1 The Contemporary Debate on Scientific Progress: What Constitutes Cognitive Progress?

> Although it is nearly uncontroversial that science makes progress of some sort or other, it is far from uncontroversial what scientific progress consists in. (Dellsén 2016: 72)

Prima facie, pre-philosophically, some developments in science have constituted intellectual advances, whereas others have not. The detection of gravitational waves, reported by Abbott et al. (2016), is a plausible example of an advance; most think that this was a 'step forward' for science qua an intellectual endeavour. The hundreds of papers on polywater in the 1970s seem instead to have been a 'step backwards' in the same dimension. Scientists now think that what was taken to be a polymer, composed of H_2O monomer units, was really something else.[1]

Some developments also appear to have been more intellectually significant than others. For instance, the discovery that polywater does not exist seems less important than the advent of Bohr's (1913) model of the atom, which gave a potential explanation of the absorption/emission spectra of hydrogen and provided a means to predict unknown spectral lines of several other elements, inter alia.[2] Thus it is natural to think that we could rank scientific changes in terms of how much of an improvement they constituted.

We might therefore adopt the following approach to exploring what kinds of changes are responsible for *cognitive progress* in science, as Laudan (1977) called it. We might trust that we can reliably identify instances when progress has occurred and that we can reliably rank them in order of significance. We might then generate lists of progressive episodes and orderings of the episodes with respect to the degree of improvement they constituted. Having done this, we might perform comparisons to determine what constitutes scientific progress. We might ask questions such as 'What are the similarities between the cases where cognitive progress occurred?' and 'What differentiates cases where substantial progress was made from situations where minimal progress was made?'

It is also natural to use thought experiments and counterfactual judgements to assist in the task. We might imagine what scientists could have done differently in a historical scenario. Or we might conceive of a hypothetical situation where

[1] Sweat contamination is commonly thought to be to blame, following Rousseau (1971). However, van Brakel (1993) argues that a reaction between water and silica was responsible.

[2] Key was Bohr's reduction of the Rydberg constant to an expression involving electron mass, electron charge, and several physical constants. For more on the background to and the significance of this episode, see Rowbottom (2019: ch. 4), Aaserud and Heilbron (2013), and Heilbron and Kuhn (1969).

scientists face a dilemma and consider which choice would result in more progress.

This is how the contemporary debate about cognitive scientific progress usually proceeds; evidence to this effect will become apparent. Questions such as 'Does science really cognitively progress?' and 'What is the source of the standards for judging what is cognitively progressive?' are only touched upon in modern philosophy of science. I believe that such meta-normative questions are pressing, for reasons that will emerge. Sections 2 and 3 of this Element will focus on them, in the hope of reorienting the ongoing discussion.

For the moment, I will say something in defence of the standard approach. All inquiries need starting points. One cannot have a view from nowhere, any more than one can have a God's-eye view. Furthermore, it is natural to begin with the assumption that something is present when it seems to be, and to see if you can learn more about it. Philosophers have done this with putative things as disparate as properties, time, knowledge, minds, and selves. And if philosophy cognitively progresses, one of the ways it does so is by charting the possibilities about such things.

In this section of the Element, therefore, I will cover the status quo on cognitive scientific progress. To prepare the stage for doing so, I will next do two things. First, I will explain and justify the precise scope of this survey. Second, I will present several key distinctions that will be useful throughout.

1.1 Survey Scope

This survey is limited in two respects. First, it concerns only literature on cognitive progress. You might wonder why this is. The short answer is that philosophers of science have focused on this kind of progress because it seems to be a distinctive feature of science. This is not to deny that other sorts of progress are possible in science or that they are worthy of philosophical study. For instance, science might progress or regress in a moral sense. But work on moral issues in science is typically labelled as 'ethics' or 'science studies' as a matter of convention. In the philosophical community, 'scientific progress' is customarily taken only to refer to a cognitive issue. To see this, consult the entry on 'scientific progress' in the *Stanford Encyclopedia of Philosophy* (Niiniluoto 2019) or the most recent article on the topic in *Philosophy Compass* (Dellsén 2018a).[3]

[3] I recommend both as further reading. Dellsén (2018a) gives a concise overview of the dominant theories of progress, and the main considerations for and against each; it is an excellent starting point, which gets to the core of the issues while minimising complications. Niiniluoto (2019), on the other hand, covers a wider time range and offers many excellent insights into the conceptual issues underlying the debate. Losee (2004) is also worth consulting to see how earlier accounts of what scientific progress consists in do not all fit into the classification system suggested by Bird (2007).

