

The Crux of Crucial Experiments: Duhem's Problems and Inference to the Best Explanation

Marcel Weber

ABSTRACT

Going back at least to Duhem, there is a tradition of thinking that crucial experiments are impossible in science. I analyse Duhem's arguments and show that they are based on the excessively strong assumption that only *deductive* reasoning is permissible in experimental science. This opens the possibility that some principle of *inductive* inference could provide a sufficient reason for preferring one among a group of hypotheses on the basis of an appropriately controlled experiment. To be sure, there are analogues to Duhem's problems that pertain to inductive inference. Using a famous experiment from the history of molecular biology as an example, I show that an experimentalist version of inference to the best explanation (IBE) does a better job in handling these problems than other accounts of scientific inference. Furthermore, I introduce a concept of experimental mechanism and show that it can guide inferences from data within an IBE-based framework for induction.

- 1 *Introduction*
- 2 *Duhem on the Logic of Crucial Experiments*
- 3 *'The Most Beautiful Experiment in Biology'*
- 4 *Why Not Simple Elimination?*
- 5 *Severe Testing*
- 6 *An Experimentalist Version of IBE*
 - 6.1 *Physiological and experimental mechanisms*
 - 6.2 *Explaining the data*
 - 6.3 *IBE and the problem of untested auxiliaries*
 - 6.4 *IBE-turtles all the way down*
- 7 *Van Fraassen's 'Bad Lot' Argument*
- 8 *IBE and Bayesianism*
- 9 *Conclusions*

1 Introduction

Some of the major discoveries in the history of molecular biology are associated with an alleged ‘crucial experiment’ that is thought to have provided decisive evidence for one among a group of hypotheses. A well-known example is the Hershey–Chase experiment (1952), which showed that viral DNA, not protein, enters a bacterial cell to reprogram it to make virus particles. Another example is the ‘PaJaMo’ experiment (1958), which showed that a certain bacterial gene produces a substance that represses the activity of other genes. In both cases, there were two major hypotheses that could explain the facts known beforehand: Either viral protein or viral DNA contains the information for making new virus particles (Hershey–Chase). Similarly, either ‘generalized induction’ (in the molecular biological, not logical sense!) or suppression of a repressor (the ‘double bluff’ theory of Leo Szilard) was thought to be responsible for the regulation of sugar metabolism in bacteria (PaJaMo). Examples such as these abound in experimental biology (see Weber [2005], Chapters 3–5). In many cases, a single experiment seems to have enabled a choice between the competing hypotheses at hand, thus strongly resembling Bacon’s ‘instances of the fingerpost’ or Newton’s ‘experimentum crucis’.

Philosophers of science, of course, have been less than enthusiastic about the possibility of crucial experiments.¹ Following Duhem ([1954]), many seem to think that a single experiment, as a matter of principle, is not able to choose among a group of hypotheses. However, as I will show, Duhem made extremely strong assumptions concerning the kind of inferences that are to be permitted. Namely, he allowed only *deductive* inferences to be used. In this paper, I will show that when crucial experiments are construed along the lines of *inductive* (ampliative) inference, Duhem’s arguments become less persuasive. Even though there are analogues to Duhem’s problems in the realm of inductive inference, these are solvable within the framework of a theory of induction based on inference to the best explanation.

I want to demonstrate the possibility of crucial experiments on a concrete historical example from molecular biology, namely the Meselson–Stahl experiment done in 1957. Even though there is an extremely detailed historical study of this experiment available (Holmes [2001]), it has to my knowledge never been subjected to a thorough methodological analysis.² ‘The most beautiful experiment in biology,’ as it has been called, is widely thought to have demonstrated semi-conservative replication of DNA as predicted by Watson and Crick in

¹ An exception is Franklin ([2007]).

² Brief methodological discussions of the case can be found in (Franklin [2007], Section E.2) and in (Roush [2005], pp. 14–6). Franklin uses it to highlight the importance of intervention and experimental control. Roush uses the episode to illustrate her tracking account of evidence.

1953. But it remains to be shown that this experiment was actually decisive from a methodological point of view.

In Section 2, I will discuss Duhem's infamous arguments against crucial experiments. Section 3 provides a brief account of the Meselson–Stahl experiment and some of the theoretical controversies that preceded it. In Section 4, I show that the evidential import of this experiment cannot be accounted for by a simple elimination scheme. In Section 5, I argue that the experiment cannot be viewed as a severe test as prescribed by the error-statistical approach to scientific inference. In Section 6, I propose an experimentalist version of inference to the best explanation (IBE) and show that it provides the most adequate reconstruction of the experiment as providing strong evidence for the semi-conservative hypothesis. My account is based on the idea that explanations in experimental biology often involve more or less detailed descriptions of *mechanisms*, which is substantiated by much recent scholarship (e.g., Machamer, Darden, and Craver [2000]; Bechtel [2005]; Darden [2006]; Craver [2007]). I add to these accounts the concept of an *experimental mechanism* and analyse the role of such mechanisms in making inferences from data (Sections 6.1 and 6.2).

The main reason why IBE provides the best account of the Meselson–Stahl experiment, I will argue, is that it does a better job in dealing with the inductive analogues of Duhem's problems. One is the problem of untested auxiliaries, treated in Sections 6.3 and 6.4, while the other is very similar to van Fraassen's 'bad lot' objection to IBE, discussed in Section 7. In Section 8, I briefly discuss the relationship of my IBE-based account to Bayesian confirmation theory.

2 Duhem on the Logic of Crucial Experiments

Duhem characterized crucial experiments as follows:

Do you wish to obtain from a group of phenomena a theoretically certain and indisputable explanation? Enumerate all the hypotheses that can be made to account for this group of phenomena; then, by experimental contradiction eliminate all except one; the latter will no longer be a hypothesis, but will become a certainty (Duhem [1954], p. 188).

This passage strongly suggests that Duhem thought of crucial experiments in terms of *eliminative induction*, in other words, in terms of the following logical scheme³:

1. $H_1 \vee H_2$
2. $H_1 \Rightarrow e$

³ I reconstruct the logical scheme for two hypotheses. It is obvious how it could be expanded for more than two.

3. $H_2 \Rightarrow \neg e$
4. e
5. From (3), (4): $\neg H_2$ [by *modus tollens*]
6. From (1), (5): H_1 [by disjunctive syllogism]

Such a train of inference faces two major problems according to Duhem. The first problem is the one that is today known as ‘Duhem’s problem’. This is the problem that auxiliary assumptions are needed to secure the deductive relation between hypothesis and evidence. Therefore, (5) will never involve a hypothesis alone; it will always be a conjunction of hypotheses that can be said to be falsified. Famously:

The only thing the experiment teaches us is that among the propositions used to predict the phenomenon and to establish whether it would be produced, there is at least one error; but where this error lies is just what it does not tell us (*ibid.*, p. 185).

But if the falsity of one of the hypotheses at issue cannot be asserted, the inference (6) does not go through. As if this weren’t enough, Duhem identifies a second problem:

Between two contradictory theorems of geometry there is no room for a third judgment; if one is false, the other is necessarily true. Do two hypotheses in physics ever constitute such a strict dilemma? Shall we ever dare to assert that no other hypothesis is imaginable? Light may be a swarm of projectiles, or it may a vibratory motion whose waves are propagated in a medium; is it forbidden to be anything else at all? (*ibid.*, p. 190).

The answer to the latter, rather rhetorical question is clear: Unlike mathematicians, physicists can never have grounds for assuming that they have exhausted the space of possible truths. In other words, there can be no warrant for a premise such as (1) in the scheme above.

Given what he sets out to prove, Duhem’s arguments are impeccable. But note that Duhem is clearly thinking in terms of *deductive* inference. What he proves is that experiments conjoined with deductive logic, together, are unable to bring about a decision for one among a group of hypotheses. Of course, he is absolutely right about that. However, Duhem’s arguments do not touch the possibility of *inductive* or *ampliative* inference enabling such a choice.⁴ An

⁴ Of course, the possibility of inductive inferences is not something that Duhem simply overlooked; he provided elaborate arguments against inductivism (mostly using Newtonian mechanics as an example). I lack the space to discuss these here.

ampliative inference rule might very well be able to mark one hypothesis as the preferable one.⁵

This proposal raises the question if such a procedure does not run into similar difficulties. It would seem that Duhem's first problem concerns only the possibility of *refuting* hypotheses. On the account that I shall give, crucial experiments do not refute the alternatives. Instead, they *positively* select one of the hypotheses as best supported by the evidence. Therefore, Duhem's first problem in its classic form seems to be irrelevant. Even so, it is clearly the case that a crucial experiment relies on *auxiliary assumptions*. If these are false, there can be no evidential support. False auxiliaries could mask the truth or the falsity of a hypothesis under test. While it might be possible to independently test some auxiliaries (see Section 6.4), it is never possible to test all of them. The reason is that each attempted test of an auxiliary assumption will require *further* assumptions, and so on. If we require that all auxiliaries be tested, there will never be any conclusive evidential support from an experiment. This is the analogue of Duhem's first problem that arises within my framework. I shall refer to it as the 'problem of untested auxiliaries'. In Section 6.3, I will show how the problem can be solved within an IBE-based framework.

