical features and landmarks (e.g., being near a brightly coloured wall, in the case of children; or in a corner in which there is a pungent smell, in the case of rats).

The hypothesis to be tested is that adults succeed in locating objects because they are able to integrate topographical and landmark information in a natural language format. So one way of using the Difference Condition would be suppression: making the natural language format no longer available. Hermer-Vasquez, Spelke, and Katsnelson achieved this by getting their subjects to fast-shadow (i.e., repeat with a very short delay) something played to them on a tape. This demanding task would fully occupy their natural-language processing, thus making them, in the relevant respect, cognitively more similar to rats and young children. Fast-shadowing did indeed reduce the 'fact-adults' to the same foil-performance on the location task as rats and young children.[26]

The interpretation of performance on the dual-task paradigm raises in a concrete fashion a general concern about the probative force of difference closure. Namely, how can we deal with confounds? It might seem that by manipulating some factor we can switch an outcome on or off, thereby confirming the hypothesis that this is the factor which causally accounts for the difference in outcomes. Yet perhaps it is not the factor we have hypothesised which is causally responsible, but something else, which is also being changed by our experimental intervention. In other words, there could be a confound, and we need to eliminate that possibility. In the Hermer-Vasquez, Spelke, and Katsnelson trial one obvious possible confound was processing load: subjects who were given a natural-language processing task were thereby also being given *more* cognitive processing to do. So perhaps it was the added processing load, rather than the specific engagement of the natural language module, which closed the difference between fact and foil. Perhaps, but such a possibility can be eliminated by sufficient ingenuity in experimental control. Hermer-Vasquez et. al. did this by getting another group of subjects to clap in time with a musical soundtrack, rather than shadowing speech. The clapping group also faced an increased demand on cognitive processing, but their ability to locate objects was not impaired in the same way as those who were shadowing speech.

On the model we have outlined it can be seen why the elimination of confounds is a major problem facing the scientific investigator — and a problem which remains after the achievement of difference closure. In so far as difference closure has succeeded, it will give us the impression of having discovered a causal 'toggle' which can be manipulated in order to control the outcome. But this leaves the residual worry that the difference we had in mind might not be the only difference which was affected. There might be some fellow-traveller which, unbeknown to us, was the factor which was really causally relevant. That is always a logical possibility. Should we therefore recognise that by the most exacting standards causal inference founded on contrastive explanation and difference closure can never be absolutely conclusive, and so the philosophical dogma *that scientific hypotheses can never be proved* is justified? In our view this is a perversely sceptical standpoint. As a general principle in the philosophy of science it seems to us no more sensible than a general principle in jurisprudence to the effect that a trial can never prove anyone guilty of a crime, on the grounds that whatever the evidence presented it always remains a logical possibility that the accused is innocent.

A possible confound is eliminated by combining Difference Closure with Confound Differentiation. The hypothesis is that factor Y is causally responsible for a difference in outcome between fact and foil. So we remove factor Y to suppress the fact-outcome or introduce factor Y to generate the fact in the foil-situation (Difference Closure). But perhaps it is factor Y's 'shadowy companion', factor Z, which is actually doing the causal work. In order to eliminate this possibility experimental intervention needs either to remove factor Z without removing factor Y from the fact-situation, or to introduce factor Z without introducing factor Y in the foil-situation (Confound Differentiation): i.e., with the intended outcome (if Z is a confound) of keeping the fact 'on', and the foil 'off'. This procedure can sometimes be iterated to eliminate a whole series of possible confounds.

Once any plausible confounds have been eliminated the onus of proof should be shifted to anyone who wishes to insist that the candidate factor is not playing the hypothesised causal role. Of course it does always remain a logical possibility that there is something else which we have not eliminated, perhaps because we have not detected or even thought of it, which accounts for the difference in outcome between fact and foil. Such logical possibilities always remain. But it can still be the case that, by using the quite familiar methods for which we have been seeking to give a general rationale, *a causal hypothesis can be established beyond reasonable doubt*. After any plausible candidate confounds have been dealt with by iteration of the eliminative procedure of Confound Differentiation, there comes a point at which any remaining doubt would have the character of a sceptical hypothesis. Here is a way to remind yourself of the far-fetched, sceptical character of the possibility of error which has been taken so seriously in the philosophy of science, and which is the constant source of anti-realisms. Go to the bathroom and turn the tap on and off. Clearly the flow of water into the sink is controlled by turning the tap and thereby opening and closing a valve. Even in this case there is a logical possibility that something else is causally responsible for filling your sink and your bath, that every time you turn the tap and open the valve, something else is going on and it is that something else which really makes the difference between water flowing out and not flowing. That is a logical possibility, but it's a logical possibility on a par with the Evil Demon and brains-in-vats. The experimenter should not be asked to eliminate those.

