Brit. J. Phil. Sci. **50** (1999), 265–282

Is Fertility Virtuous In Its Own Right?

Daniel Nolan

ABSTRACT

Fertility (or fruitfulness or fecundity) is often listed among the virtues which are desirable for scientific theories to possess. In this paper I discuss the several species of theoretical virtues called 'fertility', and argue in each case that the desirability of 'fertility' can be explicated in terms of other, more fundamental theoretical virtues.

Lists of desiderata for theories, especially scientific theories, usually include desiderata like empirical adequacy, simplicity, strength (or comprehensiveness) and coherence (both internal consistency, and coherence with other good theories). A common companion to these four is fertility, or fruitfulness.[1] Fertility is taken to be a very important scientific virtue: Kuhn, for example, says that fruitfulness is of 'special importance' ([1977], p. 322); and McMullin ([1976]) finds in an account which provides a central place for fertility 'the best argument for taking the theoretical model of the scientist realistically' (p. 395). Accounting for any cognitive virtue is a difficult task—both in trying to specify how it should operate, and *why* it should operate. It would be nice if we could reduce the general norms of theory construction (either in science, or more generally) to a small list, preferably of elegant, intuitive rules. In this paper I will show that this can be done for fertility, at least: the value of fertility, in so far as it is valuable, can be explained in terms of the other desiderata for scientific theories. This is not necessarily to *deny* that it is valuable, of course—in fact, to explain something's being valuable is to be committed to the thing in question being valuable—but it does deny fertility a role as being of distinctive value, over and above the virtues which in fact its value rests. For example, if the value of fertility is explicable in terms of the value of other standard desiderata, then alas we should not hope for a good argument, let alone the 'best argument', for scientific realism to rest on anti-realism's putative failure to account for the value of fertility.

One problem which makes any discussion of theoretical virtues so difficult is the lack of agreement about what the different putative virtues amount to. This seems to be particularly true of fertility. Some accounts of fertility seem to confuse it with strength or comprehensiveness (as with accounts which take

[1] These are essentially the five virtues selected to be discussed by Kuhn ([1977]).

fertility to be a matter of providing predictions and explanations in a wide range of areas, or which say that a theory's fertility is a matter of the extent to which it 'disclose[s] new phenomena or previously unnoted relationships among those already known' (Kuhn [1977]). Saying more than other theories (about new or old things), or saying more about a wide variety of things, is (are) a valuable feature of theories (at least *ceteris paribus*), but does not seem to be a value independent of comprehensiveness or strength. (Note that comprehensiveness, or strength, may itself have different aspects of value: a large amount of information by one or more metrics; or the virtue of breadth, of saying something about a wide range of topics; or a virtue of depth, of saying a lot about the 'hidden essence', rather than the mere regularities of observables; or the virtue of precision in what is said about a given subject matter. I think all of these, while valuable, fall under the rubric of 'comprehensiveness'—and I wish to distinguish them all from fertility). Were someone to insist that the virtue (or set of virtues) which I call 'comprehensiveness' or 'strength' *should* be called 'fertility', or that 'fertility' should be reserved as a label for one of the sub-virtues under this rubric, then I would have no more than a verbal disagreement with her. However, many who have wished to claim a special virtue for fertility have tended to see it as a valuable feature besides those which I listed in the first sentence of this paper. I will therefore take it that fertility is not merely comprehensiveness, or a variety of comprehensiveness—though again, were someone to insist that it was, I would not take myself to be in substantial disagreement. It is probably worth keeping this confusion in mind in practice, however, since it may well be that cases diagnosed as cases of fertility may be able to be redescribed as simply cases where it is comprehensiveness which provides the cognitive value.

The word 'fruitfulness' is even more liable to this ambiguity: to the extent that fruitfulness is just intended to be comprehensiveness or strength, or a variety thereof, then it is also not my concern, but if it is meant to be something distinct, I will take it that 'fruitfulness' is another name for fertility of theories (whatever that is). 'Fruitfulness' may as a matter of fact have connotations which make it more likely to be taken to be a name for the breadth of a hypothesis (perhaps the unexpected breadth at that). I do not deny that this is a Good Thing for a theory to possess. It is just that it is not the virtue I am concerned to discuss in this paper. The same remarks may be made about 'fecundity', another term for a theoretical virtue, which seems to be obviously a synonym of 'fruitfulness'. There may be those who wish to distinguish the meanings of 'fertility' and 'fruitfulness' and 'fecundity' both from comprehensiveness and from each other, I suppose: I shall have to leave it to those people, if there are any, to decide what they wish to label the phenomena I shall discuss. Enough of labels: let me move on to say what I am to take the theoretical virtue of fertility to be.

One of the most natural-sounding accounts of fertility is such that it makes the value of fertility particularly difficult to understand. This is the account of the fertility of a theory being its posing of new problems which require further investigation, or the opening up of new lines of research which result in elaborations and modifications of the original theory. On the face of it, it can be a little hard to see why the liability of a theory to require improvements, or to raise new problems, should be considered a good thing. It is as if 'Faces many problems' or 'Could do better' or 'Much room for improvement' are high praise on the report card of a theory. Surely being in difficulties is not what we would expect of an ideal theory? If we are faced with two theories which are in other respects equal, except that it is clear that one faces many more problems than the other and would be liable to require extensive modifications if we persisted with it, it is hardly crystal clear that we should prefer this theory to its rival which does not face such problems or immediate likelihood of requiring revision.