As for Duhem's second problem, it is no less relevant for inductive inference as it is for Duhem's deductive inference scheme, at least if inductive inference is to be truth-tropic. When scientists hold a hypothesis to be true on the grounds that it is the one from a group that is best supported by the evidence, they must have grounds for claiming that there are no better ones that they have not considered. If no such grounds can be had, then an inductive inference regime runs into exactly the same problem as Duhem's eliminative scheme. This, put into the context of a specific inductive principle, namely IBE, is Bas van Fraassen's 'bad lot'-argument.⁶ I shall deal with it in Section 7.

As this discussion reveals, we can expect an inductive selection regime for hypotheses to run into an analogue of Duhem's first problem, and another problem that is basically Duhem's second. I shall address these problems in due course. But right now, it is time to introduce my historical example.

3 'The Most Beautiful Experiment in Biology'

As is well known, James D. Watson and Francis H. C. Crick closed their landmark paper on the structure of DNA with the short and crisp remark 'It

⁵ A similar claim can be found in (Laudan [1990]). Duhem-type problems (and their Quinean relatives) are often discussed under the rubric of 'underdetermination of theory by the evidence.' However, this expression is ambiguous, as Laudan shows. This is why I prefer to develop the problem in terms of a reply to Duhem's classic arguments.

⁶ Van Fraassen discussed this problem specifically as a part of his argument against IBE. Stanford ([2006]) offers a more general and systematic discussion of what he calls 'the problem of unconceived alternatives' and its epistemological ramifications.

has not escaped our notice that the specific base pairing we have postulated immediately suggests a possible copying mechanism for the genetic material' (Watson and Crick [1953]). It is fairly obvious what Watson and Crick had in mind: Because of the complementarity of the base sequences of the two nucleotide chains in the double helix, a DNA molecule could be copied by first separating the two strands, and then synthesizing two new strands using the two old strands as templates. On this scheme, each newly synthesized DNA molecule will contain one strand that was already present in the parental molecule, and one newly made strand. This scheme is called 'semi-conservative replication.' However, as plausible as this scheme might seem, sceptics were quick to notice some theoretical difficulties. Here is the greatest of them all, Max Delbrück:

I am willing to bet that the complementarity idea is correct, on the basis of the base analysis data and because of the implication regarding replication. Further, I am willing to bet that the plectonemic coiling of the chains in your structure is radically wrong, because (1) the difficulties of untangling the chains do seem, after all, insuperable to me. (2) The X-ray data suggest only coiling but not specifically your kind of coiling (Delbrück to Watson, 12 May 1953, quoted from Holmes [2001], pp. 21–2).

The term 'plectonemic' referred to the topological property that, according to Watson and Crick, two DNA strands are twisted about each other so that they cannot be separated without uncoiling. The 'base analysis data' refer to the work of Erwin Chargaff, who had shown previously that the building blocks of DNA occur in certain fixed ratios. Delbrück is also pointing out that the double helix was, at the time when Watson and Crick proposed it, strongly underdetermined by the available X-ray diffraction data (i.e., other coiled structures would have been consistent with these data).

But Delbrück not only expressed scepticism about the specific kind of coiling. His point (1) also called into question the whole idea of a semi-conservative replication mechanism as suggested by Watson and Crick. The problem was that, given the plectonemic topology of the double helix, untangling the two strands requires the breaking and rejoining of the sugar–phosphate backbone of the molecule. Given the fast rate by which DNA replicates, especially in rapidly dividing bacterial cells, the molecule would have to rotate at mind-boggling velocities.⁷ This was also known as the 'problem of untwiddling'. For a while, it was a major source of scepticism about Watson and Crick's extremely elegant solution. While the structure itself became rapidly accepted thanks to

⁷ Today, it is known that this is actually what happens. There is a whole class of enzymes called topoisomerases that control the coiling of the DNA molecule. These enzymes can catalyze extremely fast breaking and re-joining of the sugar–phosphate backbone of DNA. Some enzymes can even introduce rotational strain into the molecule under the expenditure of metabolic energy.




Figure 1. Delbrück's dispersive mechanism (Delbrück [1954], p. 786). The solid lines represent parental DNA strands, the dotted lines newly synthesized material. Reprinted with publisher's permission.

the available of improved X-ray data, the semi-conservative replication mechanism continued to be doubtful for the years to come.

In the years following Watson and Crick's announcement, two alternative replication mechanisms were proposed. Delbrück ([1954]) devised a scheme under which each newly synthesized DNA molecule contains bits of the parental molecule that are interspersed with newly synthesized material (Figure 1). This became known as the *dispersive* mechanism.

Gunther Stent ([1958]) proposed that the whole double-stranded DNA molecule could serve as the template for synthesizing a copy (Figure 2). This would not require any untwisting of the parental molecule.

According to this mechanism, which was called the *conservative* mechanism, the parental molecule emerges unchanged from the replication process while the newly synthesized molecules contain only new material. The three mechanisms differ with respect to the distribution of parental and newly synthesized material that end up in the daughter molecules. Thus, in the mid-1950s there were three different hypotheses concerning the distribution of parental and newly synthesized nucleic acid chains.

Now enter two young experimentalists, Matthew Meselson and Frank Stahl, working at the California Institute of Technology in Pasadena. Using a powerful analytic ultracentrifuge, they performed a remarkable experiment in 1957.⁸

⁸ See (Holmes [2001]) for an extremely detailed account. As usual, this experiment was preceded by a long and painstaking series of failures and cul-de-sacs. Holmes, who had the complete lab



Figure 2. Stent's conservative hypothesis (Stent [1958], p. 137), showing a new DNA strand being synthesized in the major groove of the parental double helix. Reprinted with publisher's permission.

Meselson and Stahl grew *E. coli* bacteria in the presence of a heavy isotope of nitrogen, nitrogen-15. Ordinarily, DNA contains the most common isotope of nitrogen, which is nitrogen-14. But when grown in the presence of nitrogen-15, the bacteria incorporate the heavy nitrogen into their DNA. Now, DNA that contains the ordinary, light nitrogen atoms and the DNA containing heavy nitrogen can be distinguished by their weight. Of course, DNA does not occur in large enough quantities to be weighed by an ordinary balance. But Meselson and Stahl developed a highly precise instrument for determining the weight of DNA. They first dissolved the bacterial cells in a strong detergent. Then they placed the extract in a very dense solution of the salt CsCl. When a CsCl solution is centrifuged at very high speed in an ultracentrifuge for many hours, it will form a density gradient after a while. At equilibrium, the DNA molecules will float in that region of the gradient that corresponds to their own density. They form a band that can be observed with the help of UV light. Thus, the weight of the DNA molecules can be measured by determining the position of the band.

The experiment that Meselson and Stahl now did was to transfer the bacteria from a medium containing heavy nitrogen to a medium containing light nitrogen and allowing the bacteria to multiply further. At regular time intervals

records available and conducted extensive interviews with the two scientists, traces the progress of Meselson's and Stahl's work on a day-to-day basis.




Figure 3. UV absorption photographs (left) and densitometric scans (right) of the ultracentrifuge cell (Meselson and Stahl [1958], p. 675). The bands show where the DNA floats in the CsCl density gradient. What is particularly important about these data is that the band of intermediate density was located exactly in between the heavy and light bands. As both theoretical calculations and measurements showed, the density gradient was very nearly linear in the range where the DNA was floating (see Section 6.4). This allowed the inference that the intermediate band contained molecules that were composed of heavy and light nitrogen exactly in a 1:1 ratio, as predicted by the semi-conservative hypothesis. Reprinted with publisher's permission.

after the transfer, they took samples and placed them in the ultracentrifuge. What they observed is that after one generation, a band of intermediate density appeared. After another generation, the intermediate band was still present, but a new band that corresponded to light DNA appeared (Figure 3). An obvious interpretation of this pattern was that the band of intermediate density consisted of hybrid molecules composed of one heavy and one light strand (see Figures 4 and 5). Such a hybrid could obviously have been produced by the semi-conservative scheme, according to which each newly produced double helix preserves one strand from the parental molecule. In contrast, the conservative mechanism should not produce a band of intermediate density (but see Section 4). According to the dispersive mechanism, the result would look the same after one generation, but the band should shift further after subsequent rounds of replication, as the molecules would lose the heavy nitrogen bit by bit over the generations. But there were no such shifts.

