Underdetermination, methodological practices, and the case of John Snow*†

D. Tulodziecki

Version 03/2009

Contents

*Draft of paper to be presented at the *Underdetermination in Science Workshop* (Pittsburgh, 21–22 March 2009). This paper is a shortened version of a longer draft and currently still undergoing revisions; please do not cite without author's permission.

†Many thanks to Philip Kitcher for discussions of an earlier version of this draft.

1 Introduction

One realist solution to underdetermination is an appeal to the idea that there are criteria besides the empirical evidence that are supposed to have epistemic import and break ties in underdetermination scenarios (most often those going under the name of 'theoretical virtues'). Despite widespread appeal to such criteria, however, there has been little discussion of how to generate a robust set of such criteria. In this paper, I want to make some headway towards this goal.

After a brief discussion of underdetermination, the theoretical virtues, and the approach that I will take in this paper (section 2), I will examine a case in the history of medicine – that of the British physician John Snow and his reasoning about the origin and pathology of cholera – and argue that Snow used a variety of inferential and methodological practices that led him to accept various hypotheses about cholera that could not be checked at the time, and that none of his contemporaries accepted (section 3).

I will argue (in section 4) that this case-study suggests: (i) an expanded conception of epistemic criteria besides the empirical evidence; in particular, a conception that goes beyond the theoretical virtues so as to include certain inferential and methodological practices, (ii) that many of these practices are, in fact, epistemically significant, and (iii) that we can test for the success of these practices empirically by examining case-studies in the history of science.

Specifically, I will offer my reconstruction of Snow's reasoning about cholera as a concrete example of the sort of empirical research that needs to be done in order to discover what kinds of methodological rules and practices are actually of epistemic interest – it is thus the first piece of data in support of the approach outlined in this paper.

After discussing some (anti-realist) objections to this approach (section 5), I will explain how the case of Snow (and other cases like it) can help us resolve specific cases of underdetermination (section 6). I will show that this approach issues a new challenge to anti-realists, and argue that, even if anti-realists can successfully diffuse the new objections I pose, they will at most be able to do so in a piecemeal fashion, thus losing the much needed guarantee that there will always be rival cases of the required kind.

2 Underdetermination and the virtues

One of the standard ways of formulating the underdetermination argument is in terms of the following two theses:

1. Empirical Equivalence Thesis (following Kukla 1998, Psillos 1999): any theory has empirically equivalent and logically incompatible rivals, and

2. Entailment Thesis (cf. Psillos 1999): entailment of the evidence is the only epistemic constraint on theory-choice.

From these two theses it is inferred that belief in any one theory over its rivals "must be arbitrary and unfounded" (Kukla 1998: 58). In this paper, I'll mainly be concerned with the second premise. Denying this premise amounts to embracing the view that there can be criteria other than the empirical evidence that can have an effect on a theory's epistemic standing. Most often this response is given in terms of the so-called theoretical virtues (simplicity, elegance, unifying power, consilience, etc.) – criteria that are supposed to raise a theory's chances of being (approximately) true or successful. Here, for example, is a typical realist take on the theoretical virtues from Churchland (1985: 35):

> I will maintain that observational excellence or 'empirical adequacy' is only one epistemic virtue among others of equal or comparable importance [the virtues that Churchland cites later in the same paragraph are simplicity, coherence, and explanatory power].

Here's another from Psillos (1999: 171):

> We also need to take into account several theoretical virtues such as coherence with other established theories, consilience, completeness, unifying power, lack of ad hoc features and capacity to generate novel predictions.

Anti-realists, by contrast, think there is nothing epistemic about the virtues at all. van Fraassen, for example, writes:

> When a theory is advocated, it is praised for many features other than empirical adequacy and strength: it is said to be mathematically elegant, simple, of great scope, complete in certain respects: also of wonderful use in unifying our account of hitherto disparate phenomena, and most of all, explanatory. (1980: 87)

He goes on:

> Values of this sort, however, provide reasons for using a theory, or contemplating it, whether or not we think it true, and cannot rationally guide our epistemic attitudes and decisions. (87)

Thus, anti-realists hold that the virtues are only pragmatically interesting. The trick for the realist, then, is to establish a connection between the virtues and truth. Curiously, despite the fact that the virtues are often invoked by realists, there have been comparatively few attempts by realists to spell out this alleged connection in any detail.[1] In particular, there haven't been any real attempts to discuss and explain either (i) how exactly one might go about establishing this connection, or (ii) how to generate a robust set of virtues (that can that actually be used by realists to dispel alleged underdetermination scenarios).

