1-1-2014

Information versus knowledge in confirmation theory

Darrell Patrick ROWBOTTOM
Lingnan University, Hong Kong

Recommended Citation

Rowbottom, D. P. (2014). Information versus knowledge in confirmation theory. Logique et Analyse, 57(226), 137-149. doi: 10.2143/LEA.226.0.3032652

Preprint – Forthcoming in *Logique et Analyse*

Information versus Knowledge in Confirmation Theory

Darrell P. Rowbottom
Department of Philosophy
Lingnan University

darrellrowbottom@ln.edu.hk

1. Introduction: Background 'Knowledge' in Confirmation Theory

Confirmation functions are typically defined in terms of conditional probabilities, which involve three distinct types of argument: hypotheses (h), evidence (e), and background statements (b). Most obviously, this goes for Bayes's theorem; see, for example, Salmon (1990). But here are some other examples:

$$\mathrm{C}(h,e,b) = \frac{\mathrm{P}(e|hb) - \mathrm{P}(e|b)}{\mathrm{P}(e|hb) - \mathrm{P}(eh|b) + \mathrm{P}(e|b)}$$

Popper (1983)

$$r(h,\ e,\ b) = \log \frac{\mathrm{P}(h|e\&b)}{\mathrm{P}(h|b)}$$

Milne (1996)

$$l(h,\ e,\ b) = \log \frac{\mathrm{P}(e|h\&b)}{\mathrm{P}(e|{\sim}h\&b)}$$

Huber (2008)

Now many authors on confirmation-related issues (or working in formal epistemology more generally)—for example, Keynes (1921), Eells and Fitelson (2000), and Huber

(2007)—refer to *b* as 'background knowledge'.[1] But I shall argue in this paper that this should *not* be understood as knowledge in the usual sense of mainstream epistemology, i.e. as entailing belief and justification.[2] In fact, I will argue that it need concern neither justification nor belief, so that *b* does not even concern 'knowledge' on the true belief account proposed by Sartwell (1992). What I say will not tell against the rather more radical (and unpopular) notion of objective knowledge defended by Popper (1972, p. 286)—'where we take the word "knowledge" in the objective or impersonal sense, in which it may be said to be contained in a book; or stored in a library; or taught in a university'—except in so far as it may need to be conceded that *b* should be true. Indeed Allo (2010, p. 251), who advocates the kind of informational turn that I will argue for (in a limited context in the present paper), cites Popper (1968) approvingly and notes that his approach: 'immediately rules out most traditional theories of knowledge'. For present purposes, it will suffice to understand 'information' by way of a subtractive definition, in line with Dunn (2008, p. 581), as: 'what is left from knowledge when you subtract justification, truth, belief … [and] the thinker.' (We will later see that there is, however, some dispute over whether 'truth' should be subtracted.)

Although I will here focus on *b*, some might also be tempted, following Williamson (2000), to suggest that all evidence should be construed as knowledge too. (This is according to the E=K thesis, that all evidence is knowledge and *vice versa*.[3]) One interesting result of taking this route is that the distinction between *e* and *b* may appear somewhat arbitrary, provided that *b* should be understood as background

[1] Somewhat surprisingly, given his philosophical predilections, even De Finetti (2008, p. 36) writes that 'to speak of the probability of an event *tout court*, without any qualification, does not have any concrete meaning. Rather, it must be kept in mind that probability is always relative to the state of knowledge of the person who is making the judgement.' Elsewhere on the same page, however, he instead writes of changing states of *information*. And I believe he would have agreed with me, if pressed, that this is a superior formulation. This is indicated by his earlier choice of section title (p. 18) as 'Probability Depends on the Subject's State of Information', where he explains 'By the expression "state of information" I mean *all of* the previous experience, everything the person has seen, heard, read, and so forth.' (But note that 'knowledge possessed' is also referred to, again, on the same page. In short, De Finetti appears to use 'information' and 'knowledge' interchangeably, perhaps as a result of insensitivity to the difference between the notions in mainstream epistemology.) See also De Finetti (1974, §5.91), where the discussion is conducted purely in terms of information.