While it might seem obvious that these data supported the semi-conservative scheme best, there are methodological complications that I will discuss in the following sections. At any rate, the impact of this experiment on the scientific

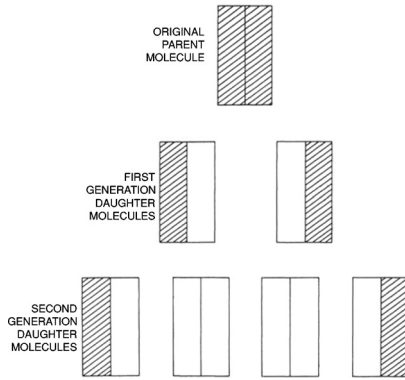


Figure 4. What the data showed according to Meselson and Stahl ([1958], p. 677). The shaded areas represent heavy nitrogen (^{15}N). Reprinted with publisher's permission.

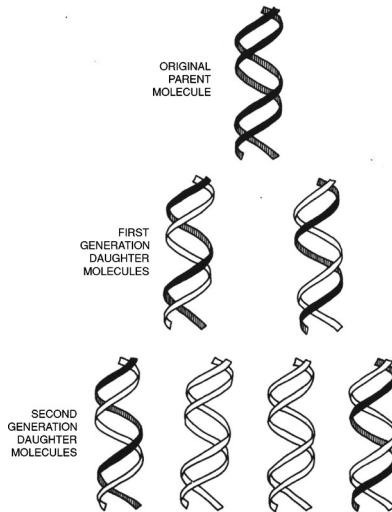


Figure 5. The most obvious interpretation of why the nitrogen was distributed as shown in Figure 4 (Meselson and Stahl [1958], p. 678). Reprinted with publisher's permission.

community at that time was considerable. Almost everyone agreed that the Meselson–Stahl experiment beautifully demonstrates semi-conservative replication. The only exception known to me is Max Delbrück, but his role in the closely knit molecular biology of that time seems to have been that of *advocatus diaboli*.

In the following sections, I shall provide a methodological analysis of this experiment and its evidential support for Watson and Crick's semi-conservative mechanism. It will turn out that the case is much more complex than it might seem at first sight.

4 Why Not Simple Elimination?

The question that I shall address first is why we cannot simply say that the semi-conservative scheme was the only one to survive the test to which Meselson and Stahl subjected it, while the alternative schemes were falsified and therefore eliminated. After all, the semi-conservative model predicted the outcome for this experiment correctly, while the two alternatives did not. The first thing to note is that this would amount to an eliminative induction, which is exactly the kind of reasoning that is not possible according to Duhem (see Section 2). If we construe the experimental reasoning like this, both of Duhem's objections can be raised. Here, I shall concentrate on Duhem's first problem.

In a Duhemian frame of mind, it could be argued that the dispersive and conservative hypotheses could still have been true because one or several auxiliary assumptions might have been false. For example, it could be that Meselson and Stahl were wrong about the molecular units that they resolved in their ultracentrifuge. Technically, what the centrifuge data showed is merely that there are three colloidal substances of different density. It does not show that these substances were simple DNA duplexes. In other words, the identification of the pattern shown in Figure 4 with the molecular entities shown in Figure 5 was a *theoretical interpretation* of the data.

This interpretation could have been false. The problem is especially acute because it is known today that Meselson and Stahl were mistaken about the length of the molecules they saw floating in their gradients. The hypodermic needles that Meselson and Stahl used to load the DNA onto the gradient must have mechanically sheared the DNA molecules into much smaller pieces—unbeknownst to these scientists in 1957.⁹ This did not alter the result because the CsCl-gradient technique separates DNA molecules according to density, not length. But this does show that there were severely mistaken assumptions about the experimental system. Meselson and Stahl were lucky that their mistake concerning length was not relevant. But what guaranteed that all the *salient* auxiliary assumptions were correct?

As these historical facts make clear, there would have been ample reason for defenders of the alternative hypotheses to blame Meselson's and Stahl's auxiliary assumptions rather than their preferred model. In fact, Meselson and Stahl, on their part, were quite cautious in stating their conclusions:

⁹ (Hanawalt [2004]). Thanks to Beatrix Rubin for bringing this paper to my attention.

The structure for DNA proposed by Watson and Crick brought forward a number of proposals as to how such a molecule might replicate [the semi-conservative, dispersive and conservative mechanisms] These proposals make specific predictions concerning the distribution of parental atoms among progeny molecules. The results presented here give a detailed answer to the question of this distribution and simultaneously direct our attention to other problems whose solution must be the next step in progress toward a complete understanding of the molecular basis of DNA duplication. What are the molecular structures of the subunits of *E. coli* DNA which are passed on intact to each daughter molecule? What is the relationship of these subunits to each other in a DNA molecule? What is the mechanism of the synthesis and dissociation of the subunits *in vivo*? (Meselson and Stahl [1958], p. 681).

As this passage makes clear, Meselson and Stahl did not even *draw* the inference from their data to the semi-conservative mechanism, at least not in their official publication. Curiously, the questions they raise toward the end of this passage are precisely those that their experiment is supposed to have answered. In print, Meselson and Stahl did obviously not want to go beyond what their data said.¹⁰ However, unofficially they showed less caution. Meselson sent J. D. Watson a little poem¹¹:

Now ¹⁵N by heavy trickery/Ends the sway of Watson-Crickery./But now
we have WC with a mighty vengeance . . . or else a diabolical camouflage.

This line strongly suggests that Meselson did think that the experiment supported Watson's and Crick's replication scheme, even though he knew that they had not really established the nature of the molecular units that they resolved in their ultracentrifuge. As Holmes ([2001], p. 329) reports, 'Meselson and Stahl were personally convinced that the experiment had proven the position that

¹⁰ While the conclusive refutation of Stent's model would have required more certainty about the nature of the molecular units resolved by the centrifuge, Meselson and Stahl were at least confident enough that their results ruled out Delbrück's dispersive mechanism ([1958], p. 681). However, they did not justify this claim on the grounds that the intermediate band did not shift any further after subsequent rounds of replication (as most textbooks have it). Rather, they reported an additional experiment with heat-denatured *E. coli*-DNA as speaking against Delbrück. When heated enough, the two strands of DNA dissociate. Meselson and Stahl denatured heavy, hybrid, and light molecules in this way and analyzed them in the ultracentrifuge. The hybrid molecule produced bands of the same density as a mixture of heavy and light DNA. Furthermore, the molecular weight of the molecules was estimated to be reduced by half by denaturing. This suggested that the DNA strands themselves were not broken and re-annealed during replication (as Delbrück's mechanism required). However, Meselson and Stahl were worried because their *E. coli* DNA dissociated at a temperature at which salmon sperm DNA did not, which led them to wonder if *E. coli*-DNA might be 'a more complex structure' than salmon sperm DNA. (To my knowledge, this could simply be due to a greater GC content in salmon sperm DNA.) At any rate, this only shows again that there were serious doubts as to what the exact nature of the molecular units was that the density gradient resolved. I will show below that this did not affect the evidential import as regards the Watson-Crick scheme (Section 6.2).

¹¹ Meselson to Watson, 8 November 1957. Quoted from (Holmes [2001], pp. 327–8).

Watson and Crick had taken in 1953 to be right.' As they had been quite sceptical of the Watson–Crick hypothesis initially, the experiment seems to have had an impact on their degrees of belief.

Of course, the scientists' own methodological judgments cannot be our ultimate standard when engaging in normative philosophy of science. Nevertheless, we should not dismiss these judgments lightly. I take part of my task to be the explication or rational reconstruction of the scientists' own reasoning. Thus, I would like to show that there is a way of making sense of the apparent discrepancy between Meselson and Stahl's official view that they expressed in the conclusion of their PNAS paper, according to which the experiment only showed the equal distribution of parental nitrogen and not semi-conservative replication, and their unofficial view, which granted their own experiment much more discriminatory power.

The reasons might have to do with the control of inductive risk. In print, scientists will take as little inductive risk as possible in order to minimize the chance of being on the record for being wrong. But on other occasions they might be willing to take more inductive risk and go a little further beyond what their data say, especially when this allows them to carry on with their research.

The alternative replication models (or at least the conservative model; see footnote 10) could not be refuted then because there was an important untested auxiliary assumption in the interpretation of the data: that the molecular units represented as bands were single DNA molecules (and not, for example, end-to-end associations of heavy parental duplexes with newly synthesized, light double strands). The evidence for semi-conservative replication was only as good as this assumption. This is the inductive analogue of Duhem's first problem mentioned in Section 2, the problem of untested auxiliaries. And this is also why the simple eliminative induction scheme fails to bring out the evidential import of this experiment.