I think we can make some headway towards both, and we can do so by examining the case of John Snow's reasoning about the origin and pathology of cholera from a certain perspective. Snow, by using a number of inferential and methodological practices, was led to accept a variety of (correct) conclusions about cholera that could not be verified at the time, and that his opponents rejected. I think his case suggests that we ought to expand our conception of epistemic criteria besides the empirical evidence in a way that goes beyond the theoretical virtues; in particular, in a way that includes certain methodological principles and rules, such as the ones used by Snow. Snow is explicit about the fact that he takes these principles to be of the epistemic kind. Of course he could be mistaken (although I hope section 3 will show that he wasn't). Nevertheless, however, this case allows us to catch a glimpse of these principles 'in action', if you will, which is more than can be said for some of the traditional theoretical virtues. What I hope the case shows is that it will at the very least be both plausible and fruitful to include them as potential candidates. In fact, I think the Snow case can do much more than this – it doesn't just suggest that we should extend our pool of criteria, I think it actually suggests that some of his principles *are* epistemically significant. After discussing the details of the case, I'll show how we can use it to make this argument more persuasively.

[1] Although there are some notable discussions of some of the individual virtues; see, for example, Glymour (1985) on explanatory power, and also Kelly's discussions of simplicity (such as his recent (2007a) and (2007b)).

A couple of provisos are in order here: First of all, I ought to stress that the argument that follows in no way depends on the fact that Snow *actually* used these principles. I happen to think that he was an excellent and careful reasoner and that it's no accident that he used the kinds of principles he did[2] – but even if it had been, that would not undermine the argument of this paper. What matters is that following certain principles in certain contexts leads to certain conclusions; this can be tested (if need be, by thought experiments) regardless of whether anyone actually happened to employ them at a particular time. Secondly, one case, of course, isn't enough to establish the epistemic significance of anything. What exactly is required, and some of the objections one might voice to this approach as a whole, will be explored below.

3 The reasoning of John Snow

Neither Snow's, nor his opponents' claims could be verified prior to the discovery of *vibrio cholerae*; there was simply no way of settling for good the question of what exactly was causing cholera and how, or indeed if, it was passed on. Snow actually inferred several different things about the disease; however, for brevity, I will focus solely on his claim that cholera is communicable. The rival theory to this claim was the miasma theory of disease, the prevalent theory of cholera at the time. According to this theory, organic matter decays, decomposes, and gives rise to new compounds that are given off into the surrounding air and atmosphere – the miasmas. Inhaling these miasmas was thought to be poisonous (decaying animal matter was thought to be particularly toxic), and it was thought that the toxic ingredients acted on the blood and disturbed the body's balance. Depending on the specific local conditions, such as location, season, humidity, barometric pressure, and weather, people then contracted one of a variety of diseases (cholera, typhoid fever, etc.). Who was afflicted, and the severity of the affliction would in turn depend on a variety of personal dispositions, such as poor nutrition, sickly constitution, bad hygiene, or, to cite some of the more interesting ones: strong emotions, especially fear, and immorality, specifically overindulgence in alcohol and sex. In fact, this was why it was thought that people were more likely to contract diseases on the weekends, because there were more

[2]For what it's worth, he didn't just excel at reasoning about cholera; he was also a pioneer in developing anaesthesia.

opportunities for dissipation and debauchery. The main reason miasmatists rejected Snow's communicability hypothesis was that cholera was unlike the classic contagious diseases, such as syphilis or smallpox. So why, with this elaborate theory in place, did Snow come to the conclusion that cholera was communicable?

3.1 Evidence I: General Considerations

One of the first things that Snow points out is that cholera always follows "the great channels of human intercourse" (PMCC: 746)[3], and that when troops were attacked with cholera, it would often remain with them through different countries and climates, and that it would often appear in villages and towns that had been free from cholera until troops passed through. In India, for example, cholera would spread along the main travel routes, but pass over villages that were a little away from these routes. This was problematic for the miasma theory: if cholera weren't communicable, and caused only by miasmas, and the presence of miasmas, in turn, depended on factors like geography and climate, then (i) cholera should disappear as soon as these factors disappear, which was not the case, (ii) it couldn't explain the correlation between cholera and the troops, which were supposed to have nothing to do at all with cholera, and (iii) one would expect it to be equally present in areas that are sufficiently similar in climate and various other relevant conditions, but there was no difference among the healthy and afflicted villages in geography or weather – the only difference being the passing through of troops.