The counterintuitiveness of taking fertility to be virtuous will perhaps encourage some philosophers of science, since philosophers of science, perhaps even more than other philosophers, seem to enjoy espousing paradoxical views (compare Popper's injunction to accept the least likely hypothesis, or the common claim that we cannot really know anything about the external world on the basis of experience, or Feyerabend's claim of the equal value of voodoo and western science for explaining the world). I trust this counterintuitiveness of seeming to take the degree to which a theory is in trouble (e.g. through facing problems, or appearing likely to require revision in the near future) as a mark of its virtue will at least be somewhat a cause of concern for many of my readers. We would like a *reason* to think that fertility, construed this way, is valuable before we take it to be so.

One sort of reason is Kuhnian (Kuhn [1977], p. 320, n. 6)). He provides an explanation for why scientists would prefer a fertile theory—because they want something to work on, an opportunity to do science, and a theory riddled with problems which require further work will be more attractive than a sterile theory which provides at most the opportunity to tinker at the edges or to be applied to another area within its purview. This will provide a reason why scientists will prefer to work with fertile theories, and will devote more time and effort to a fertile theory than a sterile one. As a psychological explanation which explains why one theory gets attention over others, this may have something to it. Fertility would then be like the other clearly pragmatic virtues—features of a theory which provide those who consider them a reason to adopt them, and give them at least practical 'acceptance'. Other instances of such pragmatic virtues are: being a theory which is likely to attract research funding for further work; being a theory which will not get a scientist into trouble with the authorities; or alternatively a theory which will gain the

scientist notoriety and therefore fame; being a theory which is easy to understand and so convenient; or alternatively a theory which is difficult enough to provide a good intellectual challenge; and the list could be extended indefinitely.

If fertility were a merely pragmatic virtue of this sort, an explanation of its attractiveness would not be far to seek: we can refer the question to the sociologists and psychologists of science, who will be best placed to tell us why scientists like doing science. (It is unlikely that there will be a single unified reason, of course, or even a handful—but hopefully there is something interesting to be said on the topic). But to stop here will be found too sociological by many. The attraction of a fertile theory is according to some accounts more than this—fertility is thought to be some sort of epistemic or methodological virtue, and is not merely attractive as providing employment prospects for scientists.

A slightly different conception of fertility, but still one which sits as well with our vague pre-theoretic notion, makes it clearer why fertility might be thought to be a methodological virtue, rather than a merely pragmatic one. Fertility, as I have claimed, must be distinguished from mere strength or comprehensiveness. Often, however, for a theoretical insight to be applied to a new area, or to be adapted to cover a previously anomalous situation, the theory is not applied as is, but is changed, adapted, and/or developed, before it can handle the new task. The project of explaining the behaviour of gases atomistically was an extremely fertile one, for example, in countless ways: but the theories of gases we employ today are hardly the same in detail as the atomic theory of gases developed by Dalton. A fertile theory, or research programme, or whatever one prefers to count as the bearer of fertility, opens up new fields, or handles problems in a new or better way, not by remaining static, but by change and adaptation.

If one considers this change to be a change of theory (so a fertile theory is replaced by a close successor when it demonstrates its fertility), the metaphor of 'fertility' becomes clear: for a theory valued for its fertility is not so valued for itself, but for its healthy vigorous 'offspring'. On the other hand, talk of adaptation or development of a fertile theory makes sense: it is intuitive to think that it is a theory's ability to survive and adapt in new areas which helps to mark out its fertility, as opposed to the infertile theory which runs out of new uses or perhaps must be discarded altogether when new challenges loom. In this second way of talking, it is numerically the same theory before and after the change—it has grown and developed, rather than been destroyed and replaced by a successor. Either way seems acceptable in normal conversation: I am of the opinion that identity conditions come in coarser and finer grains, and whether or not we want to say two stages in a theoretical development are stages of the same theory or stages of different theories depends on the context.

Did Kepler and Newton have the same theory of the motion of the planets? Yes and no. Speaking broadly enough (e.g. if the relevant contrast is between geocentric and heliocentric planetary systems, or between circular and elliptical orbits), the answer is yes. Speaking with tight conditions on what counts as the same theory, the answer is no (e.g. Kepler certainly didn't have Newton's views on gravity: so we are tempted to say that Newton's theory was a significant advance on Kepler's. This way of talking seems to presume they were different theories, since a theory shouldn't be a significant advance on itself). Where we draw the lines is often also a vague matter too, but it causes little trouble in practice. Let me talk from now on as if a theory's fertility is concerned with changes to that theory, rather than how that theory is connected to its ancestors and descendants—though if one prefers to describe the theoretical evolution in terms of replacement rather than change, it is easy enough to reinterpret my comments.

So let us say that a theory's fertility (I'll use 'theory' to pick out the bearer of fertility—those who wish to call the bearer of fertility a 'theory sequence' or a 'research programme' or some other such name are invited to substitute their favoured expression) is the extent to which modifications of it (or descendants of it) can carry out such theoretical tasks as explaining new phenomena, resolving anomalies, making unexpected unifications, and other such (for the time being I'll be deliberately vague about exactly what we want fertile theories to do).[2] This form of fertility is highly prized among scientists and philosophers of science, and it seems to appeal to more than the pragmatic interests of practitioners.