It should be clear by now that, if we want to be able to explicate Meselson's and Stahl's view that the experiment in its original form was strong evidence for the Watson–Crick mechanism, we must grapple with the problem of untested auxiliaries. Before I present my own IBE-based solution, I shall examine how another theory of scientific evidence handles this problem.

5 Severe Testing

Perhaps Meselson's hint with the 'diabolical camouflage' (see the above citation) is revealing. It indicates that Meselson thought it unlikely in the extreme that their experiment would have turned out the way it did had the semi-conservative hypothesis been false. This suggests yet another construal of the case: It could be argued that what Meselson and Stahl actually provided was a *severe test* in the sense of Mayo's ([1996]) error-statistical theory of scientific reasoning.

A severe test in this sense is a test with a low error probability, in other words, a low probability that the hypothesis passes a test in spite of being false. The term ‘error probability’ is originally a technical term from the Neyman–Pearson statistical method. However, Mayo ([1996]) argues that there is an informal analogue to such a test that also applies in non-statistical contexts. The centrepiece of this informal approach is what Mayo calls ‘arguments from error’. On her view, a hypothesis receives inductive support from some data to the extent in which it fits these data *and* some investigative procedure that was likely to detect an erroneous fit (i.e., one where the hypothesis under test is false) did not find an error.

A first problem with such a construal is to say what *justifies* the judgment that some test procedure was likely to find an error, had there been one. Specifically, how could Meselson’s judgment be justified that it was unlikely that DNA would behave as it did, had the semi-conservative scheme been false? I cannot think of a better answer than just saying that it would be a strange coincidence if Meselson’s and Stahl’s experiment behaved *as if* a semi-conservative mechanism was at work while, in fact, there was some other physiological mechanism at work. But this is just another way of expressing the *intuition* that this was unlikely; it does not really give a *justification* for it. Therefore, I think a reconstruction of the case as a severe test with an argument from error is not really helpful (unlike in cases where a formal Neyman–Pearson statistical test can be done).

But there is a second problem, and this is that there are major potential errors in the original experiment that Meselson and Stahl were not able to rule out in 1957. As already mentioned (Section 4), what the experiment showed primarily was the symmetrical distribution of heavy nitrogen in replication, not that the bands corresponded to single DNA duplexes. It was technically possible that the intermediate band represents an end-to-end association of parental DNA duplexes with newly synthesized duplexes rather than hybrid molecules composed of a light and a heavy strand (this would make the results compatible with the conservative hypothesis). This interpretation was ruled out about five years later, when Meselson’s student Ron Rolfe showed that the DNA could be broken into smaller fragments by sonication without affecting its density (Hanawalt [2004]).

According to the error-statistical approach, an experimental inquiry only supports a hypothesis to the extent in which it rules out possible errors in the data interpretation. Meselson and Stahl were unable to rule out some quite severe errors; therefore, this approach does not allow us to say that they had good evidence for the Watson–Crick hypothesis. Of course, a follower of the severe testing approach could just shrug her shoulders at this point and say: Well, in that case there really *was* no decisive evidence coming out of the Meselson–Stahl experiment, at least until Rolfe’s results were in.

The problem with this way of arguing is not merely that it does not reflect how many scientists thought about the experiment, including Meselson and Stahl themselves. The real problem is that it will *never* allow us to say that the experiment supported the Watson–Crick hypothesis; even once some additional tests had been done. Consider Rolfe's experiment, mentioned above. It might be viewed as ruling out the error of misidentifying the molecular units that worried Meselson and Stahl so much. However, it is always possible to come up with an interpretation of Rolfe's data that make them compatible with one of the alternative hypotheses. Perhaps there were covalent cross-links between the parental and newly synthesized duplexes that were resistant to the sonication treatment. So Rolfe's data were no severe test before these *other* possible errors were ruled out, and so on. When is the point reached where the data speak of the truth of a hypothesis? There are always possible errors, so the jury is never really in. For this reason, it seems to me that the error-statistical approach, far from solving Duhem's problem, makes it run amok.

The last two sections, I hope, have made it clear that we need to think about the Meselson–Stahl experiment in altogether different terms if we want to show that it was, in fact, strong evidence for the semi-conservative hypothesis.

6 An Experimentalist Version of IBE

I suggest that the Meselson–Stahl experiment selects the semi-conservative hypothesis by an inference to the best explanation (IBE).¹² In order to make this thesis good, I first need to elaborate on the relevant concept of scientific explanation. For the purposes of this paper, I shall adopt a *mechanistic* account of explanation. According to such an account, to explain a phenomenon means to describe a mechanism that produces this phenomenon. A highly influential account of the relevant concept of mechanism has been given by Machamer, Darden, and Craver ([2000]), who define mechanisms as 'entities and activities organized such that they are productive of regular changes from start or set-up conditions to finish or termination conditions'. A considerable body of scholarship exists now that shows how much experimental research in biology is organized around mechanisms in this sense (e.g., Bechtel [2005]; Darden

¹² See (Lipton [2004]) for a deep book-length philosophical study of IBE. Lipton's main example is Semmelweis's discovery of the cause of childbed fever. This is a case where IBE was used to pinpoint a causal factor ('cadaveric matter') that accounted for variations in the incidence of an infectious disease. Even though there was some controlled experimentation involved in this example, the case differs considerably from my example of molecular biology. One difference is that my example involves not merely the identification of a causal factor, but also the elucidation of a specific mechanism. Another difference is the use of a sophisticated measurement device. Problems such as Duhem's first and its inductive analogue (not discussed by Lipton) are more pressing here.

[2006]; Craver [2007]).¹³ To my knowledge, no one has yet shown how such an account of explanation could support an IBE-based account of induction.¹⁴

In order to do this, a new distinction is needed: I shall distinguish *physiological* from *experimental* mechanisms.

6.1 Physiological and experimental mechanisms

Physiological mechanisms are mechanisms that operate in a living cell. This kind of mechanism has received much attention lately. I would like to introduce a new type of mechanism: *experimental* mechanisms. In doing so, I shall leave the meaning of the term ‘mechanism’ itself the same, but allow the entities and activities as well as the changes, set-up and finish conditions to include parts of the *experimental system* (in the sense of Rheinberger [1997]). In other words, the artificially prepared materials such as the heavy nitrogen source as well as the characteristic manipulations and measurement devices used in the experiment also qualify as parts of a mechanism—the experimental mechanism. While physiological mechanisms occur in nature, experimental mechanisms require *interventions*.

In order to motivate this move a little, note that it makes perfect sense to speak of the mechanism that produced the UV absorption bands in Meselson and Stahl’s experimental setup. This mechanism includes the heavy nitrogen added to the growth medium, as well as the transfer of the growing bacteria into a medium containing light nitrogen. Furthermore, the mechanism includes the mechanical devices used to grind up the cells, extract the DNA, and transfer them onto the CsCl gradient (which, needless to say, is also part of the mechanism).

What is also important is that the physiological mechanism—i.e., the mechanism of DNA replication in this case—was somehow *embedded* in the experimental mechanism. In other words, it was responsible for some of the regular changes that constituted the experimental mechanism. Mechanisms often form hierarchical structures where particular entities and activities can be themselves decomposed into lower-level mechanisms (Craver and Darden [2001]). The lower-level mechanisms may be responsible for some of the activities that feature in higher-level mechanisms. But such a hierarchical organization is not necessary. Mechanisms may be related by one mechanism providing

¹³ The term ‘mechanism’ is sometimes used in a double sense in this literature, sometimes ontological and sometimes epistemic. In my view, the latter use should be understood as shorthand for ‘description of a mechanism’ or ‘model of a mechanism’ and the context should normally make it clear as to which of the two senses is relevant.

¹⁴ Lipton ([2004], p. 122) cites ‘mechanism’ as an ‘explanatory virtue’ (along with precision, scope, simplicity, fertility or fruitfulness, and fit with background beliefs), but the examples he discusses involve mostly just the identification of causal factors rather than the elucidation of elaborate mechanisms.

the *substrate* that another mechanism operates on. Biochemical pathways are a nice example for this. Thus, mechanisms may be *horizontally* linked. Such horizontal links exist in our present example: the heavy nitrogen is an entity of the experimental mechanism, and it is a substrate on which the physiological mechanism can act if it is provided instead of the usual substrate (i.e., light nitrogen). This entity has the *same* activity (or almost the same) within the *physiological* mechanism, but a *different* activity within the *experimental* mechanism. Finally, an important way for mechanisms to be embedded is by the physiological mechanism being a *stage* of the experimental mechanism. In other words, a product of the physiological mechanism (here: DNA) is further processed by the experimental mechanism.