This is an interesting piece of data, because it is one of the first instances in which we see Snow using causal principles; Snow's reasoning actually seems

[3] The references to Snow's works are as follows: 'MCC1' refers to *On the Mode of Communication of Cholera*, first edition, London: John Churchill, August 1849; 'MCC2' refers to *On the Mode of Communication of Cholera* (1855), second and much enlarged edition; 'PMCC' refers to *On the Pathology and Mode of Communication of Cholera*, London Medical Gazette, vol. 44, Nov. 2, 1849: 745–752/Nov. 30, 1849: 923–929; 'MPC' refers to *On the Mode of Propagation of Cholera*, Medical Times 3, Nov. 29, 1851: 559–562/Dec. 6, 1851: 610–612. Since all these texts are out of print, for the readers' convenience, I have provided quotations from the online versions at *http://www.ph.ucla.edu/epi/snow/principalwritings.html*. Unfortunately, the text for MCC2 omits Snow's original page numbers; hence, the references I will provide will refer to the divisions on this site (for example, 'MCC2: 3' indicates that a quotation was taken from the third part of MCC2).

to be based on something close to Mill's Method of Difference:

> If an instance in which the phenomenon under investigation occurs and an instance in which it does not occur have every circumstance in common save one, that one occurring only in the former, the circumstance in which alone the two instances differ is the effect, or the cause, or an indispensable part of the cause, of the phenomenon. (1843/1950: 215-216)[4]

This is a principle that Snow uses in a variety of contexts (see below); here, however, he uses it to support the communicability hypothesis: the only difference between the villages in which there was cholera and those in which there wasn't, was whether troops had passed through.

3.2 Evidence II: Snow's Many Case-Studies

Besides general considerations of the above sort, Snow also appealed to a number of cases in which people contracted cholera after having been in contact with either someone sick or something that could be connected to someone sick. Snow is very explicit about the fact that he considers these cases to be of immense importance: "the most conclusive part of the evidence, is the number of instances in which the malady has been introduced into healthy localities by persons who have been taken ill after their arrival from places where cholera prevailed" (MCP: 559a). Once again, the number of Snow's collected cases is overwhelming; however, here is an example that is fairly typical: there was a completely healthy village about which it was known that there was no cholera in an area covering a radius of 30 miles around the village. Yet suddenly one of the villages labourers, 39-year-old John Barnes, became violently ill after having suffered from cramps and diarrhoea. He was diagnosed with cholera by a local doctor, but no one knew how he might have contracted it. Barnes died the next day, and by then there were a couple of more victims, Barnes's wife, and two people who had been to visit him the day before. And while the doctors were pondering how the victims could have contracted cholera "the mystery was all at once, and most unexpectedly, unravelled by the arrival in the village of the son of the deceased John Barnes" (MCP: 559a). It turned out that the son was an apprentice to his uncle in Leeds and that his aunt, John Barnes's sister,

[4] Of course, there is never going to be only a *single* difference. For a discussion of this point, see Lipton (2004, chapter 8).

had died of cholera about a fortnight before. Since they had no children, her apparel had been sent to John Barnes without having been cleaned before – Barnes opened the box, fell sick, and then died.

As before, the miasma theory could not explain the two facts that (a) only John Barnes and no one else contracted cholera in the first place, and (b) that only those in contact with John Barnes went on to show symptoms of the disease. If miasmas had been the cause of cholera, other people in the village should have contracted cholera, not just John Barnes. The only other option was for John Barnes and the others to have left the village and been to a miasmatic region, but they were known not to have done this. And, of course, once again, the incident makes much sense on Snow's communicability hypothesis, especially after the connection to Barnes's dead sister's clothes is established. Again, in isolation these cases would not have been meaningful, but Snow points out that there are far too many cases like this to "be set down as mere coincidence" (MCP: 559a). And, once again, we find him appealing to the causal principle that similar effects have similar causes; only this time Snow spells out his reasoning for us explicitly: "And if cholera be communicated in some instances, is there not the strongest probability that it is so in the others – in short, that similar effects depend on similar causes?" (MCC1: 30).

3.3 Water and the Broad Street pump

I also briefly want to discuss the Broad Street pump episode for which Snow is now famous. Although the heroic story of Snow identifying the pump as the culprit, it being closed, and the epidemic subsiding immediately is far from the truth, there is a question of how Snow came to point his finger at the pump. It should be noted that, by this time, Snow already suspected that a water source was responsible for the outbreak.[5] In the case of the Golden Square outbreak, Snow was able to exclude the water companies; this left him with the pumps. As a first step, Snow collected water from various pumps and took it to a microscopist. However, besides general pollution, the analysis didn't show much. Even this was fairly important, however, because it meant that the pump water was generally dirty and polluted, thus also making it more likely for the cholera agent to be in the pump water,

[5] Most of the reasons for this involved complicated evidence comparing the water of different London water suppliers.

since Snow's hypothesis was that it entered the pump water through people's evacuations. If the pump water had been found to be completely clean, it would have been unclear how the agent could have contaminated the water in the first place. Since examining the water turned out to be insufficient, Snow turned to collecting data.