The exact scope of 'cognitive' is contested, however, which manifests itself in how different authors define the notion. Laudan (1977: 7) employs a broad definition, in taking cognitive progress to be 'nothing more nor less than *progress with respect to the intellectual aspirations of science*'. Niiniluoto (2019) instead introduces it as involving 'increase or advancement of scientific knowledge' or 'success in knowledge-seeking or truth-seeking'. And Dellsén (2018a) writes that 'cognitive scientific progress . . . has to do with improvement in our theories, hypotheses, or other representations of the world'. In effect, Niiniluoto and Dellsén build some of their own views on what constitutes cognitive progress into their initial characterisations of the cognitive domain. To remain more neutral on the matter, let us instead begin by assuming only Laudan's broader, less presumptuous, definition.

The second limitation of this survey is its historical scope. Cognitive scientific progress has been discussed for as long as anything resembling modern science has existed; even before the twentieth century, it was discussed by Bacon, Whewell, Mill, Mach, Duhem, and Poincaré. It has also been addressed by many scientists, as illustrated in Rowbottom (2019: ch. 4). In the past century, moreover, interest in the topic blossomed after the publication of three classics: Hanson's *Patterns of Discovery*, Popper's *Logic of Scientific Discovery*, and Kuhn's *Structure of Scientific Revolutions*. All leading philosophers of science since – such as Cartwright, Kitcher, Laudan, and van Fraassen – have had something to say about progress, even if they didn't label their work as being on the topic. But since a full survey of this literature is beyond the scope of this Element, I have elected to focus on the current debate. Bird (2007) fomented this, and the bulk of the subsequent literature criticises his position and argues for alternative views with reference to his.

1.1.1 Key Preparatory Distinctions

Three preparatory distinctions will be useful. First, a complete answer to 'What constitutes scientific progress?' would concern two interconnected matters. On the one hand, it would specify which aspects of science are pertinent when evaluating cognitive 'goodness' in science. Should we limit our attention to theories? Or should we also consider existential claims? And should we be concerned ultimately only with beliefs concerning such things? What about other components of science, such as values, methods, exemplars, models, explanations, and instruments? In short, what are cognitive *goodness bearers*? On the other hand, it would specify on what basis we should compare those items to ascertain whether, or how much, progress has occurred. Imagine, for instance, that all goodness bearers are theories. Candidate cognitive *goodness*

makers are potential properties of those goodness bearers, such as truth, approximate truth, empirical adequacy, simplicity, scope, and being known. Naturally, what goodness bearers are constrains what goodness makers can be, and vice versa. If truth is the sole goodness maker, for example, then methods are not goodness bearers (because they do not have truth values).

Second, note the distinction between *monistic* and *pluralistic* accounts of scientific progress. Monistic accounts hold that there is just one kind of goodness maker (although they often allow that there is more than one kind of goodness bearer). Pluralistic accounts hold that there are various kinds of goodness maker (but might posit only one type of goodness bearer). For instance, a pluralist might hold that theories are the only goodness bearers, but that their 'cognitive goodness' involves several dimensions, such as accuracy, simplicity, and scope.[4] Some accounts are *more pluralistic* than others, moreover, in positing more kinds of goodness makers than others. It is imperative to understand, however, that a pluralistic account may be *hierarchical*, nevertheless, in the sense that it may rank some goodness makers as more significant (or 'core') than others. For example, a pluralist might hold that increases in truthlikeness always bring more progress than increases in simplicity, while accepting that progress can occur either way.