We do not have to require that the embedded mechanism and the physiological mechanism that is under study are exactly the same.¹⁵ It might be enough if something similar to the physiological mechanism is embedded, provided that the embedded mechanism shares the salient nomological properties of the physiological mechanism. In experimental biology, mechanisms often come in families that may bear more or less resemblance to some prototype (Schaffner [1993], pp. 97–8). However, in this case, I think we can really say that the physiological mechanism itself was embedded. After all, we are dealing with an *in vivo* experiment. In other words, the experiment was done with living bacteria—at least before their DNA was extracted.

Of course, it must be assumed that the experimental mechanism is well understood by the experimenters. Generally, experimental mechanisms may be expected to be more accessible epistemically because, unlike physiological mechanisms, they were at least in part designed by the experimenter. But some of the assumptions made about experimental mechanisms may also be subject to independent test (see Section 6.4).

Why this extension of the notion of mechanism? What I would like to suggest is that the experimental mechanism is part of the explanation for the actual *data patterns* that Meselson and Stahl saw in their experiment (as shown in Figure 3). Further, I want to claim that this explanation is better than the two alternative explanations that involve the dispersive or conservative replication mechanism instead of the semi-conservative one. The experimental mechanism in combination with the semi-conservative physiological mechanism is the best explanation for the banding patterns obtained by Meselson and Stahl, at least in the group of experimental mechanisms that involve either the semi-conservative, the dispersive or the conservative mechanism and are otherwise identical. I will argue now that this explanatory relation is constitutive for the inductive support that the Meselson–Stahl experiment bestowed on the semi-conservative hypothesis.

¹⁵ I owe this point to an anonymous referee.

6.2 Explaining the data

In order to make this claim good, I need to be more specific what exactly IBE is. Lipton ([2004], pp. 58–9) has made a strong case for construing IBE as *inference to the loveliest potential explanation*. The relevant potential explanations are explanations that, if they were true, would explain the premises of the inference. (Actual explanations are actually true). The *loveliest* explanation is the one that, if true, would be most explanatory. Lipton also characterises the loveliest explanation as the one that provides ‘the most understanding.’ This should not be taken to imply that explanatoriness is a purely subjective or psychological matter, as it were, ‘in the eye of the beholder.’ Whether or not a set of propositions are explanatory with regard to some other set is a matter of their conforming to certain norms such as those of the mechanistic approach taken here. Thus, I do not mean ‘loveliness’, ‘explanatory’ and ‘understand’ in a psychological sense (i.e., as some subjective state of mind), but in a sense that is bound to strict normative standards as to what qualifies as an explanation. According to the mechanistic account that I adopt, to understand a phenomenon is a cognitive state characterized by an acquaintance with the entities and activities as well as certain patterns of counterfactual dependence (see Woodward [2002]) involved in producing the explanandum phenomenon, in particular insofar as it instantiates regularities.

The loveliest explanation should be conceptually distinguished from the *likeliest* explanation, which is the one that is most likely to be true. While likely explanations are what scientists strive for, to suggest that they infer explanations on the basis of their likeliness would be ‘like a dessert recipe that says start with a soufflé’ (Lipton [2004], p. 59). How likely an explanation is to be true is what we don’t know when we are drawing an inductive inference; what we can know is only how explanatory or lovely an explanation would be if it were true. On Lipton’s account, IBE is based on the idea that explanatory loveliness is a guide or a reliable indicator to likeliness. It is often enough the case that some complex state of affairs or some contrast (Lipton [2004], pp. 33–6) has some specific cause rather than another. A footprint in the wilderness that is exactly shaped like a bear paw is more likely to have been caused by a passing bear than, say, by surface air turbulence. That’s why it is recommended to infer the presence of a bear rather than some unusual turbulence when seeing such a structure in the wild. The passing bear is the loveliest explanation for the shape of the prints. Of course, the prints might have been caused by someone who was trying to pull a prank, but nobody claims that IBE is infallible.

The next step must be to specify what the relevant *explanandum* is, that is, the premise of the IBE. In contrast to other accounts, I suggest that the explanandum is provided by the *data* rather than the *phenomena* (in the sense of Bogen and Woodward [1988]). Applied to the present case, this means that

the semi-conservative hypothesis was the loveliest potential explanation of the *data* that Meselson and Stahl obtained, i.e., the banding pattern that they observed (see Figure 3). According to the mechanistic account of explanation that I adopt for the purposes of this analysis, a lovely explanation is a description of a mechanism, in other words, an arrangement of interlocking causal processes that together produce the explanandum facts.

It is central to my account that—unlike in Lipton—the explanandum that serves as a premise for the IBE are data, not phenomena. According to Woodward ([2000], p. S163), phenomena are ‘stable repeatable effects or processes that are potential objects of prediction and systematic explanation by general theories and which can serve as evidence for such theories’, while data are ‘public records (...) produced by measurement and experiment, that serve as evidence for the existence of phenomena.’ Phenomena and data are not related by relations of entailment, but by a hodgepodge of factual, empirical relations.¹⁶

In the realm of experimental biology, we need to refine these characterizations a little, because there are hardly any general theories. Instead, there are descriptions of mechanisms that serve the explanatory role of theories. Thus, for the case in hand, the phenomenon is the copying of DNA while the data are the banding patterns observed by Meselson and Stahl. Phenomena and data are causally connected through the experimental mechanism (see Section 6.1).

In our current example, the *phenomenon*—DNA copying—is potentially explained by all three replication mechanisms. Give or take some theoretical difficulties such as the ‘problem of untwiddling’, the three proposed mechanism schemes are about equally successful in explaining the phenomenon. Thus, the experimental IBE is not helping at this level. Where the experimental IBE argument comes into the picture is in that the experimental mechanism that contains the semi-conservative scheme provides the best explanation for the experimental *data*. The semi-conservative scheme augmented with the details of the experimental setup describes a collection of causal processes that can produce the UV absorption bands as they were observed by Meselson and Stahl, given their background knowledge about equilibrium sedimentation (see Section 6.4) and the biochemistry of nucleic acids. The alternative schemes in combination with the details of the experimental setup describe causal processes that would produce *different* banding patterns, unless they are fitted with additional assumptions that are not part of the physiological mechanism (see Section 6.3). This is why the semi-conservative scheme, together with the

¹⁶ There is much that I agree with in Woodward’s account concerning the relationship between data and phenomena. The only amendment that I would suggest is that inferences from data are sometimes guided by considerations of what mechanism(s) would explain specific data outcomes. This can help scientists to determine how reliably the data track competing claims about what causes the phenomenon in question.

details of the experimental mechanism, is the best explanation of the data and, ultimately, the reason why this experiment provided strong support for the semi-conservative hypothesis.¹⁷

I will show next that the account just outlined succeeds where the other approaches discussed in Sections 4–5 failed, namely, in solving the problem of untested auxiliaries.

6.3 IBE and the problem of untested auxiliaries

Meselson's and Stahl's data seemed to fit the semi-conservative hypothesis best, but, as we have seen, this judgment relied on the truth of auxiliary assumptions, in particular that the bands corresponded to simple DNA double strands (see Figures 4 and 5). As this assumption (together with some others) was not yet testable in 1957, how can we nonetheless justify the claim that the experiment spoke of the truth of the Watson–Crick hypothesis of replication, a judgment shared by many scientists at that time? I will argue now that IBE can be used to justify such a claim.

What needs to be shown is that the Meselson–Stahl experiment supported the semi-conservative hypothesis by its own wits, that is, without the help of additional tests that ruled out possible errors in the interpretation of the data (except the calibration of the instrument; see Section 6.4). I suggest that this is the case because the semi-conservative mechanism, in combination with what I have called the experimental mechanism, was sufficient to *explain* the data by its own wits. In contrast, the alternative mechanisms would require additional mechanisms or 'epicycles' in order to explain the Meselson–Stahl data. It would be utterly mysterious if, for example, the conservative mechanism would produce end-to-end associations of heavy and light DNA molecules, which could give rise to the pattern of bands that was actually observed (see Figure 2). With the semi-conservative mechanism, in contrast, it is absolutely clear why it is likely to produce this banding pattern; nothing is left mysterious. This is exactly the kind of difference that IBE is sensitive to.

¹⁷ An anonymous referee asked why it is not enough to say that the semi-conservative hypothesis *predicted* the outcome observed, and that it was supported by this outcome for this reason. To answer this question, it must first be noted that 'predict' is ambiguous in this context; it can mean either 'imply', or it can refer to what is known as 'novel prediction.' In either sense, prediction does not capture the methodological import of Meselson's and Stahl's data. The hypothesis did not entail the data (as the hypothetico-deductive account of confirmation would have it); rather, hypothesis and data were connected by factual, empirical relations (see Woodward [2000]). As for the second sense of 'predict,' we may have an example of a novel prediction here, although 'novel' can also mean different things. In any case, novelty cannot be the reason why the experiment *supported* the Watson–Crick hypothesis either, for it would have supported it no less if the latter hypothesis had been formulated *after* the experiment had been performed. In fact, the evidential support would have been just the same even if the hypothesis had been deliberately designed to fit the Meselson–Stahl data. This is why I think that the methodologically salient relation here is that the hypothesis, combined with the experimental mechanism, *explained* the data (see also Achinstein [2001]). Prediction is too heterogeneous a category from a methodological perspective.