IV: CONCLUSIONS

The model of contrastive inquiry can be summarised with the aid of the following flow chart:

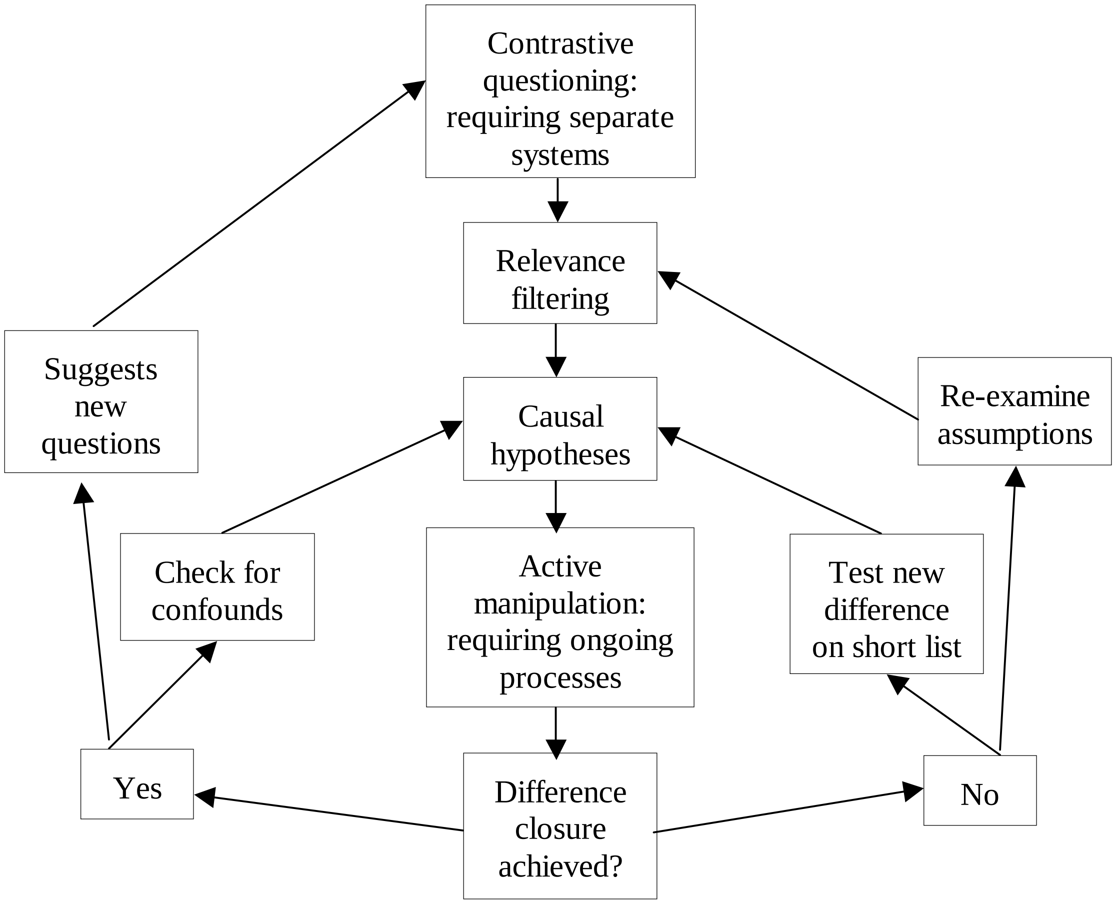

We will conclude the paper by noting the underlying causal principle upon which our model of causal inference is based; and then making explicit the consequences of the role of differential inference for the debate about scientific realism.