Fertility is thought to be important in two ways. One way is the value placed on 'forward-looking' fertility: a theory is valued if it has the prospect of being able to be developed (or replaced by a successor), in the relevant ways. The other way, more controversially, is that a theory with a track record of fertility is often more prized than a rival which has not yet had a track record of fertile development. In the vocabulary of Ernan McMullin, the distinction is between 'P-fertility' and 'U-fertility': the first is *proven* fertility—the fertility a theory has in fact displayed in the past, whereas 'U-fertility' is the *untested* fertility a theory may have (McMullin [1976]). The P-fertility of a theory is, for McMullin, a measure of 'how successful it [the theory] has been in suggesting the right modification at the right time and in allowing incorporation of new areas not originally unforseen' (pp. 400–1). The U-fertility, on the other hand, is for McMullin not as interesting—it is merely the measure of how well the

[2] This seems to accord well enough with the usage of the term by McMullin (see McMullin [1976], p. 400; Chalmers [1979], p. 229), who speaks of the 'objective opportunities for future development inherent in a programme', and Nickles' 'Heuristic Appraisal' (Nickles [1989]) seems to capture a similar notion of the potential a theory has to develop in desirable ways. I do not define fertility in exactly the same way as any of these authors, but the idea is similar enough so that I suppose we are all talking about the same thing (or at least aspects of the same thing).

theory is expected to do in this regard in the future—its 'as-yet-untested promise' (p. 400).

Let me discuss this 'forward directed' or 'U-fertility' first, since I believe explaining the attractiveness of the relevant virtue is reasonably straightforward. Exactly what kind of development is involved is, of course, important. If a high degree of fertility merely involved a high propensity for the theory to change, or require work, undergo alteration, and so on, then it is indeed hard to see why fertility is attractive for other than Kuhnian pragmatic reasons. However, mere change is usually not what people have in mind. A fertile theory *develops*, it improves in various ways—through being applied to new fields (and thus becoming more comprehensive), through changing to be able to make new predictions, or more precise predictions—and not just new or precise predictions, but crucially *correct* predictions: a theory which gets more and more glaringly wrong, even if it is changing a lot, will be unlikely to be counted fertile, and becoming more wrong and inaccurate is not really a scientific virtue by any stretch of the imagination. A fertile theory may contain in it the 'seeds' of a theory which gives a description of several apparently disparate phenomena in terms of the same underlying theoretical resources, thus offering the potential of greater theoretical unification (Watson and Crick's achievement in discovering the molecular structure of DNA provided a very fertile approach to molecular biology, to say the least: and part of the fertility at least was surely the great unifying power this new theory offered). In general, a fertile theory must not just offer the prospect of change and complication: it must offer the prospect of change or complication which is desirable in some way: whether through increasing strength, increasing predictive power and accuracy, increasing unification, or perhaps through offering all of these and other advantages too. When theorists such as Chalmers ([1979]) and Nickles ([1989]) take a future-directed characteristic like fertility to be valuable, or when examples are offered like Galileo's mechanics (Chalmers [1979], p. 229), Watson and Crick's breakthrough (Nickles [1989], p. 175), Einstein's relativistic theory as opposed to Lorentz's program (Zahar [1973], who talks of relativity's 'heuristic superiority'), or even the promise of Joseph Black's proto-caloric theory of heat (Bradie [1980], p. 11), the potential offered is no mere potential for change, but potential for *progress*, or *advance*, in respect of one or more standardly accepted criteria.

Once it is noticed that for a theory to be fertile in the usual sense is for it to have the potential to *advance*, or *improve*, or *progress*, and not merely to change or become more complicated, it becomes easier to account for the value of fertility in terms of the value of other theoretical desiderata. Forward-looking fertility, in this sense, is valuable for the same reason that a lottery ticket is valuable: not because of its intrinsic value, but because it represents the possibility or chance of an outcome valuable in its own right. The value of

'U-fertility', then, can be explained by treating it as being valuable as a means rather than as an end. Of course, this may be true of some or all of the other theoretical virtues, but the value of this forward-looking fertility is parasitic on the value of the other theoretical virtues in a simple way, since it is valuable as a means of attaining a theory with a higher degree of one or more of the other virtues. Fertility in this sense can also be clearly distinguished from strength, or comprehensiveness, easily enough: the comprehensiveness of a theory is a matter of its current breadth and content, while its fertility is the chance that a close successor or a later theory-stage will be significantly better than the current one—perhaps in respect of comprehensiveness, but possibly in respect of some other theoretical virtue or virtues instead.

This account of forward-looking fertility means that its status is not particularly philosophically problematic, but that is not to say that it is not a phenomenon worthy of philosophical and methodological investigation. One of the most interesting questions about forward-looking fertility is the question of how it might be detected or evaluated. There must be limits to how accessible a theory's degree of fertility is, of course: if judgements of fertility are ever to be very useful, it must be more feasible to make an accurate judgement of fertility of an approach than to just improve the theory to the extent an approach permits. On the other hand, forward-looking fertility does not seem to be something which can only be detected with hindsight—it is not only possible to judge with some degree of accuracy which line of inquiry will prove more fertile, it is also vital in practice to make such decisions, given the limitations of time and resources which researchers face. (One cannot simply devote unlimited resources to every avenue of approach open). Nickles ([1989]), in his discussion of 'heuristic appraisal', provides an interesting preliminary discussion of how one might investigate how judgements of forward-looking fertility are made, and their indispensability. The extent to which general principles, as opposed to educated judgement and scientific art, can help to make accurate assessments of forward-looking fertility is of course a controversial issue in its own right. Since these questions concern the detection of forward-looking fertility rather than the nature of what is being detected, however, I shall have no more to say about them here.