This explanatory sufficiency, I suggest, is the methodologically relevant difference between the semi-conservative and the alternative schemes and the reason why the experiment supported the former but not the latter. Thus, the point is not that the semi-conservative mechanism was simpler or required fewer assumptions than the alternatives; what is crucial is that it was able to explain the data pattern at all.

The same point can be put as follows. It is as if the auxiliary assumptions were hitching a free inferential ride on the experimental mechanism, powered by the latter's explanatory force. Because the hypothesis under test is augmented by certain auxiliary assumptions in explaining the data, the IBE supports these assumptions along with the hypothesis. In order to elaborate on this idea, I shall make use of Norton's ([1993]) notion of relocating inductive risk. Norton discusses a case from the history of quantum physics where inductive risk was relocated from rules to premises such that the evidence uniquely determined a theoretical claim. In a somewhat similar way, I suggest, we can think of the inference from the Meselson–Stahl data to the semi-conservative hypothesis as dividing the inductive risk equally between the hypothesis itself and some of the auxiliary assumptions needed to connect the former to the data, in particular the assumption that the bands represent single DNA duplexes. Instead of requiring that this auxiliary assumption be secured *before* or *independently of* the inference to the theoretical hypothesis in question, we can say that *both* are inferred in one fell swoop on the grounds that the *combination* of them—in form of the experimental mechanism—provides a sufficient causal-mechanical explanation of the data. In other words, the experimental mechanism is inferred from the data by IBE as a whole structure.¹⁸ Because it contains the semi-conservative mechanism as a substructure, the latter is inferred along, but so are the auxiliary assumptions (which were of little theoretical interest themselves).

The reason why this solves the problem of untested auxiliaries is that there is no comparable inference to the alternative hypotheses. Even though there are ways of fiddling with the auxiliaries to make them *consistent* with the data, this will not provide a sufficient mechanistic *explanation* of them. This can best be seen in the case of Stent's conservative hypothesis. As I already mentioned, this hypothesis can be made consistent with the Meselson–Stahl data by assuming that the DNA duplexes somehow stick together after replication, for example, as end-to-end covalent associations of newly synthesized and parental polyphosphodeoxyribose nucleotide. Why can we not say that this assumption was *also* taken for a free ride by the conservative mechanism, in the same way in

¹⁸ There were also aspects of the mechanism that were not subject to this inference, for example, the Watson–Crick structure of the DNA double helix. Such theoretical assumptions are not involved in this sharing of inductive risk.

which the semi-conservative scheme took the assumption that the bands were single duplexes along in explaining the data?¹⁹

The reason is that there is nothing in, say, the conservative hypothesis that says that the molecules should remain so associated (see Figure 2). The assumption that the bands represented simple DNA double helices is part of a lovely explanation of the data (in the mechanistic sense), while the assumption that the bands represent end-to-end associations of parental and newly synthesized double strands is not part of such a lovely explanation. It is not clear why the alternative replication schemes should produce such structures. Of course, this does not mean that the alternative explanations are ruled out with certainty. But, once again, nobody claims that IBE is an infallible inference rule. No inductive inference rule is infallible. Proponents of IBE do not claim that the loveliest explanation is always true; they only say that this is the case often enough for this rule to be epistemically useful. Our present case is certainly no counter-instance.

It seems to me that only such an IBE-type of argument can make sense of the widely shared intuition (see Franklin [2007], Section E.2; Roush [2005], pp. 14–6) that the Meselson–Stahl data provided strong discriminatory evidence for the semi-conservative hypothesis. As I have shown, two other approaches to reconstructing the scientific reasoning behind ‘the most beautiful experiment in biology’ succumb to Duhem’s first problem, or its inductive analogue. The merits of a third approach to scientific inference in reconstructing this case—Bayesian confirmation theory—will be discussed in Section 8. But first, I want to demonstrate that an IBE-type argument can also be applied to those auxiliary assumptions that were actually tested.

6.4 IBE-turtles all the way down

Not all auxiliary assumptions used by Meselson and Stahl were free riders in the sense explained in the previous section. The main example of such a free rider was the identification of the bands with single DNA duplexes (Figures 4 and 5). However, even if this is granted, Meselson’s and Stahl’s evidence could only have been as good as the correlation between the density of the DNA and the position of the bands. As we have seen, it was crucially important that the band of intermediate density was lying exactly between the heavy and light bands. But how good was Meselson and Stahl’s analytic technique to resolve molecules according to their density? Franklin ([2007], Section E.2) argues

¹⁹ This objection is due to an anonymous referee. Note that my notion of ‘inferential hitch-hiking’ is not a new methodological principle; it is merely a way of describing how IBE operates in a case like this, where the explanation inferred is not a single proposition or a systematic theory but a heterogeneous assemblage of theoretical and experimental assumptions (the experimental mechanism).

that the mere fact that heavy and light DNA produced two clearly separated bands (before the actual experiment was done) provided some evidence that the technique was reliable. I agree, except that the linearity of the gradient was an important extra element in the argument that the intermediate contained ^{14}N and ^{15}N in equal amounts. Therefore, a good calibration of the instrument was of the essence for this experiment. For this reason, I want to briefly examine how this was done.

The theory of ultracentrifugation had been worked out to a large extent by Theodor Svedberg in the 1920s. In his PhD thesis, Matthew Meselson extended the work of Svedberg to experiments with solutions of very high density, such as the CsCl-gradients that they were using. In those days, molecular biology was institutionally closely associated with physics and physical chemistry (impressively, Meselson's thesis committee included Richard Feynman and Linus Pauling). Meselson investigated in particular the conditions under which a CsCl-gradient and the macromolecules that float in it would reach a point of equilibrium. At equilibrium, the centrifugal force and the buoyant force would balance each other, tending to keep the DNA at that point where its buoyant density equals that of the solution. But there is another force that tends to displace the DNA from this equilibrium: namely molecular diffusion or Brownian motion. Meselson was able to show theoretically that, at equilibrium, these opposing forces would generate a Gaussian distribution of the molecules. Here is the relationship that Meselson derived:

$$C_{PX_n}(r) = C_{PX_n}(r_0) \exp \frac{-(r - r_0)^2}{2\sigma^2}$$

This equation describes the concentration of a charged polymer such as DNA in a linear density gradient. This is a Gaussian distribution with standard deviation σ . Meselson also obtained the following expression for the standard deviation:

$$\sigma^2 = \frac{RT}{M_{PX_n} \bar{v}_{PX_n} (d\rho/dr)_{r_0} \omega^2 r_0}$$

where M_{PX_n} is the molecular weight of the polymer PX_n ; \bar{v}_{PX_n} is the partial specific volume of the polymer PX_n ; $(d\rho/dr)_{r_0}$ is the slope of the density gradient; ω is the angular velocity; and r is the distance from the rotation axis.

The width of the distribution therefore allowed the biologists to calculate the molecular weight of the bands. The physical reason for this is that lighter molecules diffuse more rapidly; therefore they will smear more strongly when they form a band.

Meselson and Stahl checked these theoretical results against their experimental data, using DNA from bacteriophage T4 as a marker. The agreement was quite remarkable (see Figure 6).




Figure 6. Agreement of the theoretical calculation with the measured DNA concentration at equilibrium. The DNA used in this experiment was derived from bacteriophage T4. This figure appeared only in Meselson's PhD thesis (Meselson [1957]), not in the 1958 publication. Reprinted with author's permission.

This clean result may be viewed as a test that the measuring device worked properly and that the gradient was almost perfectly linear over a certain range. Thus, distances from the centre of rotation translate directly into buoyant densities. This linear CsCl gradient was an important part of what I have called the experimental mechanism, which is the centrepiece of my IBE-based reconstruction of the experiment. I would like to call the theory of how centrifugation produces a linear gradient from which the density of molecules that float in it can be read off directly the *theory of the instrument*. The final question to be discussed in this section is how this theory of the instrument was confirmed.