[2] Note that this does not presuppose a JTB-style account of knowledge. For example, Williamson (2000) takes knowledge to entail belief and justification (although he also holds that it is unanalyzable). See also Rowbottom (2010).

[3] For other criticisms of this view from the point of view of formal epistemology, see Williamson (2010, pp. 4–6).

knowledge. In short, if E=K holds then *b* would be background *evidence* in virtue of being background knowledge. And this raises the question "What's the difference between the evidence (*e*) and the background evidence (*b*)?"

One plausible way to answer this question is to appeal to the temporal order of the evidence, with *b* being evidence that is old relative to *e*, and the notion of *relative evidence*. For example, if *e* & *b* entails *h* but neither *e* nor *b* entails *h* in isolation, then we may say that *e* is evidence for *h relative to b.* Thus if *b* is our knowledge now, and we discover *e* tomorrow, then we might declare: "We have discovered new evidence for *h*!" But this would be elliptical. What we would really mean, in terms of the more fundamental notion of absolute evidence, is that the confirmation value of *h* relative to current total evidence (i.e. *e* & *b*) is greater than the confirmation value of *h* relative to past total evidence (i.e. *b*).

It may therefore be helpful to construe my following argument as that background *evidence* should not be construed as knowledge, or else that background statements should not be understood as evidence (e.g. if one insists that E=K). And since whether some evidence (or statement) is in the background rather than the foreground is merely a matter of historical contingency—e.g. the precession of the perihelion of Mercury may have been discovered after the proposal of general relativity, and therefore not have been part of *b* as it was at the time of its actual proposal—it is easy to see how what I will argue goes for *e* if it goes for *b*. In short, in what I argue below it is reasonable to substitute *e* for *b* (with a few superficial adjustments where necessary). To give a quick example, if *b* can involve information recorded in a book but not believed then so can *e*.

2. 'Background Knowledge' without Belief

When we employ confirmation functions, we are often interested in how confirmed a hypothesis is, or was, relative to the information available at a particular point in time.[4] Sometimes, it is true, we will only be concerned with what the scientists

[4] It would be wrong to dismiss appeal to 'background knowledge' simply on the basis that *b* will far outstrip what is believed by any individual scientist. Instead, one should allow for group beliefs, as well as group knowledge. One might also use an intersubjective (or interobjective) account of

involved in a scenario believe (or believed); for example, we might be interested in the estimated confirmation value of a theory with respect to views shared in a frank discussion, e.g. of Pauli's pilot wave theory relative to the debate concerning this at the Solvay conference of 1927.[5] However, other times we may be concerned not only with what scientists actually believe (or believed), but also with other information accessible to the community. (And note that what is said is *not* always what is believed. The norms of assertion are not always respected.)

Consider a scientist who has conducted many experiments to test a theory, and has painstakingly recorded the results in his notebook (or in a computer file). While he may remember what his statistical analysis of those results showed, e.g. the calculated average value of some variable, the relevant error range, and so forth, he would not remember each and every result. (In fact, he may never have possessed beliefs about many of those results. He may have typed them into a spreadsheet unthinkingly, or set up apparatus to gather them in an automated fashion.) So if his statistical analysis were flawed, he would be *mistaken* about how his results bear on the theory. Thus we might also say that he would be *mistaken* about how well confirmed (or corroborated) the theory was by his experimental results. And semantic quibbles aside, it is clear that information which has been collected as a result of human effort but never believed *can* bear a confirmation (or corroboration) relation to hypotheses in which we are interested. This is exactly why good scientists are concerned to check that they have analysed their data properly.[6]

I can only see one objection to this claim, which is that the data referred to above may not be information because it is not *propositional* in format. My response to this is twofold. First, as I will later explain, I am not convinced that all information need be propositional. Second, and more importantly at the present juncture, the foregoing

---

probability, where group degrees of belief are central, in one's favoured confirmation function; see Gillies (1991) and Rowbottom (2008a, Forthcoming).