This 'reduction' of fertility in terms of other theoretical virtues seems to me preferable to an account of fertility in terms of other virtues to the account offered briefly by Quine ([1973]). Quine briefly discusses the virtue he calls 'fecundity', and which he says is a matter of expediting further extensions of theory (p. 155). I imagine he would agree that, at least when fecundity is an advantage of a theory, the 'further extensions' must be *good* extensions (true, accurate, or otherwise desirable). With this addition, 'fecundity' seems to be functioning as another name for the virtue I have been discussing. Quine's account of it, however, is brief yet clearly different. Here it is in full:

> As for the fourth benefit, fecundity, obviously it is a consequence of the first two, simplicity and familiarity, for these two traits are the best conditions for effective thinking (Quine [1973], p. 155).

Quine is clearly offering an account of fecundity in terms of other theoretical virtues, but in a different sense I think from the sense in which I was giving an account. Quine's 'obvious' consequence relation between simplicity and familiarity, on the one hand, and fecundity on the other, seems not to be an analysis of why fecundity is a virtue: it rather seems to be a hypothesis about when in fact a theory will have fecundity (or at least what leads to fecundity in the specific case Quine is discussing, which is the advantage of the 'molecular doctrine' that our everyday objects are made up of molecules). While it is an 'explanation' of fecundity (fertility) in terms of other theoretical virtues, it seems to me best understood as a psychological hypothesis of a causal nature: what it is about a theory which enables theorists to extend it in useful ways, rather than an account of the nature of fecundity, or an account of what in virtue of which fecundity is a *virtue*. Even as this, it does not seem to me to be correct: some fertile theories are not a matter of applying familiar principles (quantum theory has been incredibly fertile, but to begin with it seemed anything but familiar), nor need they obviously be simple (though the criteria for simplicity are contested enough so that there may not be uncontroversial examples). Furthermore, some simple, familiar theories seem anything but fertile: many of the myriad quotidian 'folk' theories of phenomena are exceedingly simple and paradigmatically familiar: but a theorist seeking a fertile theory would be better advised to examine speculations which go well beyond these theories.

The other aspect of fertility which is sometimes thought to be relevant in assessing the relevant merits of theories is what McMullin ([1976]) refers to as 'P-fertility': the past record of a theory's having been developed and improved in fruitful ways (its 'proven fertility', for McMullin). McMullin argues that the extent of a theory's P-fertility is relevant to confirmation: *ceteris paribus*, a theory with higher P-fertility, a proven track record of development if you like, is to be preferred to a theory with less or no P-fertility. (It is important to stress the *ceteris paribus*: there is no reason to suppose McMullin thinks that fertility is the only feature relevant for epistemic appraisal.) This cannot be justified on the same grounds I provided for forward-looking U-fertility, of course: any more than we should have any particular reason to hold onto our tickets after the lottery has been drawn and prizes claimed. If we take a past history of successful developments and modifications to be relevant to confirmation (as McMullin claims, though as he points out much of the logicist tradition is against him), some other explanation is needed. Alternatively, we could take past, or 'backward looking' fertility to be a *sui generis* theoretical virtue, not in need of explanation and perhaps not possible to explain: but it strikes me at

least that it is not an obvious candidate for such a status. Better that, if we accept it is a virtue, we can explain its desirability in terms of other desirable features of theories. I believe this can be.

Lots of real theories come to have high P-fertility: molecular theories of gasses and quantum mechanical theories of radiation, to name two theories still accepted, and phlogiston theory (according to Bradie [1980]), and ether mechanics to mention two theories which, while exhibiting a fair degree of fertility in their day, are no longer accepted. However, actual cases need to be described fairly carefully, and often possess features which cloud specific issues. So instead I will only discuss schematically a series of versions of a theory V_1, V_2, V_3, and so on. For the theory to display a fertile development, each new stage needs to be able to display some new desirable feature, such as offering explanations or predictions in a new area, or resolving an anomaly, or providing some unification, or whatever. Notably fertile theories, of course, often evolve to possess higher degrees of all of these characteristics.

Let us suppose that, in the theory of a certain phenomenon, we begin a phase of theoretical development with theory V_1. New evidence comes to hand incompatible with V_1, but can be accommodated by altering V_1 to produce V_2 (V_2 can be thought of either as a later stage of the theory V, or as a close successor to V_1, depending on taste). Then, let us suppose phenomena in a new area can be predicted if the theory can be adjusted to remove what is now seen as some previous idealisation, to produce V_3. V_3 is then vindicated, when the new prediction matches subsequent observations. V_3 now has a record of P-fertility, and according to McMullin, so far as I can gather, this means that if we were forced to decide between V_3 and a new theory, hot off the drawing board (U_1, let us say), which was as virtuous as V_3 in all other respects (it is just as unified, it is compatible with previous results and makes as many interesting new predictions, it is just as simple, coheres just as well with the rest of our theoretical picture, etc.), then V_3 is to be preferred.[3] This is one of the historicist, diachronic elements which McMullin wants to add to our theory of confirmation.

Some will find this implausible. Surely if two theories do as well now, accidents of history do not signify? After all, the theories are now being judged against the same evidence, and it might be thought that the V-series is in no better shape than in the counterfactual case where a proto-U theory was

[3] What implications flow from one theory being preferable to another in terms of what is then done with the two theories is of course not a straightforward matter. It does not mean that the less preferred one may or should be instantly assigned to the dustbin: it may be worth a community continuing to work on both theories, and rational partisans may even wish to concentrate more effort on the flagging rival to try to improve it. Lakatos's conception of methodology as having virtually nothing to say to practising scientists (Lakatos [1971], p. 174) goes too far—but the grain of truth in it is that it is far from straightforward how the status of rival theories should flow through to what further work is done, or which theories it is rational to accept or reject.

formulated in the past, underwent modification etc. to yield today's U_1, while our V_3 had just been proposed? Intuitions and other considerations are divided here.

