Confirmation: two types, both logical not historical.

JOHN WORRALL

(forthcoming in C.Cheyne and J. Worrall (eds) *Rationality and Reality: Conversatons with Alan Musgrave.* Kluwer – Australasian Studies Series, 2005. NOT FOR QUOTATION IN ITS PRESENT FORM))

There are very many topics in philosophy of science on which Alan Musgrave and I see eye to eye. So it has not been easy to do the decent Popperian thing and pick a (friendly) fight with him. However, thinking again about his much-read 1974 paper on theory-confirmation ('Logical versus Historical Theories of Confirmation') solved my problem. Despite having much of its heart in most of the right places, both the argument of that paper and the position it ends up endorsing are, I believe, importantly off-beam. In this paper I shall explain why and clarify what I think is the correct account of the issue that he addressed. I shall then take the opportunity to contrast my views on confirmation with those of Deborah Mayo, who was herself much influenced by Alan Musgrave's paper and who has developed her own view on the issues it raises. Although Alan's paper was published in 1974, the problem it faces has not been given a satisfactory resolution – at least not one that has met widespread acceptance. So it remains a live issue within current philosophy of science.

Musgrave begins his paper with a sharp formulation of the prediction versus accommodation issue: is there some epistemic premium on *predictive* success? That is, does a theory obtain, *ceteris paribus*, more confirmation from a piece of evidence that it correctly predicts than it does from an otherwise 'equivalent' piece of known evidence that it correctly entails?

He takes it that a 'purely logical' account of confirmation must answer 'no' to this question. Any such account sees confirmation as entirely based on the logical relationships between the theory, T, and the piece of evidence, e, at issue; and hence must entail, whatever the details of the logical relationships it highlights, that whether or not e was already known to hold when T was proposed is entirely irrelevant to confirmation. All logical accounts have their difficulties - in particular, in Musgrave's view, they supply no satisfactory answer to the 'paradox of confirmation.'

An historical (or more accurately – as he allows - a "logico-historical") account, on the other hand, sees confirmation as a relationship, not just between T and e, but also a third variable: 'background knowledge', b. All variants of the historical view entail that T fails to be confirmed by any e that is in b, even if T (of course, in conjunction with appropriate initial conditions and auxiliaries) entails e. All variants of this account are indeed at least partly historical since the answer to the question 'does e confirm T?' may be different in two different historical epochs, because these will be characterised by different states of background knowledge.

But which evidential results exactly should belong in 'background knowledge' and hence fail to be possible confirmers of new theories? Musgrave distinguishes three versions of the historical approach, based on three different answers to this question.

According to the first account, 'the strictly temporal view', background knowledge contains 'all the relevant experimental results, hypotheses, *etc.*, which are "known to science" when [the] theory [in question] was proposed' (p.8). This entails that a theory T is *only* confirmed by facts that were unknown at the time of T's initial proposal and cannot be confirmed by any evidence that was already known to hold. Musgrave points out *both* that this suggestion flies in the face of quite clear-cut intuitions about some particular cases (e.g. that GTR was confirmed by getting the already well-known details of the precession of Mercury's perihelion right) *and* that it seems difficult to see a general rationale for giving such a crucial role to purely temporal considerations.

On the second account, the 'heuristic view', the relevant background knowledge for assessing the confirmation of theory T is restricted to those known facts and results that *were involved in the development* of T. This gives scope for the recapture of some of the intuitive judgments about particular cases: GTR may be confirmed by the details of Mercury's orbit, for example, provided that those details played no role in the construction of GTR. However, aside from the issue of whether there is a convincing rationale for this account, it is, argued Musgrave, altogether too person-relative: ' [i]f different scientists take different routes to the same theory, then the evidential support of that theory as proposed by one of them might be different from its evidential support as proposed by the other.'(p.14) And he, not unreasonably it would seem, regards this as in effect a *reductio ad absurdum* of the account.