Third and finally, heed the distinction between changes that constitute progress and changes that promote progress. The difference is sometimes easy to see. For example, drastically increasing universities' research funding would probably result in cognitive progress but would not itself constitute a cognitive 'good'. Part of what makes this obvious, though, is that increasing research funding is not a change of a cognitive or intellectual kind. When we instead consider changes internal to science which feature centrally in what scientists do – the development of instruments such as the scanning electron microscope (SEM) and lab techniques such as gene splicing, for example – it is harder to determine what we should say. A complicating factor is the possibility, especially from a pluralistic perspective, that some kinds of change might simultaneously promote progress and constitute progress. Consider again the development of the SEM. One might argue that this constituted cognitive progress by enabling us to do new things and promoted progress by leading to the discovery of new truths (about things investigated using the SEM).

Having presented these distinctions, I will now turn my attention to the dominant extant accounts of scientific goodness makers and goodness bearers, with a special focus on the most prominent defenders of each. In doing so, I will follow the thread of the debate initiated by Bird (2007) in a broadly

[4] Each of these items appears on the list of theoretical virtues proposed by Kuhn (1977: ch. 13).

chronological fashion. In the interests of economy, I will rarely hereafter use 'cognitive' to delimit the scope of the claims about progress.

I will cover the four dominant monistic (or near-monistic) accounts: epistemic (knowledge-based), semantic (truth-based), functionalist (problem-based), and noetic (understanding-based). Each has different variants. As we will see, pluralistic accounts tend to combine elements of each.

1.2 Bird's Epistemic Account: Progress As Increasing Knowledge

Bird (2007) reinvigorated the debate on scientific progress by championing a simple monistic view thereof: scientific progress occurs if and only if scientific knowledge increases. As Bird concedes, the idea is not novel. For instance, as Charkravartty (2017) notes, realists typically 'regard theories as furnishing *knowledge* of both observables and unobservables [and] arguably the most important strains of antirealism have been varieties of empiricism which, given their emphasis on experience as a source and subject matter of *knowledge*, are naturally set against the idea of *knowledge* of unobservables' [my emphasis]. That is to say, realists and anti-realists disagree about the *scope* of the knowledge science provides. But both sides typically agree that science provides some knowledge. It is therefore natural, prima facie, for both sides to accept that science makes cognitive progress when scientific knowledge increases.

However, Bird's epistemic account is bold because of its monistic character; no one else has defended such an epistemic account in detail.[5] This is partly why it is so interesting. It is also no accident that Bird's account came in the wake of Williamson's (2000) knowledge-first approach to epistemology, with which it dovetails. According to this approach, which is defended by Kelp (2014), finding knowledge is the constitutive aim of inquiry. And to this, to reach a position approximating Bird's epistemic account, one need only add that: (a) science is a form of inquiry; and (b) to achieve (or approach achievement of) the aim of an activity is the only way to make progress therein. As Bird (2007: 65) puts it: 'Given that science is an epistemic activity it seems almost tautologous to suggest that its success and so progress should be measured by epistemic standards.' In due course, however, we will see that many philosophers of science have resisted this seemingly 'almost tautologous' claim.

Bird (2007) proceeds by comparing his epistemic account with the most prominent prior accounts of progress. I will now present these comparisons and summarise the subsequent exchanges concerning these in the literature.

[5] There are, however, precursors that Bird (2007) does not cite. Cohen (1980: 491), for example, argued: 'in science . . . the objective is not truth but knowledge'. See also Barnes (1991).

I will begin by using Bird's (2007) sketches of alternative accounts of progress. Some respects in which these are inadequate will emerge subsequently.

1.2.1 Bird on the Epistemic View versus the Semantic View

Bird (2007: 65) initially attacks the semantic account of progress, which he characterises as 'progress is the accumulation of true scientific beliefs [or] a matter of increasing verisimilitude . . . or nearness to the truth'. His concern is that this semantic view fails to save our intuitions concerning episodes where luck is involved in new true beliefs arising or in the verisimilitude of scientific beliefs increasing.