Our account of causal inference, as well as Mill's structurally similar Method of Difference, is founded on a fundamental causal principle. It is the one sometimes casually expressed as 'Like causes, like effects'. For well known reasons, that principle is only acceptable if 'like causes' include the totality of causally relevant antecedent conditions. We could therefore state this principle more perspicuously as the Principle of Nomological Causal Determinism (PNCD):

> (PNCD) For every event, there exists a set of antecedent causal conditions such that if those conditions are repeated an event of exactly the same kind will occur.

Hume was quite definitely committed to this principle, which he called the 'Doctrine of Necessity'. The impression has arisen that, if we are to follow Hume's insights, we should regard the basic form of inductive inference as induction by repetition, or enumerative induction (i.e., the projection of some experienced regularity into the future). Yet the present treatment of causal explanatory inference could just as well be regarded as neo-Humean. Indeed, the Difference Principle is anticipated in the sixth of his 'rules by which to judge of causes and effects':

> 'The difference in the effects of two resembling objects must proceed from that particular, in which they differ. For as like causes always produce like effects, when in any instance we find our expectation to be disappointed, we must conclude that this irregularity proceeds from some difference in the causes.'[27]

In our view it is singularly unfortunate that induction by repetition has come to be taken as the standard form of causal inference in the philosophical tradition deriving from Hume, whereas differential inference has been comparatively neglected. Suppose that we lived in a world in which the laws of nature were directly accessible to observational experience, and things which were taken by us to be of the same kind were homogeneous in all causally significant ways. Call this the Manifest World. In such a world induction by repetition would serve us well, since observed, relatively localised, regularities really would be lawlike. Suppose instead we lived in a world in which the laws of nature were deeply hidden from ordinary human experience, and things which were taken by us to be of the same kind were extremely heterogeneous in causally significant ways. Call this the Secret World. In the Secret World laws of nature would not be revealed to us in regularities we experienced. So we would repeatedly have to rely upon difference closure, inferring some causally significant difference from a difference in outcome. The world we actually live in is not the Secret World. But it isn't the Manifest World either. It's somewhere between the two, a quasi-secret world, and so contrastive inference is of at least as much use to us as straightforward inductive projection.

If our model of Differential Inference is broadly correct, then it does constitute at least one way in which we can reason actively from the phenomena. If it is correct, then folk and scientists alike knew all along, albeit tacitly, how to employ this form of inference. We have tried to make its structure and presuppositions explicit. It must be allowed that those presuppositions are themselves ultimately defeasible. Once differential inference has been used to build a body of causal knowledge vast enough to support modern science it is quite possible that the Principle of Nomological Causal Determinism itself should be brought into question.

How is the cause of scientific realism aided by the model presented in this paper? Primarily in three ways. First, this model bolsters a Hacking-inspired emphasis on the connection between experiment and scientific realism. The fundamental idea is that we must accord real existence to those things with which we interact. The tendency of such argument to be overly permissive (allowing the inference to phlogiston, for example) is counteracted by emphasising the ongoing and defeasible nature of the inferential process. Second, the thesis of underdetermination is undermined by developing an account of the relation between data and theory, which emphasises the extent to which causal notions come to penetrate interpretations of the data. Third, underdetermination is further challenged by noting that ongoing experiment can eliminate confounds. Radical scepticism notwithstanding, we can indeed know that the causal factors inferred are those that really do make the difference.

[1] Hacking reports (1983, pp.22-23) a conversation with a scientist who explains that the way to alter the charge on niobium balls is either to spray them with positrons (to increase the charge) or with electrons (to decrease the charge). In the concluding section of his book he remarks: "The 'direct' proof of electrons and the like is our ability to manipulate them using well-understood low-level causal properties. … The clincher is when we can put a spin on the electrons, polarize them and get them thereby to scatter in slightly different proportions." By contrast he expresses a degree of scepticism about black holes, entities we have difficulty in interacting with, even if there is theoretical justification for postulating them. Hacking (1983), pp.274-275.

[2] 'The positive argument for realism is that it is the only philosophy that doesn't make the success of science a miracle.' Putnam (1975), p.73.

[3] See Laudan (1981)

[4] This argument is presented by van Fraassen in chapter 2 of van Fraassen (1980). He subsequently reinforces the case against realism by objections to the realist's use of Inference to the Best Explanation in chapter 6 of his (1989).