However, I believe that I can to some extent sidestep this debate here. For what I shall try to show is that the intuitive support for P-fertility can be explained by appeal to another historicist candidate for addition to confirmation theory. This candidate is also controversial, if a little more familiar to many. If this candidate is accepted as a theoretical virtue, it can explain the virtue of P-fertility, whereas if it is not accepted, it should seem strange to suppose that P-fertility should be relevant to confirmation. Depending on one's views about this standard historicist approach to confirmation in the end, one should either think that P-fertility can be explained, and so is not fundamental, or is not a virtue at all (which is incompatible with its being a fundamental virtue, presumably).

It is often thought that a theory is confirmed more by evidence which comes forward after it is proposed, or that a theory is confirmed more by confirmatory evidence which it was not originally designed to account for. This is sometimes described as the value of *novel* predictions (and there is a dispute over what it is to count as novel).[4] It is not uncontroversial that confirmation of 'novel' predictions, or confirmation gained from results which were not available or considered when the theory was constructed is indeed a virtue (see, for example, Mayo ([1991])).[5] However, the thought that a theory which predicts correctly and is vindicated is to be preferred to a theory cooked up to fit already known results is intuitive to some extent, and if this is a mistake, then it should at least be an understandable mistake. To take the perspective of hindsight, what this doctrine comes to is that a theory with a history of 'novel confirmation', as I shall put it, is better confirmed, and thus to be preferred *ceteris paribus*, to a theory which lacks this track record.

[4] Worrall ([1985]) discusses several accounts in the course of arguing for his preferred version.

[5] As an aside, what Bayesians should say about the confirmatory value of novel evidence *vis-à-vis* 'old evidence' is an interesting issue, and a point of contention among Bayesians. One straightforward Bayesian approach, which measures the support evidence gives to a hypothesis by the change in probability of the hypothesis before and after the evidence came in which it is obvious that novel confirmation is more valuable than 'old evidence', available before the construction of the hypothesis, but in which the problem is to see how old evidence could provide any confirmatory support to theories at all (see Glymour's well-known objection to Bayesianism on this point in (Glymour [1980], pp. 85–6)). Another common Bayesian approach to judging the value of evidence is to look at how one's rational confidence in a give theory *would* have been altered if the theory had been entertained first and then the evidence had come in (see Howson and Urbach [1989], pp. 270–5 for an example). On this approach there does not appear to be room for distinguishing the support provided for evidence which is actually posterior from evidence that was available and even used during actual theory construction. This latter approach seems to deny the particular value of confirmation of 'novel' predictions, since in effect the value of all evidence is judged as the support it would have provided if it were novel. These do not exhaust the options for a Bayesian by any means (for a discussion of alternatives, see Earman [1992], Ch. 5).

P-fertility is not just another name for this virtue of novel confirmation.[6] Indeed, a P-fertile theory may have very little confirmation of its novel predictions. Consider the case of V_2: it had some P-fertility, since a previous version had been able to be modified to account for a new variety of phenomenon. Had we been using V_2 all along, it would have received some novel confirmation, but as it happens it was only developed in response to the evidence, rather than predicting it in advance. For even more obvious cases, consider cases where a theory's P-fertility has manifested in there being a process of improvements other than improvements in prediction. A theory may have been able to be generalized to yield greater unification of apparently disparate phenomena, or may have been able to be modified to explain what were previously thought to be anomalies. In both cases the development of the theory gives it a positive track record in the a way which counts for McMullin's P-fertility, at least, but in none of these cases need there be novel confirmation.

Conversely, a theory may have a great deal of novel confirmation without much in the way of P-fertility at all. For if a theory has had a string of successes which more than account for the evidence which was employed in constructing it, and indeed that it has been so successful that its initial formulation was sufficient to make the predictions, without needing adjustment or alteration along the way, then the theory has no P-fertility at all: it has not been expanded or altered, but has merely continued in its original form. (Of course this may be unlikely to happen very often in practice, but it is hardly impossible.) P-fertility and novel confirmation are both diachronic types of confirmation, but they are not the same virtue.

Nevertheless, despite their apparent lack of connection, I want to argue that P-fertility can be accounted for in terms of novel confirmation. This will not be an unqualified vindication of P-fertility, since there is the further question of whether there really is such a virtue as 'novel confirmation'; even so, reducing the task of accounting for P-fertility to the task of accounting for novel confirmation (or the task of eliminating P-fertility to the task of eliminating novel confirmation) is still a step forward. One way to account for P-fertility in terms of novel confirmation is to consider meta-hypotheses which receive confirmation through the progress of first-order scientific hypotheses. Hypotheses like 'some atomic theory of gases is true', or 'some caloric theory of heat is correct', or 'the best available theory is going to lie in the tradition of V' are

[6] McMullin makes the point that they are not the same thing (see McMullin [1976], pp. 425–6). Fertile developments are not 'logically contained in' the original theory in the way that novel predictions are (p. 426). Of course, novel predictions are to be distinguished from novel confirmation: novel predictions, in the Popper–Lakatos tradition, are valuable in a theory in their own right, though the corroboration the theory receives by virtue of having novel predictions vindicated (or resist falsification, at least), corresponds pretty well to one common understanding of what it is for a theory to receive novel confirmation: and the sort of corroboration gained by vindication of novel predictions is especially valued by both Popper and Lakatos.

hypotheses about theories, rather than the world directly, but are the sort of things which get confirmed or disconfirmed often by the same sorts of evidence which confirm or disconfirm atomic theories of gases, caloric theories of heat, or theories in the series V.