Bird uses one highly abstract thought experiment, and another which appeals to a historical episode. The former involves a 'scientific community that has formed its beliefs using some very weak or even irrational method M, such as astrology' (Bird 2007: 66). A natural reaction to such a thought experiment, however, is that a community fails to be scientific if its 'true beliefs are believed solely because they are generated by M' (66).[6] The subsequent literature has therefore focused on his other thought experiment, which involves a real scientific community. Here it is:

> Réné Blondlot believed in the existence of what he called N-rays for what it is clear were entirely spurious and irrational reasons. Imagine for sake of argument that we were to discover that there are in fact hitherto unobserved entities answering to Blondlot's description of N-rays. So Blondlot's belief in N-rays would have been true but unjustified and not knowledge. The semantic approach would have to regard Blondlot's belief (which was widely shared in France) as constituting progress. That is clearly wrong. (Bird 2007: 67)

However, Rowbottom (2008) argues that matters are not so clear, because the history of science shows that Blondlot's (and the community's) belief in N-rays was based on several other false beliefs.[7] Hence, he argues that the episode might be construed as *regressive overall* on the semantic view, even if N-rays exist, provided the introduction of those numerous false beliefs was detrimental enough to outweigh any improvement derived from the introduction of some true beliefs about N-rays. To hold this, one need only reject the crude notion that progress *only consists in the accumulation of true beliefs*. One should grant also

[6] For instance, Niiniluoto (2017) writes: 'the primary application of the notion of scientific progress concerns successive theories which have been accepted by the scientific community. Some sort of tentative justification for such theories is presupposed even by a radical fallibilist like Popper . . . Irrational beliefs and beliefs without any justification simply do not belong to the scope of *scientific* progress.' Rowbottom (2015a) and Dellsén (2021), on the other hand, argue that such 'tentative justification' falls short of the kind of justification required for knowledge.

[7] Rowbottom (2008) cites Lagemann (1977) and Nye (1980). The history also casts some doubt on whether the reasons were 'entirely spurious and irrational'.

that *avoiding false beliefs* is significant. This idea is old. One finds it, for example, in James (1896: §vii):

> There are two ways of looking at our duty in the matter of opinion, – ways entirely different . . . *We must know the truth*; and *we must avoid error* . . . they are two separable laws.

James (1896) notes also that eliminating a false belief need not involve substituting it with another belief. His discussion is richly suggestive of the consequence that, ceteris paribus, removing a false belief is progressive and introducing a false belief is regressive when it comes to doing 'our [epistemic] duty'. No modern defender of the semantic view disagrees with this sentiment (although not all think in terms of duties). And as we will see, Cevolani and Tambolo (2013) show that this consequence is 'built in' to the verisimilitude-based variant of the semantic view.[8]

Rowbottom (2008) continues by presenting several variants of the N-rays-based thought experiment, aimed at showing the significance of the value problem in epistemology. Broadly, the value problem concerns a question which Plato tackles in *Meno*: 'Why is knowledge more valuable than belief?' Rowbottom (2008: 278) argues that appeal to justification does not suffice to answer satisfactorily; he asks, rhetorically: 'Wouldn't we be better off having an unjustified true belief that "N-rays exist" and being neutral concerning whether there is any evidence to that effect rather than having a justified true belief that "N-rays exist" based on justified false beliefs in a great deal of evidence to that effect?' Rowbottom (2008) continues by suggesting that justification is only *instrumental* in achieving progress and hinting that knowledge may be in a similar boat.

Bird (2008) responds to Rowbottom's (2008) point about false beliefs as follows. First, he makes an empirical claim about what drives our response to his thought experiment: 'It takes a certain amount of reflection to see that Blondlot must have had some false beliefs, even if his theory is true . . . I do not think that this reflection is in fact present when we carry out the thought-experiment' (Bird 2008: 208). However, Rowbottom (2008) did not deny this claim; rather, he suggested that the thought experiment appears to tell against the semantic view only if one mistakenly thinks that the semantic view is committed to the claim that the N-ray episode was progressive overall.[9] Second, Bird (2008: 280) claims that the semantic view of progress doesn't entail anything about false beliefs, and therefore that Rowbottom is proposing a revision to it:

[8] Indeed, Bird (2007: 85) himself notes that 'the truth view of the aim of belief is typically modified, so that the aim of belief is characterised as the complex aim of achieving truth subject to the proviso that falsity is always avoided'.