[5] See Cushing (1994) for more on this.

[6] What I say here gels with the view of Allo (2011, p. 420) on being informed: 'what is… the difference between being informed that *p* and having a true belief that *p*? The main reason for keeping these distinct is that the latter includes a mental state whereas the former doesn't.' On such a view, we may say that we are primarily interested in *how the scientist is informed*, with a view to calculating confirmation. (I have some worries, however, about whether *mental states* are ultimately the issue. For example, one might take a betting interpretation, or a dispositional interpretation, of degrees of belief. See the further discussion in Rowbottom (2012, Forthcoming).)

argument goes through even if all information *is* propositional. This is because we may imagine a scientist writing complete sentences in his notebook, i.e. explicit propositions, even if one wants to insist that recording numbers in a table does not constitute a kind of short-hand (as I think it does).

Actually, it may even be possible to advance an argument that we *ought* to consider all the information *accessible* to us (from a purely epistemic, not pragmatic, point of view), even when some of it is believed by no-one, on the basis of Carnap's (1962, p. 211) 'Requirement of Total Evidence… [that] the total evidence available must be taken as a basis' when a probability is calculated. Normally, this requirement is only considered in a synchronic and/or local sense, with respect to background beliefs. And it is easy to see why it holds in this case. Electing to disregard an experimental result that is favourable for a theory which you dislike—without making any attempt to explain the result away via appeal to experimental error, or mistaken auxiliary hypotheses used in generating predictions—is an epistemic no-no. Carnap (ibid.) offers further examples:

> If a judge in determining the probability of the defendant's guilt were to disregard some relevant facts brought to his knowledge… or if a scientist pleading for a certain hypothesis omitted in his publication some experimental results unfavourable to the hypothesis, then everybody would regard such a procedure as wrong.

Although Carnap neglects to mention this, however, his very own examples suggest that the principle may also be appropriate in a diachronic and/or 'extended' sense. Return to our aforementioned scientist. His published reports on his tests are supposed to take into account all of the relevant experimental evidence at his disposal. But he need not—thank goodness!—memorise the results of each experiment he (or his automated system) conducted.

Perhaps it is true that we are interested in what such a scientist is *disposed* to believe; in his *dispositions to believe*, as explained by Audi (1994). (Before I wrote this sentence, I had a *disposition to believe* that '145 multiplied by 0.9 is 130.5' due to my understanding of elementary arithmetic. As I write *this* sentence, I occurrently believe

that '145 multiplied by 0.9 is 130.5'. And for quite a while after writing—until I forget—I will have the dispositional belief that '145 multiplied by 0.9 is 130.5'.[7]) But then, to reiterate, we are not merely interested in what he *actually* (occurrently or dispositionally) believes, at the point in time in question, when we calculate a confirmation value. We take the information in the notebook, which is not believed by anyone (unless we entertain the highly controversial notion of an 'extended mind'), to be relevant to the state of science (as opposed to the scientist). In short, the notebook is like an external memory. The information in this external memory can be read by anyone with the right hardware and software (e.g. perceptual apparatus and language skills).

There is one remaining worry. Didn't I admit at the start of this section that *sometimes*, at least, we will be interested only in the beliefs of scientists? So might we not conclude that there are two kinds of confirmation: one objective, not concerning only beliefs, and one personal or interpersonal, which necessarily concerns beliefs?[8] We could. But I do not think we should. Rather, I think that we consider only beliefs, of necessity, when we are interested in determining a person's or group's confirmation *estimate*—or best confirmation estimate—for some theory. (The mathematical formulae used may stay the same. The interpretation of the values can differ.) But even if I am wrong, it does not follow that any personal or interpersonal sense of confirmation involves knowledge. As we will see, the relevant beliefs may be unjustified or false.