Looking at all the deaths from cholera, he realised that most of them were close to the Broad Street pump. However, there were several 'anomalies', areas where the mortality from cholera should have been high, but wasn't. This data seemed to disconfirm Snow's theory, and there were two cases of particular interest: the brewery and the workhouse, in both of which there were virtually no deaths. Both of these were close to the pump, and, in the case of the workhouse all of the exacerbating circumstances (according to Snow) such as bad hygiene, lack of light, and overcrowding, were present, so, according to Snow's theory, the mortality rate there should've been high. It's clear how this was problematic for the miasma theory: the miasmas presumably didn't distinguish between those in the workhouse and the surrounding streets, and both the brewery workers and the inhabitants of the workhouse were thought to have low morals. But what about Snow? After making enquiries, he found out that the inmates of the workhouse didn't get their water from the Broad Street pump. What about the brewery? The Board of Health had in fact publicly advised people that alcohol was connected to cholera, so again there was a problem for the miasma theory, since the brewery workers should've been dying in large numbers. Once again, Snow made enquiries. He called on the proprietor and was told that the brewery obtained its water from the New River Company, and, in addition to this, also had a deep well inside. However, more importantly, the proprietor also pointed out that the workers were "allowed a certain quantity of malt liquor, and Mr. Huggins believes they do not drink water at all" (MCC2).

Once again, we can see causal reasoning at work in Snow. Where earlier he used the principle that similar effects have similar causes, he now appeals to the principle that similar causes result in similar effects. This principle makes him rule out miasmas in the Golden Square outbreak, since, if miasmas had been the cause, the effects should've been similar – in our case this would have entailed a fairly uniform distribution of cholera incidents all over the affected neighbourhood, which was found not to be the case. Again, Snow spells out his thought process explicitly:

> I suspected some contamination of the water of the much-frequented

> street-pump in Broad Street, near the end of Cambridge Street; but on examining the water, on the evening of the 3rd September, I found so little impurity in it of an organic nature, that I hesitated to come to a conclusion. Further inquiry, however, showed me that *there was no other circumstance or agent common to the circumscribed locality in which this sudden increase of cholera occurred and not extending beyond it, except the water of the above mentioned pump.* (MCC2, my emphasis)

Snow discounted miasmas precisely because they should have extended beyond the area in question; however, they did not single out the affected neighbourhood, as a true cause should have done.[6] We can also see him appealing to Mill's Method of Difference again (cf. p. 7 above). Snow points out that the only difference between those people who fell sick and those who didn't, is that the former group drank the pump water, with this being corroborated especially by evidence arising from unexpected cases:

> The result of the inquiry then was, that there had been no particular outbreak or increase of cholera, in this part of London, *except among the persons* who were in the habit of drinking water of the above-mentioned [i.e. Broad Street] pump well. (MCC2, my emphasis)

Another principle that Snow uses at several points is probable reasoning. The basic structure of this principle is to think about what the odds of something would have been against the present circumstances, if not for a certain cause C. For example, he invokes something of this kind when talking about the expected mortality rate in the workhouse and its unexpected lowness, although there are many more cases. In MCC1, he concludes:

> These opinions respecting the cause of cholera are brought forward, not as matters of certainty, but as containing a greater amount of probability in their favour than any other, in the present state of our knowledge. (MCC1: 29).[7]

[6] Note how close this line of argument is to Mill's Method of Agreement: "If two or more instance of the phenomenon under investigation have only one circumstance in common, the circumstance in which alone all the instances agree is the cause (or effect) of the given phenomenon" (214).

[7] It's interesting to note that Snow's reasoning at this point actually appears to be an early version of the no-miracle argument, in thinking that the phenomena concerning cholera would be a miracle unless his hypotheses were true.

3.4 Snow and the traditional theoretical virtues

As we can see, there are several principles of reasoning at work in Snow's case. Snow is much concerned with explanatory considerations, a variety of causal principles, and probable reasoning. I mentioned at the outset that I think the case of Snow shows that we ought to consider certain principles, such as the above, as excellent candidates for epistemic criteria besides the empirical evidence. However, Snow's case can actually also shed some light on some of the more traditional theoretical virtues. Although Snow is less explicit about these, I think it's nevertheless worth mentioning how we can use his case not just for checking whether certain rules are successful, but also for checking whether some of the more traditional virtues play a role. Psillos's list of virtues included avoidance of ad hoc modifications, the generation of novel predictions, and consilience, and it turns out that all of these are virtues that Snow recognises as confirmatory.

3.4.1 Avoidance of ad hoc modifications

The idea behind this virtue is that, in the case of true theories, new (or different) phenomena can be explained by elements already present in the theory, without any need to introduce additional claims specifically designed for the evidence in question but otherwise uncalled for. An example that Snow cites in this context concerns the question of why Scotland was afflicted with cholera during winter, but not England. The miasma theory, unsurprisingly, has nothing to say about this, and although it's not immediately obvious how Snow can account for this, it turns out that his theory can account for what's going on beautifully. Snow's explanation is that the English don't drink much unboiled water, except when it's warm, and that in winter they drink mostly tea, coffee, and malt liquor. The Scottish, however, use unboiled water all the time, either for mixing with whiskey, or if they drink the whiskey straight, they often get thirsty afterwards, and drink large quantities of water then, in this way constantly exposing themselves to the cholera poison.