Musgrave is inclined to endorse the third variant ('for my money it is the best version of the historical approach to confirmation'(p.19)) This holds that the relevant 'background knowledge' for T consists only of the 'touchstone theory' for T – in effect T's most plausible rival. A theory T is then confirmed by any correct piece of evidence e that it entails provided that e is not also entailed by its 'touchstone' T'. Clearly there will in general be two types of such evidence: evidence that contradicts the touchstone T' and evidence on which T' is simply silent. On this account, GTR is confirmed by getting the details of Mercury's perihelion correct, since its rival, classical physics, gets those details wrong. Ditto with the Special Theory and, say, the Michelson-Morley result.[1] On the other hand, neither STR nor GTR is confirmed by any result that classical physics also correctly entails.

But *why* would anyone want to endorse these judgments delivered by the third variant? Of course, scientists will be especially interested in the question of whether Relativity theory is *better* confirmed, gets *greater* empirical support from the total evidence, than classical physics, and this will direct special attention to those pieces of evidence that are entailed by Relativity but not also by classical physics. But all that we can justifiably say is that those phenomena that they jointly entail fail to distinguish between them or supply no basis for a preference for one of them; it does not of course follow that the jointly entailed phenomena fail to support the new theory. Surely the reasonable judgement is that *both* classical physics *and* relativity theory are confirmed by, say, the precession of the equinoxes since they both entail the relevant data - this is why that phenomenon, unlike, say, the precession of Mercury's perihelion, is irrelevant to the *comparison* of the degrees of evidential support of the two theories.

---

[1] Of course it depends which *particuar versions* of classical physics we are considering – therein lies much of the tail that will unfold in this paper.

Indeed, Musgrave's preferred alternative not only fails to have an obvious general rationale, it would lead to any number of intuitively extremely awkward consequences. Consider what the account says, for example, not about the confirmation of the newer theory in some case of inter-theoretic rivalry but about the confirmation of the older one. The account seems clearly committed to the view that, for example, Newton's theory was confirmed by, say, the precession of the equinoxes in, say, 1900 (when its 'touchstone theory' was what? Galileo's (very partial) mechanics? or Aristotle's more comprehensive but hopeless system?); but by 1914, when nothing had changed either in the theory or (of course) in the phenomena, it was no longer confirmed by those phenomena because a new theory, Relativity, had arisen that equally well entailed a correct description of them. (The alternative suggestion within this approach – that the right way to view the empirical support gained by a theory is always to take as background knowledge that theory's chief rival *at the time it was introduced* - is worse, *much* worse. This alternative would, admittedly, allow Newton's theory to retain its support from phenomena such as the precession of the equinoxes or the existence of Neptune, but it would of course have the consequence that, as well as there being empirical phenomena (like the Michelson-Morley experiment, and the precession of Mercury's perihelion) that support Relativity but not Classical Physics, there are also phenomena (like the precession of the equinoxes or the existence of Neptune) that support Classical Physics but not Relativity. And this, despite the fact that Relativity entails correct descriptions of these phenomena too!

We saw that Musgrave castigates the other two alternative construals of the historical approach as lacking any obvious rationale, but, as we have just now seen, his own preferred version certainly does no better. This surely raises the general issue: Why should 'background knowledge' *in any form* be a factor in empirical support? There is of course, as just remarked, an obvious rationale for taking it to be a factor in *increased* support: if we are interested in why new theory T is *better* supported than its earlier rival T', then results that T' has already either predicted or explained will be in background knowledge and will drop out of the equation – if T' had already got e correct, then, even should T also yield e, then this trivially will provide no reason to prefer it over T'. But it seems difficult to see why the fact that an empirical result e is already in background knowledge *in any sense* should by itself totally rule out e as support for some newly proposed theory T, in the *non-incremental* sense of support. (In the context of the Bayesian account of empirical support this is of course the 'problem of old evidence'.)