My proposed answer to this last question is that the theory of the instrument was also supported by an IBE-type argument, and that explanation is best understood in the mechanistic sense. Here, the relevant experimental mechanism contains the DNA molecules, the caesium and chloride ions, as well as the water molecules. These entities interact by electrostatic forces and weak chemical bonds (hydrogen bonds). Further, this experimental mechanism involves the centrifuge itself with its rotor and the cell containing the DNA/CsCl solution. Together with physical laws²⁰ (Newton's laws, Coulomb's law, and the laws of thermodynamics), this mechanism explains why, under suitable conditions, DNA molecules will reach a sedimentation equilibrium, in which they are distributed in accordance with a Gauss curve where the mean is a linear

²⁰ Some proponents of a mechanistic account of explanation have argued that laws are redundant; all the explanatory work they were once thought to do can be captured by activities (Machamer, Darden, and Craver [2000]). I have criticised this view in a previous work (Weber [2005], Chapter 2). Woodward ([2002]) gives an account of mechanisms based on his counterfactual account of causal regularities. These differences are of no relevance for the present discussion.

function of density and the width an inversely linear function of molecular weight, which is what was actually observed. It is this explanatory relation that provided grounds for thinking that the analytic ultracentrifuge is a reliable instrument for determining the density of certain biopolymers. In other words, it's IBE-turtles all the way down.

7 Van Fraassen's 'Bad Lot' Argument

In the previous section, I have shown how the IBE approach combined with a mechanistic account of explanation solves the inductive analogue to Duhem's first problem. But we still have Duhem's second problem to cope with, which is the claim that scientists can never have rational grounds for believing that the set of available hypotheses includes one that is true. There is a more recent version of this argument that pertains directly to IBE, namely van Fraassen's 'bad lot' objection (van Fraassen [1989], p. 142ff.). According to this argument, IBE can perhaps rank a set of hypotheses with regard to their explanatory merits, but it cannot provide grounds for accepting one of them as true. For the best explanation could still be a very bad one; it affords no epistemic distinction to be the best of a bad lot.

The most direct way of answering this challenge would be to show that the list of alternatives was, in fact, exhaustive. This is what Roush ([2005], p. 15) suggests in her brief discussion of the Meselson–Stahl case:

It is hard to argue with the claim that all, some, or none of an original strand appears in a daughter molecule, and all, some or none exhaust the possibilities. The genius of the investigation, perhaps, was to have pitched the question at a level of description where this exhaustiveness could be achieved in a simple way.

According to Roush, the level of description chosen in this case rules out that there are alternatives that have not been considered. If she were right, this would be a direct counterexample to van Fraassen's and Duhem's claims that there are always unconsidered alternatives (see Section 2). In a similar vein, John Norton ([1993]) has argued that there are theoretical claims in physics that are completely determined by a body of evidence. Might something like this work here as well? Unfortunately, I do not think so. Roush's argument for the exhaustiveness of the three replication schemes is not successful; there are other conceivable schemes. For example, it is at least logically possible that the original molecule is degraded completely in the process; in other words, that both molecules are newly synthesized.²¹ There are other conceivable alternatives.²²

²¹ I wish to thank Eric Oberheim for pointing this possibility out to me.

²² One could also think of mechanisms that use some kind of intermediate (e.g., RNA or protein) for copying the DNA molecule. To use an analogy, in the early days of molecular biology, there were

To be sure, such a mechanism is not supported by Meselson and Stahl's data, but it is enough to cast doubt on the exhaustiveness of the list of alternatives. Even though Roush's formulation ('all, some or none of an original strand') makes it look as if there were a complete disjunction involved, this is a result of an ambiguity in this way of describing the alternatives. We will thus have to settle with something weaker than a complete disjunction.

At least Roush's argument points us in the right direction. The three hypotheses about DNA replication might not exhaust the space of all possible replication mechanisms, but they may be the only ones that satisfy certain *mechanistic constraints*. All three schemes of DNA replication had to incorporate some quite stringent constraints. Most importantly, the schemes had to explain how DNA molecules with the same nucleotide sequence as an existing molecule could be synthesized. Thus, explanatory considerations were already involved in the construction of the various hypotheses. This fits nicely with Lipton's ([2004]) two-filter strategy, according to which the generation of a number of 'live options' of candidate hypotheses is followed by a selection of the 'loveliest' one and where explanatory considerations enter at all stages of the research process, i.e., in both 'filters.' The main difference to my account is that I propose to base these explanatory considerations on a mechanistic account of explanation.

This mechanism-based view puts very stringent constraints²³ on what qualifies as a live option. Suitable candidate hypotheses must incorporate a considerable body of knowledge from organic chemistry and molecular biology. In my example, the double helix model was such a constraint. It incorporated a great body of knowledge from organic chemistry, the physical chemistry of colloids, and crystallography. Furthermore, it was already fairly clear at that time that the sequence of bases in DNA was biologically highly significant (see Crick [1958], who could already cite a considerable body of evidence that supported this idea). Therefore, the replication mechanism had to preserve the nucleotide sequence of DNA. The complementarity of base pairing provided a lovely explanation for how a mechanism of DNA synthesis could achieve this. Hence, it was set that either single- or double-stranded DNA had to serve as a template

ideas around that proteins are assembled directly on the DNA molecule (Gamow's 'diamond' hypothesis). Later, it was shown that protein synthesis requires RNA as an intermediate (Judson [1979], p. 252). The hypothesis that DNA replication might require an intermediate was, to my knowledge, never seriously entertained; but it cannot be ruled out a priori, which makes our case vulnerable to Duhem's and van Fraassen's arguments.

²³ A detailed list of such mechanistic constraints can be found in (Craver [2007], Chapter 3, Section 2). He distinguishes component constraints (given by the stock of available entities), spatial constraints (pertaining to the possible spatial organization of mechanisms), temporal constraints (time courses and necessary sequences of events), and active constraints (given by invariant relationships between intervention variables). I would add in particular *functional* constraints, i.e., considerations on what biological task the mechanism must perform (here: copying of the genetic material).

for the (then still putative) DNA polymerase. Indeed, all the three major replication mechanisms that were considered as live options during the mid-1950s incorporated this template idea. The great open questions were whether the template was single- or double stranded, and the extent to which the template was conserved in the process.

Thus, background knowledge imposed a set of mechanistic constraints on the space of possible solutions to the replication problem. There was simply no alternative that could satisfy all these constraints *and* explain the Meselson–Stahl data by their own wits. Only the Watson–Crick model passed both of these IBE-filters—this, I suggest, is what made this experiment so compelling. Possibly, many alleged ‘crucial experiments’ in and out of biology owe their strength to this kind of logic.

8 IBE and Bayesianism

So far, I have argued that IBE does a better job in reconstructing ‘the most beautiful experiment in biology’ than eliminative induction and the error-statistical approach to scientific inference. In this last section, I want to show that the IBE approach is not in conflict with the current mainstream theory (or family of theories) of confirmation, which, of course, is Bayesianism. Okasha ([2000]) and Lipton ([2004], Chapter 7) have already shown that IBE need not be in conflict with Bayesian constraints on personal degrees of belief. That is, proponents of IBE need in no way challenge the Bayesian’s credo that the only rational way of assigning probabilities to hypotheses that are subject to empirical confirmation is by conditionalizing on the evidence in accordance with Bayes’s theorem. But this does not make IBE superfluous. Far from it, IBE can provide a way of *realizing* the Bayesian formalism in concrete cases. The formalism as such makes no prescriptions as to how the prior probabilities and likelihoods ought to be set; it only says that once these have been set, the posterior probabilities are set as well, on pain of incoherence. This is quite a weak constraint. What IBE can do here is to provide some further constraints, for example, on the prior probabilities and likelihoods themselves. Sometimes, estimating how likely a hypothesis makes some piece of evidence, that is, estimating $p(e|h)$, might involve considerations as to whether h is able to explain e . In our present idiom, this means that there must be a mechanism whereby the state of affairs described in e is produced and where this mechanism, or parts of it, are described by h . Furthermore, explanatory considerations may be used to set prior probabilities, on which Bayesianism imposes no constraints.

This way of reconciling IBE and Bayesianism is not new. What I would like to do, briefly, is to show how the present case study illuminates this reconciliation and thus adds credibility to it.

The semi-conservative mechanism does not entail the Meselson–Stahl results; as I have argued (Section 6.1) the two are connected by the factual assumptions that describe the experimental mechanism. Since this mechanism is plausible given the background knowledge (see Section 7), we could say that its description h bestows a high probability on e (the banding pattern observed). Since no comparable mechanism explains e given the other hypotheses, we can say that the latter make e less likely. Assuming that the hypotheses do not differ too much in prior probability, this already means that the semi-conservative hypotheses had a higher posterior probability after the experiment was done. Alternatively, it could be argued that explanatory considerations afford the semi-conservative scheme with the most favourable likelihood ratio (Roush [2005]).