3.4.2 Generation of novel predictions

Snow predicted a whole host of new data that was found to actually be the case. For example, he predicted a variety of statistical distributions, the contamination of the pump, where there would and would not be high mortality rates, that where the mortality rate was low, there were other water sources, and so on. Although the events this data was based on had already occurred they ought to nevertheless count as predictions: first of all, they are the only kinds of prediction theories like his admit of, secondly, what seems crucial is whether the predictions were already known to be the case or not, not whether they had already obtained or not.

3.4.3 Consilience

The term consilience goes back to Whewell, who explains it this way:

> [T]he evidence in favour of our induction is of a much higher and more forcible character when it enables us to explain and determine [i.e. predict] cases of a kind different from those which were contemplated in the formation of our hypothesis. (1858: 87–88)

Snow's theory also exhibits this criterion. Snow is able to explain a variety of different classes of evidence (not all of which we could examine in detail, unfortunately). To give a quick example, Snow could explain and determine several kinds of statistics (such as the distribution of deaths during the Soho outbreak, and certain numbers about the average populations and the durations of epidemics) despite the fact that he did not initially use statistical evidence in order to arrive at his hypothesis. Snow, in contrast to the miasma theory, can explain completely different and *independent* kinds of evidence – evidence from pathology, epidemiology, and statistics.[8]

4 The testability of epistemic significance

The point of the case of Snow then is to illustrate how we can test empirically for the success of certain principles and rules. As it is only one case, it can never establish the epistemic significance of any rule; however, it is exactly the sort of case we require, and so provides an illustration of the kind of data that we need in order to establish the epistemic significance of certain practices (or any other non-empirical criteria). The case of Snow allows us to see how we can use case-studies to check whether certain styles of reasoning were successful: we can do so by examining certain principles, the conclusions

[8]Snow considers the statistical evidence a "circumstance strongly confirmatory of the communication of cholera" (MCP: 559a).

they lead to, and by determining whether or not the conclusions turned out to be correct. Moreover, checking whether or not a certain state of affairs obtained after it was predicted on the basis of a certain rule is an entirely empirical undertaking. The case in this paper shows that (i) there are cases where we can do this, and (ii) exactly what kind of analysis is required in order to carry out this enterprise.

While not establishing epistemic significance, it shows us how we can do so – by obtaining more data of the same kind. Consider what there is to the epistemic significance of a certain rule (or any other criterion, for that matter): what's required is that using it leads to success, and failing to use it doesn't.[9] This, however, is exactly what we can test for by using cases of the above kind. One case does not establish the epistemic significance of anything, but it is a piece of data that can either confirm or disconfirm the entirely empirical hypothesis that a certain rule is truth-conducive.[10]

This means that the hypothesis of the epistemic significance of the criteria themselves is empirically testable and admits of such testing in the same way other empirical hypotheses do. In the usual scientific cases (or even in 'non-scientific' everyday empirical cases) we begin with a small number of instances, and, as time goes on and we collect more data, we find more and more occurrences of the phenomenon, ending up with more and more evidence, thereby increasing support for our hypothesis.

As more positive instances come to light, hypotheses get confirmed over time (in the case in which they are supported, that is; obviously not in cases where the hypothesis is a bad one – in fact, this is one of the factors that distinguishes good from bad hypotheses). This also goes for our hypothesis: what's required to establish truth-conduciveness is a growing number of cases where inferences were made on the basis of certain rules and practices and where the inferred hypotheses are later confirmed, as is the case in the example of Snow. As we examine cases in the history of science, we find that certain rules and criteria are consistently and constantly successful. As we gather more cases, we can then generate a robust and increasingly solid

[9]It's sufficient for it to do so in sufficiently high proportions; after all, a rule like this is supposed to be conducive to the truth, not guarantee it. For some objections related to this, see section 5 below.

[10]Of course this leaves open some of the details as to what exactly is required in order to establish epistemic significance, for example questions about how many cases are required, and so on. I deal with these questions in section 5 when I discuss objections to this approach.

set of epistemic criteria that can guide us in theory-choice. Of course, it's still possible that I'm mistaken in thinking that the criteria mentioned above actually are truth-conducive; that may be, but what counts is that – either way – this is an empirical question that we can (and should) test.

5 Objections

One might have several objections to the project of showing that certain inferential practices (or non-empirical criteria in general) are epistemically significant. The first objection someone might voice is that there is no way of knowing whether it's not just a coincidence that the rules, in those particular cases, led to the right conclusions. To illustrate this with an example: how can this approach differentiate between 'bogus'-rules and actual epistemically significant rules? For example, it might be the case that we could successfully use the rule 'choose the theory that has the most anagrams of my name' to select a hypothesis that was later found to be true. Likely, this will be the case for some theory and some name, but clearly we don't want to call this rule truth-conducive. But how do we know that our rule is not of the anagram variety? For all we know, it could have been a mere coincidence that our rule led to success. The way to show that more than coincidence is at work here is through showing that there is a systematic correlation between the rules and success. And the claim that there is such a correlation is an empirical claim that is subject to testing and confirmation (or disconfirmation).