Various remnants of Popperianism in Alan Musgrave's paper suggest that he holds that a justification for giving background knowledge this central role might be developed by considering which bits of evidence do or do not supply a proper *test* of the theory concerned. Now it is true, of course, that if we already know that e holds rather than some alternative result of the experiment or observation it describes, the fact that it turns out that some new theory T entails e rather than any of the alternatives will not have us on the edge of our seats wondering if the theory might turn out to be refuted. In that sense there is no test from old data. But why should that sense have the slightest epistemic relevance? The new theory is by no means *a priori* guaranteed to correctly entail all the phenomena correctly entailed by its predecessor. (Indeed if Popper's account of new theories as 'bold conjectures' were true, it would be a miracle if this happened in a field where the old theory had had any considerable degree of empirical success). Still less is there an *a priori* guarantee that the new theory will get right all known phenomena – whether or not dealt with successfully by its predecessor. And indeed few, if any, theories do get *all* known phenomena correct (at least when first proposed). There is, then, a clear sense in which such a theory *was* tested by the already known data: it might have entailed different data that contradicts that actually recorded, but in fact it didn't. To the extent that a theory might have got some already known phenomenon wrong, but in fact got it right, it seems perverse to rule ahead of time, that this success fails to count as surviving a 'test', and so cannot yield any degree of empirical support for that theory.

Hence Alan Musgrave's solution is wrong; and, as so often in philosophy, this is because he has got the problem wrong.[2]

The real problem: prediction versus *accommodation*

The problem is not whether new evidence counts more than old – it doesn't (at any rate it doesn't just because it's new). The problem is *adhoc*ness (indeed the real problem is perhaps seeing that the *adhoc*ness problem is the *only* problem in this area). In the early 19th Century, the classical wave theory of light predicted the results of various diffraction or interference experiments. Intuitively these results told very strongly in favour of this theory against its rival - the emission or corpuscular theory of light. Yet, as we would predict on Duhemian grounds, the emissionists by no means immediately surrendered. Duhem emphasised that single 'isolated' theories such as the corpuscular theory have no empirical consequences of their own, but achieve them only when conjoined with both specific assumptions (what velocities do the light-corpuscles have? And what masses? Most importantly what forces are they subjected to in particular circumstances?) plus further auxiliary and instrumental assumptions.

It follows that there is always logical leeway for holding onto the central theory in the light of experimental 'anomalies' and looking to modify either a specific or auxiliary assumption. 18th and 19th century corpuscularists duly obliged – some postulated, for example, a force of diffraction, exercised on the light-corpuscles as they passed the edges of any 'gross' opaque object; others considered the possibility that the fringe phenomena that wave theorists attributed to interference and/or diffraction were in fact physiological phenomena. Although in this case it was never achieved, it clearly has to be possible in principle for the emissionists to have given themselves an expression for the 'force of diffraction' with so many parameters that, any given particular fringe phenomenon could have been accommodated. Certainly by appealing to (unknown) physiological facts about vision an entirely cheap corpuscularist 'explanation' was suggested at the time and could have been developed in some detail.

[2] Before outlining what I see as the real problem and its solution, however, I need to note his characterisation of 'independent testability'. Musgrave takes it that the idea that what scientific theories need to be accepted is not just testability, but *independent* testability is captured by the third variant of the historical approach: T is independently testable through any of its empirically checkable consequences that are not also consequences of its 'touchstone' T'. As we shall see, independent testability is the key both to the correct construal of the problem here and its solution, but Musgrave's account of independent testability is the wrong one. **Probably just a footnote**

Or consider another case where this sort of dodge definitely works ('works' in the sense that it does produce a theory that yields the accommodated data, not of course in the sense that it produces a scientifically respectable theory.)

The fossil record looks like strong confirmation of the Darwinian theory of evolution. (Of course the situation is rather complex in this example because that theory does not actually deductively entail any particular aspect of the fossil record, but this is inessential to the point at issue.) As is well known, however, it is trivially easy for the "scientific" creationist to "match" this success. All that she needs to do is follow Gosse and assert that God decided, when creating the Universe in 4004 BC, to include some pretty pictures in some rocks that look awfully like the marks of the skeletons of now extinct organisms but are in fact *just* pretty pictures, and to include some buried bone-*like* objects that seem to fit together to form the skeletons of impressive and now extinct creatures but are in fact just artefacts, and so on. She will thus create a version of "scientific" creationism that entails the correct facts about the (now alleged) "fossil record", but clearly it would be absurd to hold that this means that the view that this record supports the Darwinian theory over its rival must be abandoned.