[5] According to Thomas Kuhn, of course, the regular business of 'normal science' is attempting to deal with *anomalies*, some of which are apt to prove highly resistant to all efforts to smooth them out (Kuhn (1962), esp. Chs. 2-6). Philosophers of science may consider empirical adequacy a more modest aim than truth. But it will not seem that way to a normal scientist who is confident her paradigm theory is true, but who must recognise that her whole career may fail to secure its empirical adequacy in relation to just some of the problems within its domain.

[6] Feyerabend emphasised the claim that no theory is, or can reasonably be expected to be, in agreement with all the facts in its domain. Cf. Feyerabend (1981), p.106; Feyerabend (1975), pp.55-68. Clearly Feyerabend would have regarded empirical adequacy as an aim that would put an end to genuine scientific progress. But his position on this does, of course, depend upon claims about incommensurability which few are prepared to buy.

[7] van Fraassen (1989), pp.142-143.

[8] Cf. Fodor (1983), p.107 and pp.116-119; Fodor (2000), Ch. 3.

[9] So the data set includes observations which are indeed 'theory-laden'. But the theories involved are not the same as those which the data is used to support. So this kind of 'theory-ladenness' of observation, so far from being a problem for the empirical check upon theorising, actually strengthens that check by introducing causal structure which an acceptable theory will need to explain, or at least be consistent with.

[10] Of course, the explanation needs to be correct: the security of the inference is proportional to the security of the explanation. Ptolemy was able to 'explain' planetary retrogression by movement around epicycles. But since the explanation was wrong, we should not infer epicyclical motion. We hope our account of Differential Inference shows how one can sometimes determine whether an explanation is correct or not.

[11] We could also have called this form of inference ‘Contrastive Inference’, which is the label Lipton adopts for what appears to be this kind of reasoning in Chapter 5 of his seminal (1991/2004). However, Lipton appears to view contrastive inference as a species of IBE, even calling it ‘Inference to the Best Contrastive Explanation’ (IBCE) at one point (p.78/p.73). In this paper we make the case for distinguishing between IBE as inference to a theory which accounts for a diverse body of phenomena, and particular abductive inferences from a single phenomenon.

[12] Though we would not endorse the use he makes of this principle in *Meditation III*.

[13] van Fraassen (1980), Chapter 5.

[14] Garfinkel (1981), Introduction and Chapter 1.

[15] Lipton (1991/2004), p.42. The formulation in the book differs from that in the 1990 paper in using the phrase “in the case of not-Q” rather than the phrase “in the history of not-Q”. We take it that the change is intended to accommodate factors which are still present.

[16] The key notion of ‘correspondence’ in the Difference Condition has received some attention in the literature. However, despite alternative suggestions, it seems that the best understanding is the most obvious: corresponding features in the histories of P and not-Q are those which are of the same type (where that type must be of a kind which could be taken as potentially causally relevant to the differing outcome).

[17] Lipton (1990), p.256. A clarificatory note on a potentially confusing feature of the Difference Condition: the locution ‘not-Q’ is used for good reason, given that the contrast is, by definition, a state which has not been realised. In terms of the present example, ‘Q’ refers to life being on Mars; thus the actual state of affairs of a lifeless Mars, the state of affairs whose causal history is examined in explanatory investigation, is ‘not-Q’.

[18] In his original 1990 paper and in both first and second editions of *Inference to the Best Explanation* Lipton tries to demonstrate that contrastive explanation is not reducible to “straight” or non-contrastive explanation of some fact: e.g., that the answer to ‘Why P rather than Q?’ is not just the same as the conjunction of the answers to ‘Why P?’ and ‘Why not-Q?. The argument for this (that sometimes it is easier to supply contrastive explanations than corresponding non-contrastive explanations, and *vice versa*) might appear to involve a commitment to the existence of non-contrastive explanations, though Lipton actually describes himself as “agnostic about the existence of non-contrastive why-questions” (Lipton (1991/2004), p.50). We are inclined to be sceptical about the existence of non-contrastive why-questions, but will be pursuing that issue in a future paper.