Presumably scientists often accept these higher order claims when they accept a relevant first-order theory, and indeed may employ these hypotheses when deciding how to modify their views (the latest caloric theory has some trouble, but if one still accepts that some sort of caloric theory is right, then looking for a new modification seems to be the thing to try, whereas someone already convinced that some kinetic theory of heat is going to turn out to be right is presumably less likely to try to rescue the caloric tradition). And a sequence of theories or theory-stages which have each had to be discarded may still provide confirmation for the meta-hypothesis: V_3 might still have trouble or face disconfirmation, but the hypothesis that the V-series of theories is on the right track (or some V-like theory is true, or close to the truth, or the best theory of the phenomena available) may be doing very well nonetheless, and may even be confirmed by the evidence which refutes specific versions. (E.g. new information comes in which means that new phenomenon P can be explained only if parameter values are fixed differently and a previously-unnoticed idealization is dropped, so V_3 must be replaced by V_4: but that the V series provides the sort of framework for integrating P might count in favour of the meta-hypothesis that *some* V-variety theory is correct.)

In general, the move to invoke higher-order theories in this sort of context will not be terribly relevant: for when it comes time to compare a theory with a record of P-fertility with a theory without such a record, the latter theory will also be associated with a meta-hypothesis which will also have a lot of support: the theory that some U-like theory is true will receive a lot of support from the support which the new theory U_1 has, and it is not at all clear that V_3, in the case of the schematic example above, will better support the meta-claim that some V-like theory is true. (And given that U_1 and V_3 were stipulated to be otherwise equally confirmed, apart from any confirmation the track-record of V theories had, it is likely that as far as synchronic confirmation goes, the theory 'some V-like theory is true (or has high verisimilitude, or will be the best available theory)' is as well confirmed as the theory 'some U-like theory is true (or etc.).')

The situation becomes asymmetric, however, if novel confirmation is to count for more than confirmation by 'old evidence'. For the meta-theory concerning V-like theories had presumably been developed back with V_1, and the fact that V-like theories got better and better counts as novel confirmation for the theory that V-like theories are good (that V theories will survive and flourish is what one would expect if they were basically on the right track; though there need be no necessary correlation between these of

course). In retrospect, the evidence and other theoretical background may have also been favourable to the hypothesis that U-like theories were on the right track too (after all, the history meant that one could end up with U_1, a theory which is the equal synchronically to V_1). But the confirmation for the meta-claim about U theories is, as a matter of fact, old evidence by the time U_1 is formulated, since there is not in fact a history of proto-U theories in the same way that there is a history of V theories prior to V_3. So the confirmation (or much of the confirmation) for the meta-claim concerning V is novel, and counts for more than the confirmation for the meta-claim U. It is this disparity which can explain why theories are to be preferred *ceteris paribus* if they have high P-fertility. For having high P-fertility provides a theory with a better confirmed meta-claim: V_3 gets preferred to U_1 because the meta-claim that a V-style theory is the correct one is better confirmed than the meta-claim that a U-style theory is the correct one. V_3 thus inherits some of the 'track record' of success which the V-like approaches have had, while U_1 has yet to stand the test of time.

Identifying the value of P-fertility as being a matter of the value of novel confirmation of the more general meta-claim would explain why it might be thought to be desirable, without having to take P-fertility to be an unanalysed or inexplicable theoretical virtue. It also I think accords with our intuitions about why a theory which is a survivor of experiment and modification is to be preferred, *ceteris paribus*, to one just cooked up. It makes P-fertility very similar to novel confirmation, and appeals to the same special value we place on theories having a record of having stood up to new tests and challenges. So I think the analysis of P-fertility in terms of novel confirmation of a 'programmatic' meta-hypothesis is to be preferred to taking the putative value of P-fertility to not be further explicable. The 'meta-hypothesis' account is not the only way of tying P-fertility to novel confirmation, however, and an alternative can be constructed which has similar features. One worry people might have about the 'meta-hypothesis' account is that it makes the virtue one which is too reflective: we think that the virtue of P-fertility is something to do with the connections between the theory (and its predecessors) and the evidence, and it is somewhat of a worry that the 'meta-hypothesis' requires a dog-leg detour through a theory about theories.