There is a long tradition in science of deeply engrained distrust of such *ad hoc* moves. We surely require an account of the confirmation of theories by evidence that underwrites the judgement that the interference effects continued to give grounds to prefer the wave theory in the early 19th century even once it had been indicated that emissionist accounts could be constructed, and similarly underwrites the judgment that the fossil record continues to give good empirical reason to prefer the Darwinian theory even after creationists have availed themselves of the 'Gosse dodge'. But how *exactly* are we to capture these judgments within a generally defensible account of confirmation?

The obvious initial suggestion is to say that no theory can be confirmed by evidence that it has simply accommodated in this *ad hoc* way, where the advocates of the theory have taken the evidence at issue as given and used it to produce a specific version of their favoured theory that yields that evidence. These are, at least when the notion is used liberally, all exercises in parameter-fitting. The idea behind the 'diffracting force' emissionist account of fringe-phenomena was to start from a very complicated expression for the force as a function of the distance from the diffracting object (allowing this to be attractive at some distances and repulsive at others) and then use particular fringe measurements to fix those parameter values so that the required phenomena are entailed. Similarly, the Creationist's general theory – that God created the Universe in 4004 B.C. 'essentially' as it now is - effectively gives him a whole series of 'free parameters' specifying how *exactly* it was that God chose to create the universe: if you observe particular patterns in some rocks, then that specifies one part of God's creation, you tie this 'parameter' value down on the basis of the observation and, unsurprisingly, produce a specific theory that entails the observed data - the theory being of course that God created the Universe not just any old how but in particular with these patterns in these rocks.

The positive side of the account would then be that a theory is confirmed by any piece of data a correct description of which it entails, provided that the evidence was not used in the construction of the specific version of the theory that entails it, *whether or not* the data was already known. There appears to be, then, an important methodological distinction between accommodation and prediction in the general sense in which it is often used in science (meaning simply that some evidence follows from a theory without having needed to be accommodated within it)[3].

This 'heuristic account' is essentially Alan Musgrave's second variant of the historical view. Any version of the historical view says that elements of the relevant 'background knowledge' fail to confirm, even when entailed by the theory concerned and the second version of this view specifies that the particular elements of 'background knowledge' that fail to support a theory are just those elements that were used in the construction of the theory. This account, which I have before defended myself in a number of places,[4] has a couple of immediate advantages: *first* it accords with a range of intuitive judgments about particular cases (one such is the precession of Mercury's perihelion and the General Theory of Relativity) where 'old evidence' is taken to provide strong support for a theory and *secondly* (and of course relatedly) it relegates the time-order of theory and evidence *in itself* to what it should be – namely, a complete historical irrelevance (what possible *general* justification could there be for old evidence always to count less?).

However the heuristic account has been alleged to face at least two fundamental objections of its own. The objection that Musgrave himself cites, as we already noted, concerns the fact that the account seems to make theory-confirmation an unacceptably relativistic (enquirer-relative) affair:

"If different scientists take different heuristic routes to the same theory, then the evidential support of that theory as proposed by one of them might be different from its evidential support as proposed by the other. In short, Zahar's [heuristic] view makes confirmation a person-relative affair." (*op. cit.*,p.14)

An evenly more frequently voiced criticism of the heuristic view is that, just like the purely temporal view that it replaces, it flies in the face of deeply held intuitions about particular cases. Nickles, Mayo, Howson and others[5] have all pointed to cases in which evidence e was used in the construction some theory T and yet where e was, it is claimed, taken to provide (strong) support for T. As Colin Howson, for example, claimed, the idea that evidence used in the construction of a theory cannot be used in its support "makes nonsense of quite basic and eminently reasonable scientific appraisals" and is on that and other accounts "entirely bogus".[6] In the next two sections, I address these objections in turn and in the process give what I hope is an important clarification of exactly what the heuristic view does and does not claim.[7]