This brings us to a related objection. The problem is that no matter how many cases we come up with, one can always argue that the (alleged) connection between the criterion and empirical success is a mere coincidence, and not a sign of an actual and substantial connection. We might have simply observed an unrepresentative subset of cases. This is true, of course, but, once again, this is a general objection, not to our hypothesis about rules, but one that can be voiced with respect to almost any hypothesis whose confirmation depends on repeated occurrences. In that case, however, one is committed to a fairly thorough-going scepticism about confirmation in general. Of course there is no guarantee that we are observing representative subsets; in addition, there is always the possibility of coincidences or miracles. But taking this route commits one to a scepticism that goes much beyond scepticism about theoretical virtues and rules. Secondly, succumbing to this kind of scepticism seriously thwarts the pursuit of the scientific enterprise. Scientific epistemology becomes impossible if one takes seriously the thought that all of our best correlations, our cause- and effect-relationships, for example, are coincidences. Doing scientific epistemology requires rejecting a certain kind of academic scepticism. Note, by the way, that the same goes for our everyday beliefs. As before, denying an epistemic connection between our inferential practices and successful theories on the grounds that it might be a coincidence (and given evidence to the contrary), on pain of consistency, leads to the same results in cases of everyday reasoning. Thus, if in everyday cases of reasoning about mundane objects, I infer certain claims, say when I infer that I had a mouse in my apartment based on odd scratchy noises, and the finding of mouse-droppings, who's to say that the fact that I was right (later confirmed by an actual sighting of the mouse) wasn't just a coincidence, and did not, in fact, have anything at all to do with my employing a successful rule? However, if one rejects this argument in the case of mice, one ought to also reject it in cases of scientific reasoning.

One might also object that, even granting that it's possible in principle to establish the general epistemic significance of certain rules, it will never be possible to eliminate the possibility of coincidence in specific cases. For each individual case, there is nothing that tells us whether it was, in fact, *that* rule that was responsible for success, and whether it was really that rule that was epistemically significant in that particular case (even if the rule in general is known to be truth-conducive). This is true, but once again note that this also goes for instances of everyday reasoning: it's simply impossible to show, for example, that certain causal principles connecting mouse-droppings to mice were truth-conducive in the case of my mouse. Even if it's true that the underlying principle, in general, is taken to be a successful and epistemically significant inferential practice, it's not clear how one could ever show that this particular inference to the mouse was based on it and epistemically significant. The possibility still remains that success in this particular case was a pure coincidence. Note, however, that this doesn't undermine the general truth-conduciveness of the rule, once established. While it might be true that we will never be in a position to establish that specific rules were responsible for success in particular cases, we can establish that a rule is truth-conducive in general (and as a corollary, of course, making it more likely that it was actually responsible for success in that case too, the probability of this depending on the degree of truth-conduciveness of the rule). And that's all that's required. What confirms our hypotheses are all the instances taken together and not isolated occurrences, such a 'global' view being required in order to even make sense of the concept of truth-conduciveness.

Another objection goes like this: we have seen that Snow used a variety of principles in arriving at his conclusions. One might think that this provides additional material for an objection in the following sense: one might claim that the fact that there are different rules at work in different cases means that there is no guarantee that there will ever be any significant overlap among the different cases and that this will prevent us from ever establishing the truth-conduciveness of any of our practices and rules, since their use is simply too isolated. There are several things to be said in reply here: the first and most important one is that the ubiquity of the rules, just as everything else, is an empirical matter. This means that we have to examine a number of cases, discern what rules might have been in play, and see whether there exists a general pattern (if we are unlucky, there will be no pattern at all). However, once again, we don't have to work blindly at this juncture. I suspect that we will find more or less the same patterns and distributions of rules that realists expect to find. And while this is, of course, merely speculation, it can nevertheless provide a useful starting point for enquiry. Certain rules will be more prevalent than others, certain other rules will apply only to very specific circumstances and, as such, not always be applicable, and so on. And while this is, ultimately, only a conjecture, it is a conjecture that can be tested. And the very least that the case of Snow does is to be suggestive of such connections.