An alternative way of making the connection between P-fertility and novel confirmation is via an appeal to first-order theories (theories about the phenomena in question, rather than theories about theories of the phenomena). As well as quite specific theories, there are also less-content filled, but more general theories. As well as Dalton's hypotheses about the atomic nature of gases, there is also the much more general theory that the behaviour of gases in many cases can be explained by reference to the atoms (or at least molecules) which make them up, and the interactions between these. Dalton's specific

theory was false and is no longer taken seriously: however the more general atomic(/molecular) theory of gases is still believed. Specific theories (or theory-stages, depending on how we wish to talk) are often quickly discarded, but a lineage of these specific theories often has a more general, less content-filled but still significant general theory in common to them: witness the basic idea which runs through atomic(/molecular) theories of gases. (These general hypotheses bear a resemblance to Lakatos's 'hard core' of research programmes, obviously.) Taking P-fertility of a specific theory as a matter of the support given to it by a more general theory which has benefited from novel confirmation can explain why a theory's failed but fertile predecessors make it a more attractive option than rivals who lack such a lineage. The more general theory (e.g. that gasses are made up of atoms(/molecules), and the behaviour of gases can be explained by the behaviour and interactions of these) receives support through providing unificatory power, predictive success, and so on, even though the specific theories incorporating it fall by the wayside. Its confirmation through the history of different specific versions of the general theory provides support for the latest version over rivals which deny the general claim (that is, its realising of some parts of its U-fertility); and the fact that this general claim was advanced before the history of specific versions renders it able to have novel confirmation, over the general claims which are part of a rival sort of theory straight off the drawing board, without the same history of fertility. So V_3 is better off than U_1, to revert to the schematic example, not because of synchronic features of the two (*ex hypothesi* they were otherwise pretty much equal), nor because either has received novel confirmation in their own right, for they are both new versions of theories, but because the general theory running through V_1 to V_3 has a history of novel success (and thus provides confirmation for theories incorporating it over theories which deny it), while the general theory or theories embodied in U_1 lack such a history.

It may be that the 'meta-hypothesis' approach and the 'general theory' approach are different ways of presenting the same connection: at any rate, they are clearly closely connected. And it is again worth emphasizing that it may be that there is no special value for novel as opposed to non-novel confirmation: it is just that it is intuitive that it is, and so any intuition we might have that P-fertility is valuable can be explained with reference to this underlying intuition, whether or not the intuitions are vindicated. If confirmational virtues are all synchronic, then it seems that neither novel confirmation nor P-fertility are to be accepted as genuine virtues, and so the issue of whether the value of P-fertility can be explained in terms of other virtues will not arise. On the other hand, if there *are* diachronic confirmational virtues, which is intuitive to many, we do not need to see P-fertility as being a separate fundamental virtue from the more usual one of novel confirmation, since an explanation of why P-fertility would be valuable in terms of novel confirmation is available.

If I am right about the connection between P-fertility and novel confirmation (and let us assume for the moment that there is some special value to novel confirmation), then a specific theory with high P-fertility is thereby confirmed through another closely connected theory having been confirmed by past novel confirmation (either a meta-hypothesis about what sort of theory will be successful, or a general hypothesis of which the theory and its predecessors are more specific versions). I have not tried to say much about what might be valuable about novel confirmation, and it in turn may not be a fundamental virtue—capable of further explication. Since I have been unspecific about what it is and how it works, there may be other ways of tying P-fertility to novel confirmation other than the way I have mentioned. One may even think that P-fertility is after all a sub-species of novel confirmation, if the value of a theory's being novelly confirmed can somehow be directly passed on to successor theories related in a certain way. This would depart somewhat from our usual way of thinking about novel confirmation: it could turn out that a theory receives 'novel confirmation' from evidence gained before its creation and which formed part of the basis of its construction, if it is enough for evidence to count as 'novel confirmation' in this sense if it has supported a predecessor in the more usual way. Extending the notion of 'novel confirmation' to allow this virtue to be transmitted to successor theories constructed in the light of the evidence which novelly-confirmed their predecessors would make P-fertility of the relevant sort into a version of 'novel confirmation'. I mention this as a possible way of carving up the terrain, but cannot recommend it because of the violence it does to our ordinary notion (indeterminate though it probably is) of what has to be the case for *novel* confirmation to occur. The substance of such a view would be in agreement with mine, however, as opposed to the form in which it would be expressed.

There is a potential difficulty with accounting for P-fertility in terms of novel confirmation in the sort of way(s) I have advocated. Novel confirmation only seems to be applicable in cases where we have *confirmation*, but there seem to be theoretical virtues which, while desirable, are not such that a theory with them is better confirmed. The most obvious of these is strength or comprehensiveness: the bolder a theory is, the more it says and the more precise its predictions are, the more comprehensive it is: and this is surely something we want in theories. But confirmation seems to be inversely proportional to strength: one way to have a theory which is very well confirmed is to not go much beyond the evidence at all: the limit case is a list of data, without any attempt to postulate projectible generalizations. Conversely, a theory which gives one an awful lot of information about the world is often more risky, since there are every so many more places where the world can disoblige. There may be other virtues which have nothing to do with confirmation as such: doing well on heuristic measures, so that a theory is easy to

understand and easy to pass on, may be too anthropocentric a virtue to give any indication as to whether the theory gets it right. And there may be others. In my discussion of U-fertility, I said that the relevant potential had to be that the theory was to become better (or be replaced by an appropriately connected better successor theory), but I did not require that this improvement necessarily had anything to do with confirmation. Indeed, one of the paradigm cases of 'fertility' is the case when a theory turns out to be able to be adapted so as to have applications in a field other than the one for which it was formulated, and in this case the 'improvement' seems as much an improvement which yields a gain in comprehensiveness, as in any virtue which yields confirmation. So it would be well, in the case of U-fertility, not to tie the improvements necessarily to other virtues involving confirmation. (Especially since the chance at any sort of improvement is derivatively valuable, not just chances at improvements in confirmation of our theories.) Might the same be tempting in the case of P-fertility? Might a track record of improvements, regardless of their nature, be evidence that the theory is on the right track?