First objection against the heuristic view: used data sometimes (strongly) confirms

It will be best to address first the direct criticism – that there are clear cut cases where used data confirms (even strongly confirms) the theory in whose construction it was

---

[3] quote from French
[4] ref
[5] ref
[6] ref
[7] My treatment here follows and builds upon that given in my [2002] – actually written for a conference in 1999.

used; the response to this will then indicate the answer to Alan Musgrave's 'overly enquirer-relative' criticism.

Allan Franklin once gave a seminar talk at the LSE under the title '*Ad hoc* is not a four letter word'. Underneath the surface correctness of this title, there lies a somewhat deeper but no less correct point: scientists entirely legitimately use data all the time in the construction of their theories. If general theoretical considerations leave the value of some important parameter open, then how else would a scientist tie down that parameter's value except by using data? The only other alternative that seems open would be to conjecture a value and then test – but this attempt to find a needle in a (generally nondenumerably large) haystack would be madness. Here is one simple but canonical instance.

Suppose a mid-19th Century scientist already accepted the *general* wave theory of light - the theory that light from any particular source consists of waves of some wavelength or other transmitted through the luminiferous aether. This general theory does not specify the wavelength of any particular kind of monochromatic light - say light from a sodium arc. The scientist would like a more detailed theory that *does* specify that wavelength. Rather than attempt to conjecture a value, she would 'deduce' the specific theory, involving the specific value of the wavelength, 'from the phenomena'. She would look for some consequence, e, of her general theory T, where e characterises some observable magnitude (fringe separation in some particular experiment, say) as a one-to-one function of the wavelength. She would perform the experiment using light from a sodium arc, measure the magnitude at issue - here, the fringe separation (call the result of this measurement e') - and infer to a more specific theory T'. So for example, subject to a couple of idealisations, it follows from the general wave theory that, in the case of the famous two-slit experiment, the (observable) distance X from the fringe at the centre of the pattern to the first fringe on either side is related to (theoretical) wavelength λ, via the equation $X/(X^2 + D^2)^{1/2} = \lambda/d$ (where d is the distance between the two slits and D the distance from the two-slit screen to the observation screen - both of course observable quantities). It follows analytically, of course, that $\lambda = dX/(X^2 + D^2)^{1/2}$. But all the terms on the right hand side of this last equation are measurable. Hence particular observed values will determine the wavelength, and so the more specific theory T', with the parameter that had been free in T now given a precise value. Far from being scientifically questionable, this is, to repeat, entirely standard (and patently legitimate) scientific procedure.

Several of most celebrated episodes from the history of science involve using data (often anomalous data for an earlier theory) to construct a new theory. For example, Adams and Leverrier used the data from Uranus's orbit that had proved inconsistent with the initial Newtonian account essentially as follows. They took it that the basic Newtonian theory (of mechanics plus universal gravitation) was correct, and then worked backwards from the Uranian data to work out what assumptions would have to be made about a further trans-Uranian planet, such that, when that further planet's gravitational interaction with Uranus was taken into account (along of course with the gravitational interaction with the sun and the other, already known planets), the overall Newtonian theory would ascribe the correct orbit to Uranus. This manoeuvre, as is well known, led to the discovery of Neptune - one of Newtonian theory's greatest successes and indeed one of the most impressive confirmations of any theory in the history of science.

So how in light of facts like these, could anyone have held the ‘heuristic account’ of confirmation which seems committed to the view that evidence used in the construction of theories *cannot* confirm those theories? In the wave theory case in particular, there is a very clear sense in which e’, the fringe data used in the construction of the more specific wave theory T’ supports that theory: given that the general theory T has already been accepted, e’ deductively entails T’, and what better support could there be than deductive entailment?