[9] Plausibly, we should also not be interested in snap judgements. We should reflect on our responses to thought experiments, formulate arguments as a result, and reach judgements on that basis. See, for instance, Deutsch (2015).

> to accommodate Rowbottom's response . . . the defender of the semantic view will have to make changes to (S) [the semantic view]. This is because (S) says nothing about false belief; it does not yet say that an episode is not progressive if it involves a considerable quantity of false beliefs. All it says is that there is progress when there is an increase in true beliefs – which there is in my hypothetical example.

Bird continues by suggesting that such revisions would be problematic, and that the epistemic view requires no such machinations: 'one has to complicate matters by saying that belief aims at truth and at avoiding falsehood . . . The epistemic view of the aim of belief and of progress can avoid all this' (280).

As Cevolani and Tambolo (2013) point out, however, it follows from Bird's (2007: 65) own definition of the semantic view, on which progress may be 'a matter of *increasing verisimilitude . . . or nearness to the truth*' [my emphasis], that the elimination of false beliefs may result in progress (and that the introduction of false beliefs may result in regress). Verisimilitude 'represents the idea of approaching comprehensive truth' (Popper 1963: 237). To illustrate, imagine a future in which we have the comprehensive scientific truth. We have all the true claims in science's domain of inquiry: truths about the fundamental constituents of the world, about the laws governing the behaviour of those things, and so forth. Now consider how the degree of verisimilitude of our science could decrease. We might forget some truths. Or we might replace some true claims with false ones. And the latter kind of regressive change might sometimes be worse than the former kind. Avoiding falsehood is important. Thus, no revision is required to the semantic view for Rowbottom's (2008) argument to go through. To return to Cevolani and Tambolo (2013: 925):

> verisimilitude is a 'mixture' of two ingredients, truth and content. If truth were the only ingredient, then all truths, including the tautology, would be equally (and maximally) verisimilar; and, vice versa, if only content were relevant, then a plain contradiction would be closer to the truth than any other theory. Thus, devising highly verisimilar theories is a 'game of excluding falsity and preserving truth'. [quoting Niiniluoto (1999: 73)]

So the semantic view is pluralistic to a small extent; it admits two dimensions of progress. Hence, Cevolani and Tambolo (2013: 930) opine that 'Bird's attack . . . is apparently based on a misunderstanding . . . First of all, verisimilitude should not be conflated with approximate truth; and secondly, the accumulation of (approximate) truths does not guarantee increasing verisimilitude.'[10]

[10] On the second issue, Cevolani and Tambolo (2013: 931) explain that 'accumulation of (approximate) truths is neither a necessary nor a sufficient condition for increasing verisimilitude'. It is not necessary because abandoning a false claim could serve that end. It is not sufficient because adding a true (or approximately true) statement to a false hypothesis can result in a less

Cevolani and Tambolo (2013) add that Bird's N-ray thought experiment involves a scenario where *estimated* progress and real progress come apart, and that such cases are accepted as commonplace by advocates of the verisimilitude-based view of progress. Niiniluoto (2014) agrees with this verdict, labelling Bird's account of the semantic view 'incomplete and misleading', and adds that estimates of progress must be justified. It is therefore reasonable to conclude that Bird's (2007) thought experiment only refutes a simplistic variant of the semantic approach – dubious since James (1896) – which takes progress to occur only via accumulation of true or approximately true statements.