Moreover, it is to be expected that we would find a variety of rules across a variety of conditions, some more frequent than others, perhaps. And this should not come as a surprise. Just as in cases of everyday reasoning, we use different principles in different circumstances, this will also be the case for scientific circumstances. The reason this is not surprising, is that certain rules are simply not applicable in certain cases. Thus, for example, when I'm wondering about why my white laundry came out pink, principles of statistical reasoning simply aren't relevant, whereas some sort of causal reasoning involving a pair of red socks is. It might well turn out that certain practices and rules are particularly prevalent in certain scenarios, and this suggests an interesting agenda for future research. It would be useful to find out, first of all, whether this is actually the case, and, if so, secondly, what rules are prevalent in what kinds of cases. We might find, for example, that reasoning by analogy plays a bigger role in reasoning in biology say, than in reasoning in physics. If this were the case, and it could be shown that the various sciences self-select to employ the principles most successful with respect to their specific kind of data, and most advantageous in reaching conclusions of certain kinds, this would shed much light on the way in which scientists and science proceed.

6 Underdetermination again

Let's return to underdetermination. What's the upshot of the Snow-case in this context? It's clear that, if we can establish the epistemic significance of certain rules or other criteria besides the empirical evidence, we have a way to break ties in underdetermination scenarios. In the case in which empirical evidence was the only constraint on theory-choice, what was required of the anti-realist in order to establish underdetermination was only empirical equivalence. Clearly, once other criteria come into play, what the anti-realist would need to establish in order to establish underdetermination is equivalence with respect to these additional criteria, in addition to empirical equivalence. What's more, it is not enough for the anti-realist to show that there is *sometimes* epistemic equivalence of this kind; he needs to show that this kind of epistemic equivalence *always* obtains.

However, the anti-realist might point out that all he needs to do in order to establish this new underdetermination argument is to show that, in each case in which a certain rule was used, there were other rules that would have yielded the same hypothesis (these might even be rules traditionally considered as non-epistemic, such as choosing hypotheses that are anagrams of someone's name, those selected on a specific date, etc.). Thus, the anti-realist might claim that, because of these 'rival' rules the realist is forced into a new underdetermination argument that he cannot escape. Should he invoke further criteria, the anti-realist will do likewise, and so on.

However, while anti-realists might be able to show that there are other rules which would have yielded the hypotheses in question, here's what they wouldn't be able to show (or so I conjecture): that there is a 'rival' rule that could have been used instead of the actual rule *and* that has been empirically shown to be truth-conducive and successful over a variety of cases and circumstances at different times. That is, they will not be able to come up with a rival rule that has empirical support in the way that the realist's rule or set of rules has it.

Consider what exactly it is that the anti-realist would need to show. The argument of this paper has been that there is a set of methodological practices and rules, members of which are routinely used to select certain scientific hypotheses over others. What the anti-realist would need to show in order to maintain a wholesale argument against underdetermination then is this: that *every* possible theory could have been arrived at by such a rival rule. And in order to do this, he would have to show one of the following two:

1. either that there is a rival set of practices and that, for every theory, it is possible to select an empirically equivalent, logically incompatible rival theory by means of members/rules from that set,

2. or, that for every theory, selected by members of the epistemologically significant set, there exists a rival theory that is empirically equivalent to T, logically incompatible with it, and that was also selected by different members or different combinations of members from that same set.

And it is extremely hard to see how this would be possible. But notice that even if it was possible to show that there is a rival selection mechanism for each individual case (which is already highly doubtful, to my mind), what's almost impossible to see is how this could possibly be established in advance by means of a general argument. In the case of empirically equivalent theories, this was always possible by appealing to sceptical hypotheses, but this course is no longer available.

In fact, there are two requirements for the anti-realist that make it hard to see how he could ever produce the required rival rules: (i) first of all, he needs to come up with a rival set that applies to any (possible, future) theory whatsoever (just as, in the standard underdetermination scenario, he needed a way to generate empirically equivalent rivals to any theory whatsoever), (ii) secondly, he needs a way to *generate* such a rival set for any theory whatsoever: he needs an algorithm for doing this. Why the second requirement? The point is that, absent an algorithm to generate such rivals, anti-realists would need to generate their rivals on a case to case basis. And while we might be able to see how this could be accomplished, this is not enough for the anti-realist. As soon as he makes the move to producing rival theories case by case, he loses the *guarantee* that there will always be some rival equivalent both empirically and with respect to the rules. But it is precisely this guarantee that is required in order for the anti-realist to pull through his wholesale argument against realism. Absent any mechanism for producing rivals to the rules, the realist can always argue that it's possible for the rules to actually break the tie, and this, as we have seen, is already enough to undermine the general anti-realist conclusion. Thus, while we can see how to establish equivalence by means of algorithms in the case of evidence, it is not at all clear how to do this in the case of epistemically significant rules.

Another difference is this: realists, at this point, have not completely established the set of epistemically significant methodological rules and practices, but as more data comes in, the set solidifies to the extent that we are increasingly able to capture what its members are. In that sense, we can predict, that in the next successful case, certain members of that set will be used. This is where an additional problem for the anti-realist comes in — he cannot tell us in advance what the members of the rival set would be. He might resort to claim, for example, that there is a weird disjunctive set of rival practices (anagrams *or* a specific date *or* the shortest *or* the easiest to pronounce *or*... and so on). But in that case, he will never be able to obtain the complete set until the end of the universe, or until the last theory is ever proposed, and it's unclear in what sense it's then a *rival* set. It's certainly not rival in the sense that the members of that set have been *well tested over long periods of time*. This last option will be unavailable in the disjunctive case. And note, by the way, that a similar argument is possible with respect to any other criterion that might be established as epistemically significant.