3. 'Background Knowledge' without Justification

Now let's imagine that the argument made in the previous section fails (although I insist that it does not). Does it follow that *b* should be understood as 'background

[7] One could develop a similar account of dispositions to know, as opposed to dispositional knowledge. If one is disposed to form a true belief in a proposition in a reliable way, for instance, then one may have a disposition to know that proposition. Thus some might think I possessed a disposition to know that '145 multiplied by 0.9 is 130.5' before writing this paper, and now possess dispositional knowledge of that proposition.

[8] I am grateful to an anonymous referee for raising this interesting possibility. He/she agrees with me, however, that a personal or interpersonal sense of confirmation should not concern only true beliefs, and hence not only background knowledge.

knowledge' rather than 'background beliefs' (or 'background information that is believed')? Not if *b* may include *unjustified* beliefs.

The argument that it may is again based on possible applications of confirmation theory *to actual situations*. For in many actual situations, something is believed to be true without any justification in either an internal (i.e. personally accessible) or external sense.[9] (And although we will not discuss this here, information that has never been the object of belief can be classified as true, e.g. in a computer file, by unreliable processes or for poor reasons.)

Consider how often memories are appealed to in science and beyond, and how easily testimony based on memories may spread through a community. Now if we are to trust the findings of psychologists on how easily false memories may be created—see, for instance, Roediger and McDermott (1995)—then we must conclude that appeal to memory is *unreliable* in a significant class of circumstances (some of which will be relevant to memories invoked in scientific contexts). Hence it is not externally justified on a reliabilist account, in said class of circumstances, and it is difficult to see how it could be on any other externalist view of justification.

Alternatively, consider the widespread belief in N-rays in early twentieth-century France. These were posited by a well-respected senior physicist, Blondlot, and over 300 papers were written on the topic, by around 120 scientists, between 1903 and 1906.[10] But Blondlot was eventually discredited in an embarrassing fashion; for he claimed to be able to detect them even when an allegedly vital component of his apparatus—an aluminium prism—had been secretly removed by a visiting American physicist, Wood. Now the bulk of serious historical scholarship on this episode agrees that it is false, as Bird (2007, p. 67) has recently claimed, that 'Blondlot believed in the existence of what he called N-rays for what it is clear were entirely spurious and irrational reasons.'[11] (And even if Bird were right, then I could use this example, instead, in support of the view that beliefs unjustified in an *internal* sense can be

[9] In the words of Chisholm (1977, p. 17), on an internalist account: 'we can know what it is, on any occasion, that constitutes our grounds, or reasons, or evidence for thinking that we know.'
[10] The figures are reported by Lagemann (1977).
[11] Langmuir (1985) agrees with Bird's sentiments. Price (1975) and Nye (1980) do not.

widespread in science.[12]) Rather, it seems that Blondlot and the other scientists concerned were relying on naked eye observations in a class of scenarios in which these are unreliable. Here's what Bauer (2002) says about the episode:

> [P]resumably what was pathological here was a reliance on visual observation under conditions—a darkened room—where optical illusions readily occur. (One modern test for glaucoma is to note over what field of view one can detect flashes of light on a dark background. Anyone who has taken such a test knows that one 'sees' some number of flashes that are not actually there.) But Blondlot was a distinguished member of the French scientific establishment. He had been particularly praised for showing that X-rays moved at the speed of light *which he had established by the same method of visual observation*, in that case variations in the apparent intensity of electric sparks. Blondlot was therefore very unfortunate; but how can he be blamed for continuing to use a technique that had been so successful? "The curious error of N-rays is much more a sort of mass hallucination, proceeding from an entirely reasonable beginning" (Price 1975, p. 159). Moreover, the *facts* Blondlot reported were confirmed by a number of his fellow scientists, not only in his laboratory but also elsewhere in France; which gave Blondlot good reason to think his discovery a genuine one.