This brings us to Bird's (2008) response to Rowbottom's (2008) suggestion that justification is only instrumental for progress. Bird agrees that adding justification has a negative effect on progress in some contexts, but explains that this need not go for knowledge, provided one adopts an appropriate account of knowledge, such as Williamson's (2000). Bird (2008: 280) writes that the epistemic view 'says knowledge constitutes progress, and nothing short of knowledge. It does not imply that justification constitutes progress (or some weaker progress-like good); even less does (E) imply that justification (or even knowledge) will cause future progress'. Bird adds that Williamson's view of knowledge nonetheless entails that justification is not merely instrumental in value, because justification is necessary for knowledge. In short, knowledge entails justification (and truth) on Williamson's (2000) view.

Rowbottom should have anticipated this reply, because Bird (2007: 72) states that truth and justification are not, on his view, 'jointly sufficient for a new scientific belief adding to progress'. Rowbottom (2010) acknowledges this, although, as we will see, he uses this relationship to underpin a different attack on Bird's epistemic view. However, Rowbottom (2010) also points out that an epistemic view need not suppose a Williamsonian view of knowledge. Like the semantic view, the epistemic view has different variants. Justification's instrumental value tells against some such accounts.

In any event, Bird (2007) does not rely entirely on thought experiments. He argues separately that the semantic account has no advantages over the epistemic account. Most notably, he writes:

> The notion of verisimilitude lacks a worthwhile characterization in place of a definition. It is in less general use than the concept of knowledge. It is not

verisimilar hypothesis. Their example, which is artificial but nonetheless illustrative, is as follows. Imagine we have the false theory, *T*, that Mont Blanc's height is either 1,000 or 4,809 metres. If we now discover the true claim that 'Mont Blanc's height is either 1,000 or 4,810 metres' and we add this to *T*, we conclude that 'Mont Blanc's height is 1,000 metres'. We have been led away from *the truth* by discovering *a true claim*. The underlying issue here is that elimination of error, discussed previously, is important.

> obviously explanatorily significant. And, most importantly, it is difficult to see how its use can be helpfully extended beyond the simple cases we do apply it to. (Bird 2007: 75)

No one has responded to the claim that verisimilitude is in less widespread use than knowledge and is of dubious explanatory significance. But a possible response is as follows: (1) even if a knowledge-based account of progress would be easier for non-specialists to grasp and 'get behind', this does not indicate it is correct; and (2) an adequate epistemic account will be more complex than first appearances suggest, because some folk conceptions of knowledge are unfit for the task, as Bird concedes in ruling out justified-true-belief accounts. This is just a programmatic sketch of a response, however; this aspect of Bird's argument deserves more attention than I can give it.

This brings us to Bird's claim that it is difficult to extend the verisimilitude account beyond simple cases (without it amounting to a cumulative view). He takes the most significant problem to be as follows:

> Let it be that a science adds to the set of its generally accepted beliefs just one new belief that is . . . closer to the truth. There is now not even an intuitive sense in which the science as a whole is now closer to the truth than it was – unless that sense is identical to the thought that this science includes *more* (approximate) truth. (Bird 2007: 75).

However, as we have already seen, advocates of the verisimilitude view of progress are not committed to the view that progress always occurs in such circumstances. Adding new approximate truths – which most consider to be falsehoods – is only a potential way to increase verisimilitude.[11] So there is no problem here for Bird's target; any appearance to the contrary stems from conflating approximate truth with verisimilitude. To reiterate, Cevolani and Tambolo (2013: 933) warn that 'no simple principle of the form "add to T a true, approximately true, or verisimilar, belief" can guarantee truth approximation through belief change . . . [Although] if T' is obtained from a true theory T by adding a new truth to it, then T' will be more verisimilar than T'. See Footnote 10 and Niiniluoto (1999: 201–3) for more detailed discussion of this point.

1.2.2 Bird on the Epistemic View versus the Functionalist-Internalist View

Bird (2007: 67) also launches an attack on 'the functionalist-internalist' view of scientific progress, which he associates with Kuhn, Laudan, and, to a lesser extent, Lakatos (1978). He calls it functionalist because it takes progress to occur when science performs the function of 'solving problems

[11] For an alternative view, though, see Rowbottom (2022).