What is the advantage of supplementing the Bayesian account with explanatory considerations in this manner? I suggest that, in addition to providing a way of assigning likelihoods, IBE can also illuminate Bayesian solutions to Duhem's problem. Some Bayesians have argued that untested auxiliary assumptions could still have a prior probability that is sufficiently high to allow confirmation or disconfirmation of a hypothesis under test, thus taking the sting out of Duhem's problem (Dorling [1979]). The problem with this approach is that it implies that most experimental tests will only be conclusive for some people (namely, those who give a high prior probability to all the auxiliaries) and not for others, which is undesirable. Here, IBE can help: Explanatory considerations such as those discussed in Section 6.3 can be used to set bounds for belief in the auxiliaries, thus rendering evidence more objective.

9 Conclusions

I have argued that an experimentalist version of IBE permits a reconstruction of the Meselson–Stahl experiment according to which the latter provided decisive veridical evidence (Achinstein [2001]) for the semi-conservative hypothesis, while the two alternatives remained without such support. This is pretty close to what crucial experiments were always supposed to do, except that I am of course not claiming that such an experimental demonstration can reach the apodictic certainty of deduction (as Duhem required, see Section 2). In contrast, eliminative induction and the severe-testing approach fail to exhibit the evidential support of the experiment for the semi-conservative hypothesis.

In contrast to Lipton's ([2004]) account of IBE, I have used a mechanistic account of explanation. An advantage of such an account is that it does justice to actual explanations in molecular biology. Another advantage is that it makes explanation an objective relation between explanans and explanandum, which

means that the evidential relation can also be objective. I have also introduced the notion of an *experimental mechanism*, which is like the physiological mechanisms discussed by philosophers of biology and neuroscience, except that it includes experimental manipulations, instruments and artificially created entities and activities. Experimental mechanisms connect data and phenomena via causal processes. Physiological mechanisms are embedded in experimental mechanisms. This notion allows IBE to be extended to show how inferences can be drawn from experimental data such as the banding patterns observed in an analytic ultracentrifuge.

Finally, I have shown that the two predicaments that Duhem identified for crucial experiments (though on the assumption that all inferences would have to be deductive) as well as van Fraassen's well-known 'bad lot'-objection to IBE can be solved in the IBE-based framework that I have used. As regards Duhem's first problem, the experimentalist variant of IBE allows (fallible) inferences to hypotheses about mechanisms even if there are untested experimental assumptions.²⁴ The crucial move is the recognition that an experimental mechanism containing untested auxiliaries that is *sufficient* for explaining the data is better supported by these data than schemes that are not sufficient. As it were, the untested auxiliaries hitch a free inferential ride on the experimental mechanism. In addition to the free riders, there were experimental assumptions that were actually tested in this case, and these tests also involved IBE.

Van Fraassen's 'bad lot' problem (which I take to be basically Duhem's second problem as applied to ampliative instead of deductive inference) can be handled by showing how an extensive body of background knowledge provided a host of stringent material constraints on the candidate hypotheses. Mechanistic-explanatory considerations are involved in the construction of such candidates as well as in the selection of the best one by a crucial experimental test.

Finally, I have shown that this case study strengthens the view that proponents of IBE and Bayesians can be friends.

Acknowledgements

Versions of this paper were presented at the conferences 'Confirmation, Induction and Science' (London School of Economics, March 2007) and 'Generating Experimental Knowledge' (University of Wuppertal, June 2007). I am grateful to the organisers and audiences of these conferences. Further, I wish to thank Daniel Sirtes for critically reading several drafts and suggesting the title, and

²⁴ Note that the justificatory status of a proposition does not affect its explanatory force, at least not with respect to potential explanations (where the explanans need not be true), which is the relevant sense of explanation in IBE.

Matthew Meselson for giving me permission to reproduce a figure from his PhD thesis. This research was supported by the Swiss National Science Foundation.

Science Studies Program and Department of Philosophy
University of Basel
 Missionsstrasse 21, 4003, Basel
 Switzerland
 marcel.weber@unibas.ch

References

- Achinstein, P. [2001]: *The Book of Evidence*, Oxford: Oxford University Press.
- Bechtel, W. [2005]: *Discovering Cell Mechanisms: The Creation of Modern Cell Biology*, Cambridge: Cambridge University Press.
- Bogen, J. and Woodward, J. [1988]: 'Saving the Phenomena', *Philosophical Review*, **97**, pp. 303–52.
- Craver, C. F. [2007]: *Explaining the Brain. Mechanisms and the Mosaic Unity of Neuroscience*, Oxford: Oxford University Press.
- Craver, C. F. and Darden, L. [2001]: 'Discovering Mechanisms in Neurobiology: The Case of Spatial Memory', in P. K. Machamer, R. Grush and P. McLaughlin (eds), *Theory and Method in the Neurosciences*, Pittsburgh: Pittsburgh University Press, pp. 112–37.
- Crick, F. H. C. [1958]: 'On Protein Synthesis', *Symposia of the Society for Experimental Biology*, **12**, pp. 138–63.
- Darden, L. [2006]: *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*, Cambridge: Cambridge University Press.
- Delbrück, M. [1954]: 'On the Replication of Deoxyribonucleic Acid', *Proceedings of the National Academy of Sciences of the United States of America*, **40**, pp. 783–8.
- Dorling, J. [1979]: 'Bayesian Personalism and Duhem's Problem', *Studies in History and Philosophy of Science*, **10**, pp. 177–87.
- Duhem, P. [1954]: *The Aim and Structure of Physical Theory*, Princeton, NJ: Princeton University Press.
- Franklin, A. [2007]: 'Experiment in Physics', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, <plato.stanford.edu/entries/physics-experiment/>
- Hanawalt, P. C. [2004]: 'Density Matters: The Semiconservative Replication of DNA', *Proceedings of the National Academy of Science of the United States of America*, **101**, pp. 17889–94.
- Holmes, F. L. [2001]: *Meselson, Stahl, and the Replication of DNA: A History of 'The Most Beautiful Experiment in Biology'*, New Haven, CT: Yale University Press.
- Judson, H. F. [1979]: *The Eighth Day of Creation: Makers of the Revolution in Biology*, London: Jonathan Cape.
- Laudan, L. [1990]: 'Demystifying Underdetermination', in C. Wade Savage (ed.), *Scientific Theories*, Minnesota Studies in the Philosophy of Science, Vol. XIV, Minneapolis, MN: University of Minnesota Press, pp. 267–97.
- Lipton, P. [2004]: *Inference to the Best Explanation*, 2nd edition, London: Routledge.

- Machamer, P., Darden, L. and Craver, C. F. [2000]: 'Thinking about Mechanisms', *Philosophy of Science*, **67**, pp. 1–25.
- Mayo, D. G. [1996]: *Error and the Growth of Experimental Knowledge*, Chicago: University of Chicago Press.
- Meselson, M. [1957]: *Equilibrium Sedimentation of Macromolecules in Density Gradients with Application to the Study of Deoxyribonucleic Acid*, PhD thesis, Pasadena, CA: California Institute of Technology.
- Meselson, M. and Stahl, F. W. [1958]: 'The Replication of DNA in *Escherichia coli*', *Proceedings of the National Academy of Science of the United States of America*, **44**, pp. 671–82.
- Norton, J. D. [1993]: 'The Determination of Theory by Evidence: The Case for Quantum Discontinuity, 1900–1915', *Synthese*, **97**, pp. 1–31.
- Okasha, S. [2000]: 'Van Fraassen's Critique of Inference to the Best Explanation', *Studies in History and Philosophy of Science*, **31**, pp. 691–710.
- Rheinberger, H.-J. [1997]: *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*, Stanford, CA: Stanford University Press.
- Roush, S. [2005]: *Tracking Truth: Knowledge, Evidence, and Science*, Oxford: Oxford University Press.
- Schaffner, K. F. [1993]: *Discovery and Explanation in Biology and Medicine*, Chicago: The University of Chicago Press.
- Stanford, P. K. [2006]: *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford: Oxford University Press.
- Stent, G. S. [1958]: 'Mating in the Reproduction of Bacterial Viruses', *Advances in Virus Research*, **5**, pp. 95–149.
- van Fraassen, B. C. [1989]: *Laws and Symmetry*, Oxford: Clarendon.
- Watson, J. D. and Crick, F. H. C. [1953]: 'Molecular Structure of Nucleic Acids. A Structure for Deoxyribose Nucleic Acid', *Nature*, **171**, pp. 737–8.
- Weber, M. [2005]: *Philosophy of Experimental Biology*, Cambridge: Cambridge University Press.
- Woodward, J. [2000]: 'Data, Phenomena, and Reliability', *Philosophy of Science (Proceedings)*, **67**, pp. S163–S179.
- Woodward, J. [2002]: 'What Is a Mechanism? A Counterfactual Account', *Philosophy of Science (Proceedings)*, **69**, pp. S366–S377.