[19] Lipton (1991/2004), p.36/p.34

[20] To illustrate fallibility, consider the possibility that the process which brought Earth into a solar orbit conducive to life may have displaced Mars’ position with regard to such an orbit. To illustrate the continuum of compatibility, consider that it is unlikely, though sometimes happens, that the Prize be shared in a given year.

[21] Mill (1843/1973), p.391. Mill interestingly goes on to remark that "It is scarcely necessary to give examples of a logical process to which we owe almost all the inductive conclusions we draw in daily life. When a man is shot through the heart, it is by this method we know that it was the gunshot which killed him: for he was in the fullness of life immediately before, all circumstances being the same, except the wound."

[22] Mill (1843/1974), pp.881-882. Mill there states that in order to apply the Method of Difference "we require to find two instances, which tally in every particular except the one which is the subject of inquiry."

[23] A case discussed by both Hempel (1966, pp.3-8) and Lipton (1991/2004, ch.5).

[24] Salamon (1992).

[25] Hermer-Vasquez et al (1999)

[26] In a follow-up study Shusterman and Spelke (2005) have attempted to close the difference in the other direction, giving child subjects specific coaching in locational aspects of language to see if their performance can thereby be made closer to that of adults. This raises the interesting question: what, if anything, is gained by closing the difference in both ways? One thereby gains the added assurance of experimental replication. But there is more to it than that, for closure of the difference in both directions involves asymmetries which restrict the range of possible confounds.

[27] Hume (1739/1989), I.III.xv.

References

Feyerabend, P.K. (1975) *Against Method* (London: New Left Books)

Feyerabend, P.K. (1981) *Problems of Empiricism: Philosophical Papers Volume 2* (Cambridge University Press)

Fodor, J.A. (1983) *The Modularity of Mind* (Cambridge, Massachusetts and London: The MIT Press).

Fodor, J.A. (2000) *The Mind Doesn't Work That Way* (Cambridge, Massachusetts and London: The MIT Press).

Garfinkel, A. (1981) *Forms of Explanation* (London and New Haven: Yale University Press)

Hacking, I. (1983) *Representing and Intervening* (Cambridge University Press).

Hempel, C.G. (1966) *Philosophy of Natural Science* (Englewood Cliffs, N.J.: Prentice-Hall)

Hermer-Vasquez, L., Spelke, E., and Katsnelson, A. Sources of flexibility in human cognition: dual task studies of space and language. *Cognitive Psychology* 39, 1999, pp.3-36.

Hume, D. (1739/1989) *A Treatise of Human Nature*, ed. P.H. Nidditch (Oxford: Clarendon Press).

Kuhn, T.S. (1962) *The Structure of Scientific Revolutions* (Chicago and London: University of Chicago Press).

Laudan, L. (1981) "A confutation of convergent realism", *Philosophy of Science*, 48, pp.19-48.

Lipton, P. (1990) "Contrastive explanation", in: D. Knowles (Ed), *Explanation and Its Limits*, pp.247-266 (Cambridge University Press).

Lipton, P. (2004, 2nd edn.) *Inference to the Best Explanation* (London: Routledge)

Mill, J.S. (1843/1973) *A System of Logic Ratiocinative and Inductive Books I-III*, in Robson, J.M. (Ed) *Collected Works of John Stuart Mill Volume VII* (Toronto and Buffalo: University of Toronto Press; and London: Routledge & Kegan Paul).

Mill, J.S. (1843/1974) *A System of Logic Ratiocinative and Inductive Books IV-VI*, in Robson, J.M. (Ed) *Collected Works of John Stuart Mill Volume VIII* (Toronto and Buffalo: University of Toronto Press; and London: Routledge & Kegan Paul).

Putnam, H. (1975) *Mathematics, Matter and Method: Philosophical Papers Volume I* (Cambridge University Press).

Salamon, S. (1992) *Prairie Patrimony: Family, Farming, and Community in the Midwest,* University of North Carolina Press.

Shusterman, A. & Spelke, E. (2005) "Language and the development of spatial reasoning", in: Carruthers, P., Laurence, S. & Stich, S. (Eds.), *The Innate Mind* (Oxford University Press).

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

van Fraassen, B.C. (1989) *Laws and Symmetry* (Oxford: Clarendon Press).