Colin Howson likes to emphasise a still more general sort of case - standard statistical examples such as the following We are given that an urn contains only black and white balls though in an unknown proportion; we are prevented from looking inside the urn but can draw balls one at a time from it. Suppose that a sample of size n has been taken (with replacement) of which k have been found to be white. Standard statistical estimation theory then recommends the hypothesis that the proportion of white balls in the urn is $k/n \pm \varepsilon$, where ε is calculated as a function of n by standard confidence interval techniques. The sample evidence is the basis here of the particular hypothesis constructed and surely also supports it at least to some degree - the evidence for the hypothesis just is that a proportion k/n of the balls drawn were white.

Deborah Mayo cites and analyses in more detail the same case and also cites the following "trivial but instructive example" ([]p.271). Suppose one wanted to arrive at what she describes as "a hypothesis H" about the average SAT score of the students in her logic class. She points out that the “obvious” – in fact, surely uniquely sensible - way to arrive at H is by summing all the individual scores of the n students in the class and dividing the result by n. The "hypothesis" arrived at in this way would clearly be "use-constructed". Suppose the constructed “hypothesis” is that the average SAT score for these students is 1121. It would clearly be madness to suppose that the data used in the construction of the "hypothesis" that the average SAT score is 1121 fails to support that hypothesis. On the contrary, as she writes:

“Surely the data on my students are excellent grounds for my hypothesis about their average SAT scores. It would be absurd to suppose that further tests would give better support.”[8]

Exactly so: the data provide not just excellent, but, short of some trivial error, entirely *conclusive* grounds for the "hypothesis" - further tests are entirely irrelevant. (This is exactly why it seems odd in the extreme to talk of a "hypothesis" at all in these circumstances – a point to which I will return *below* in my more extensive consideration of Mayo’s views.)

Does the admission that these sorts of ‘deductions from the phenomena’ provide clear-cut cases of theories that are supported by data used in their construction spell the end for the heuristic account of confirmation?

---

[8] ref

To start to see that the answer is 'no', consider again the "Gosse dodge" within "scientific" Creationism, or indeed any of the other standard cases of blatantly *ad hoc* moves in defence of a theory that have been cited in the literature. In all these cases, the specific theory is 'deduced from the phenomena' – meaning, as always, of course deduced from the phenomena *plus already accepted general principles*. This, as Newton emphasised, is a very powerful technique in the case where the necessary general principles are indeed *generally* accepted and therefore, presumably, themselves have strong evidence in their favour. In the case, however, where some at least of the 'background' general principles necessary for the deduction of the specific theory concerned from the phenomena are simply accepted by some particular group and cannot themselves legitimately claim strong empirical support then a much more cautious methodological judgment is in order.

While it is indeed it is true that if you *were already convinced* of the general Creationist claim that God created the Universe "essentially" as it now is in 4004 B.C. then the data that your irritating Darwinian supporters insist on calling the "fossil record" do of course deductively entail the more specific version of your theory that says that part of God's creation was some pretty pictures in the rocks and buried bone-like artefacts, and so on. Those data thus give you not only good but conclusive reason to accept that particular version of the general theory that you already accepted on other grounds. In this regard the case is surely no different from the (intuitively more scientifically respectable) case of the early 19th Century optical scientist, who, being already convinced of the general wave theory, deduces the more specific version with specific wavelengths for light from particular monochromatic sources from the phenomena: again, *given* that she accepts the general wave theory, T, the fringe data, e', gives her entirely conclusive reason to accept the particular version of the theory T', involving a now fixed value of an initially free parameter.

But the intuitive reaction to the Creationist/Gosse dodge case is surely that while the 'fossil record' data may indeed give you reason, in fact conclusive reason, to adopt the particular Gossefied version of Creationism, this is an ineliminably conditional judgment – the evidence gives you absolutely no reason to have adopted the general Creationist view in the first place. If you are going to be any sort of Creationist at all, then this data gives you as solid a reason as there could possibly be for being a Gosse-dodge-Creationist, but it gives you absolutely no reason to be any sort of Creationist at all! There is no reason to think that the general underlying theory itself obtains any empirical support just because the specific version of it entails the correct empirical data.