However, it's also important to note the limitations of this argument. Providing an argument establishing the (possible) epistemic significance of our methodological rules and practices, and showing that these rules can, in principle, function as epistemic tie-breakers, does not in fact solve the problem of underdetermination or even preclude its possibility. Even if the rules can break the ties in some cases, this does not mean that they can break the ties in all cases. There are different reasons for this: firstly, there might be underdetermination with respect to *them*, and secondly, we might simply not know how to use them as tie-breakers in particular cases, even though there might be an answer. Again note, however, that this will be the case for *any* criterion whatsoever that the realist puts forward. If the realist can show, for any criterion C, that it is epistemically relevant or significant, there is always the possibility that theories might tie with respect to it, whatever C is. The relevant question then centres on the sense of 'possibility' that's at play here. It is certainly the case that, no matter what C is, there will always be the logical possibility of underdetermination. This means that no argument that establishes the epistemic significance of any criterion will ever be able to show that underdetermination is no longer possible. However, depending on the criterion in question, what it might do, is significantly narrow the possibilities under which underdetermination can obtain.

What this discussion brings out, on a more general level, is that realists and anti-realists pursue different aims: anti-realists want to establish that underdetermination is *always* possible, and realists want to show that underdetermination is rare, but, of course, they want to leave open its possibility. Put another way, realists aren't interested in showing that underdetermination is *never* possible, since, they think, the question of whether there ever is underdetermination or not is empirical. As such, any attempt to try to show that underdetermination is either always or never possible is misguided. What the case of the rules – and the debate about theoretical virtues in general – shows, is that establishing their epistemic significance will never be sufficient to completely rule out arguments from underdetermination, since a new kind of underdetermination with respect to the rules is always possible. However, it also makes it hard to see how anti-realists could ever establish a revised in-principle underdetermination argument that takes such criteria into account – but this is exactly what is required of anti-realists, if they wish underdetermination to remain their wholesale argument against realism.

References

[1] Churchland, P. (1985), 'The Ontological Status of Observables: in Praise of Superempirical Virtues', in P. Churchland & C. Hooker, eds, *Images of Science*, Chicago: University of Chicago Press, pp. 35–47

[2] Churchland, P. & Hooker, C., eds (1985), *Images of Science*, Chicago: Chicago University Press

[3] Friend, M., Goethe, N., & Harizanov, V. (2007), eds, *Induction, Algorithmic Learning Theory, and Philosophy Series: Logic, Epistemology, and the Unity of Science*, Vol. 9, Dordrecht: Springer

[4] Glymour, C. (1985), 'Explanation and Realism', in P. Churchland & C. Hooker, eds, *Images of Science*, Chicago: University of Chicago Press, pp. 99–117

[5] Kelly, K. (2007a),‘ A new solution to the puzzle of simplicity’, *Philosophy of Science*, 74, pp. 561–573

[6] Kelly, K. (2007b), ‘How simplicity helps you find the truth without pointing at it’, in M. Friend, Goethe, N., & Harizanov, V., eds, *Induction, Algorithmic Learning Theory, and Philosophy Series: Logic, Epistemology, and the Unity of Science*, Vol. 9, Dordrecht: Springer, pp. 111–143.

[7] Kukla, A. (1993), ‘Laudan, Leplin, Empirical Equivalence, and Underdetermination’, *Analysis*, 53, pp. 1–17

[8] Kukla, A. (1998), *Studies in Scientific Realism*, New York: Oxford University Press

[9] Lipton, P. (1994/2004), *Inference to the Best Explanation*, second revised edn, London: Routledge

[10] Mill, J.S., (1843/1950) *Philosophy of Scientific Method*, New York: Hafner Publishing Co.

[11] Psillos, Stathis (1999), *Scientific Realism: How Science Tracks Truth*, London: Routledge

[12] Snow, J. (1849a), *On the Mode of Communication of Cholera*, J. Churchill, London

[13] Snow, J. (1849b), ‘On the Pathology and Mode of Communication of Cholera’, *London Medical Gazette* 44, Nov. 2, 745–752; Nov. 30, 929–932

[14] Snow, J. (1851), ‘On the Mode of Propagation of Cholera’, *Medical Times* 3, Nov. 29, 559–562; Dec. 13, 610–612

[15] Snow, J. (1855), *On the Mode of Communication of Cholera*, second much enlarged edn, London: J. Churchill

[16] van Frassen, Bas (1980), *The Scientific Image*, Oxford: Clarendon Press