While it would make P-fertility and U-fertility more alike to say that P-fertility is a matter of a record of improvements of any sort, I do not think it is plausible that P-fertility, conceived of as a confirmatory virtue (as McMullin at least seems to do), can be a matter of any old improvement. For if a given virtue does not provide confirmation (e.g. comprehensiveness, heuristic appeal), surely having acquired it in the past, or being the descendent of a theory which had it, cannot be a *confirmatory* virtue. Coming from a line of comprehensive theories does not itself make a theory more likely to be *correct*, for example. So P-fertility, if it is to play the role in confirmation which McMullin, at least, supposes, cannot be as broad as U-fertility.

Since P-fertility can be accounted for in this way, it seems that for many purposes its significance will be no greater than the significance of novel confirmation. In particular, McMullin's claim ([1976], p. 402, [1968], pp. 395–6) that fertility provides a good argument for realism in the philosophy of science looks dubious. Some non-realists will of course not accept novel confirmation or P-fertility: logicist verificationists, for example, let alone that brand of falsificationist (e.g. the early Popper) who have no truck with *confirmation*, novel or otherwise. But it is not at all clear why many non-realists will be worried. Instrumentalists, for example, may accept that there are all sorts of 'confirmation': it is just that 'confirming evidence' does not point to the truth of the theory, but only the usefulness, or perhaps how much confidence we should have in the theory's observable predictions. And constructivists in general do not seem to be committed to anything which is incompatible with thinking that fertility is one of the marks of a theory people take into account when deciding whether or not to believe it (or be committed to it, or whatever the appropriate attitude might be). At the very least, the prospects of an

argument which is not merely a rewording of one of the standard realist/non-realist debates looks very slim. McMullin's assertion that the 'surplus content' of models which makes them fruitful 'is our assurance that the model-structure has some sort of basis in "real world"' (McMullin [1968], p. 395) is so far as I can see either blatantly question-begging, or embodies an assumption that confirmation of a theory has to be confirmation of that theory interpreted realistically (an assumption anti-realists are surely likely to deny).

Fertility, then, can be taken from the list of fundamental theoretical virtues: since both the prized forward-looking fertility, and the arguably confirmationally relevant history of past fertility can have their charms accounted for. Of course, nothing I have said here strictly proves that one could not think that one or the other was a *sui generis* virtue, which turns out interestingly to be necessarily accompanied by other attractions which do much of the same methodological work: but this multiplication of virtues would at least stand in need of some further justification, one that I have never seen. Of course, in explaining fertility, I am not explaining it away: fertility may not be a fundamental virtue, but is arguably virtuous none the less (forward-looking fertility seems clearly to be, whereas P-fertility's status hangs, I believe, on more general issues concerning diachronic confirmation). And merely to show that it is not fundamental is of course not to solve all of the interesting philosophical problems surrounding this under-explored aspect of evaluating theories: I have already mentioned the interesting issue of how one is to detect U-fertility, but there are a host of others. Hopefully, however, a better understanding of fertility will help with the other philosophical issues also.

Acknowledgements

Thanks to Greg Restall, Caroline West, and especially James Chase for discussion and advice.

Department of Philosophy
Macquarie University
North Ryde NSW 2109
Australia
dnolan@laurel.ocs.mq.edu.au

References

Bradie, M. [1980]: 'Models, Metaphors and Scientific Realism', *Nature and System* , **2**, pp. 3–20.

Chalmers, A. [1979]: 'Towards an Objectivist Account of Theory Change', *British Journal for the Philosophy of Science*, **30**, pp. 227–33.

Earman, J. [1992]: *Bayes or Bust?* Cambridge, MA: MIT Press.

Glymour, C. [1980]: *Theory and Evidence*, Princeton: Princeton University Press.

Howson, C. and Urbach, P. [1989]: *Scientific Reasoning*, La Salle, IL: Open Court.

Kuhn, T. [1977]: 'Objectivity, Value Judgement and Theory Choice', in T. Kuhn, *The Essential Tension*, Chicago: University of Chicago Press, pp. 320–39.

Lakatos, I. [1971]: 'Replies to Critics', in R. C. Buck and R. S. Cohen (*eds*), *P. S. A. 1970: In Memory of Rudolf Carnap*, Dordrecht: Reidel, pp. 174–82.

McMullin, E. [1976]: 'The Fertility of Theory and the Unit for Appraisal in Science', in R. S. Cohen *et al.* (*eds*), *Essays in Memory of Imre Lakatos*, Dordrecht: Reidel, pp. 395–432.

McMullin, E. [1968]: 'What Do Physical Models Tell Us?' in B. van Rootselaar and J. F. Stall (*eds*), *Logic, Methodology and Philosophy of Science III*, Amsterdam: North-Holland, pp. 385–96.

Mayo, D. G. [1991]: 'Novel Evidence and Severe Tests', *Philosophy of Science*, **58**, pp. 523–52.

Nickles, T. [1989]: 'Heuristic Appraisal: A Proposal', *Social Epistemology*, **3**, pp. 175–88.

Quine, W. V. O. [1973]: 'Posits and Reality', in R. E. Grandy (*ed.*), *Theories and Observation in Science*, Englewood Cliffs: Prentice-Hall, pp. 154–61.

Worrall, J. [1985]: 'Scientific Discovery and Theory-Confirmation', in J. Pitt (*ed.*), *Change and Progress in Modern Science*, Dordrecht: Reidel, pp. 301–31.

Zahar, Elie [1973]: 'Why Did Einstein's Programme Supersede Lorentz's (II)'? *British Journal for the Philosophy of Science,* **24**, pp. 223–62.