What is sauce for the goose is sauce for the Gander. Exactly the same judgment is valid in the (intuitively scientifically respectable) wave theory case: the fringe data, e', give you solid (indeed conclusive) reason to believe T' (the wave theory with a specified wavelength for monochromatic light from a sodium arc), *provided* that you have already accepted the general wave theory (with free parameter), but give you absolutely no reason to accept the general wave theory in the first place. Both in this – seemingly legitimate – case and in the, apparently illegitimate case of the Gosse dodge, the right judgment seems, then, to be twofold: *first* that, if the general underlying theory is taken as given, then if e is used in the construction of a specific version of that general underlying idea, e gives very strong (perhaps conclusive) support for the specific theory; however, *secondly*, there is no support from that evidence for the general, underlying theory itself.

A similar remark also applies to Colin Howson's statistical examples: so long as the basic theory or 'model' is given (basically in his urn case, that we are dealing with a 'Bernouilli process' with fixed, but unknown parameter p (the proportion of white balls in the urn)), then the evidence that k/n of the sampled balls were white gives support (in this case of course not conclusive) for the specific theory that estimates p as lying in the interval $k/n \pm \varepsilon$. But that data gives no conceivable reason for having greater faith in the idea that this is the correct model. (Indeed this is not an issue that would normally even arise in that case.)

Reverting to the cases of scientific theory, the intuitive judgment seems to be that in some but not all of the cases there is already good empirical reason to accept the general underlying theory concerned in these 'deductions from the phenomena'. There was already good reason to accept the general wave theory with the free parameter, ahead of any measurement of fringe distances with light from the sodium source. Hence, when evidence e turns out deductively to entail the specific theory T' (complete with filled-in value for the wavelength of light from the sodium arc) *given* T, we can intuitively " discharge the antecedent" and say that e gives us reason to accept T' full stop. In contrast, in the Gosse dodge case, exactly because there is no independent reason to accept the underlying general Creationist account, the fact that the fossil record entails the Gosse dodge variant of Creationism, justifies only the conditional judgment that e gives us reason to accept the Gosse dodge variant only to the extent that we already have reason to accept the general theory.

But how exactly can these general underlying theories earn their independent empirical support, as, in some cases it seems they do? Duhem, after all, taught us that these general theories do not have directly checkable empirical consequences of their own. All empirical tests of the wave theory of light, for example, are tests of the general wave theory *plus* particular assumptions. It seems, then, that if this whole approach is to be at all coherent, there must be a 'contrast class' to the sorts of cases we have considered so far. That is there must be empirical tests, the results of which not only confirm the specific version of the theory that entails their results, *but also* confirm the underlying general theory. Scientists do, in other words, sometimes take it that the empirical success of some particular version of a general theory gives good reason to accept the general theory itself - and in particular good reason to seek to develop another specific theory for a different field of phenomena based on that same general theory. So, for example, the empirical success of Fresnel's specific wave theory of diffraction was taken to provide good reason to develop another specific theory based on the same general elastic medium wave theory to deal with the phenomena of polarisation and crystal optics.

In sum, scientists do not restrict themselves simply to judgments of the conditional kind that we just highlighted – that against the given background of some general framework theory, some piece of evidence e gives strong support to some specific version of the general theory. They also sometimes see the general framework theory as empirically supported. Yet, as Duhem showed us, such support must always be achieved, not directly, but *via* specific versions of the general theory (i.e. not the general theory alone but that theory plus some further assumptions).[9]

What kinds of evidence perform this trick? The answer, I think, is two kinds, of which the more straightforward is the following. A scientist starts with some general theory T, uses e to fix some parameter in T, to create (by 'deduction from the phenomena') the more specific theory T'; T'then goes on to make some further independent prediction e'. So, to take again the classic case, Newtonians ..

[9] not holism – forward reference to Copernicus example.