Problems with memories, or even with simple observational beliefs, are also significant when we consider internal accounts of justification instead. For example, it is hard to see how the mere fact that I possess a memory that *p* makes me justified in believing that *p* when I am painfully aware of the ways in which my memory formation procedure may be flawed.[13] Consider, in particular, the research in empirical psychology indicating that we often have false memories concerning the

[12] The *possibility* of widespread unjustified beliefs in science is also defended in Rowbottom (2008b) and Rowbottom (2010), and appears to be accepted by Bird (2008) (in line with the quotation in the main text).

[13] As Steup (2006) notes, 'Some internalists think that sensory experiences, in and by themselves, constitute evidence… Other internalists, however, think that S's sensory experiences constitute evidence only if S can coherently view them as a reliable guide to truth.' The first route seems highly dubious, in so far as justificatory support, if it is to exist, only obtains between *propositions* or statements; and clearly we must not confuse the statements I believe to be true as a result of my experiences with the experiences themselves. (For more on this, see Popper's discussion of what he calls 'Fries's trilemma', which is better thought of as a dilemma-inducing dilemma, as shown by Bartley (1984, app. 3). For some interesting background on the trilemma, see Floridi (1996, pp. 112–117).)

sources of—and sometimes, therefore, of the reasons for which we came to hold—our (true) beliefs.[14] (And note that this research does not only concern odd scenarios. On the contrary, several are commonplace.)

Besides, there are also actual situations in which theories, and predictions made by those theories, are uncritically maintained in the face of contradictory evidence. Think of poor Mpemba, the Tanzanian student who was ridiculed by his teacher and classmates for claiming that warm ice cream mix can freeze faster—or at the very least can *start to freeze faster*—than cold mix placed in the same freezer, and later claiming that the same was true of water. Mpemba certainly had the last laugh—see Mpemba and Osborne (1969), and Esposito et al. (2008)—although he was fortunate to encounter a physicist willing to take his claim seriously enough to test it (and displayed admirable courage in raising the question in public in spite of the aggressive dogmatism he had repeatedly been confronted with).[15] In short, anything that (superficially) looked to challenge Newton's law of cooling was typically dismissed by those better versed in physics than Mpemba. And the worry is that belief in said law, gained by testimony, was not sufficient to grant good reasons for believing that hot water would never start to freeze faster than cold, *in the face of testimony to the contrary by someone who had done an appropriate experiment*.

A natural objection is that the source of the testimony—and in particular, the (perceived) expertise of the testifier—is relevant. Why take a school child seriously, in the face of the theory of a genius, and the might of the scientific establishment? My answer is twofold. First, the experiment is extremely easy to conduct; it was well within the abilities of Mpemba, as it is most teenagers. (Think of what justification there was, if any, for the teacher to fail to repeat it at least once. Failure to repeat, with ample opportunity to do so easily, plausibly removed any initial justification for believing Mpemba to be wrong.) Second, there are independent physical reasons not to be confident that Newton's law of cooling is gospel. As Osborne puts it, recounting his initial reaction to Mpemba's claim: 'everyday events are seldom as simple as they

[14] See, for example, Schacter et al. 1984, Jacoby et al. 1989, and Kelley and Lindsay 1993.
[15] The dogmatism was not confined to Mpemba's school. Osborne writes: 'At the University College in Dar es Salaam I asked a young technician to test the facts. The technician reported that the water that started hot did indeed freeze first and added in a moment of unscientific enthusiasm: "But we'll keep on repeating the experiment until we get the right result."'

seem and it is dangerous to pass a superficial judgment on what can and cannot be … it was possible that the rate of cooling might be affected by some factor I had not considered' (Mpemba and Osborne 1969, 173–174).[16]

Obviously I have not done full justice to the intricacies of accounts of justification in epistemology, or even come close, in the above. But the claim that unjustified beliefs are *sometimes* present in science *and more particularly the background assumptions employed by scientists*, is hardly objectionable. This is all I require in order for my overarching argument to go through.

I should add that even justified true beliefs are not normally taken to constitute knowledge, due to the Gettier (1963) problem. This is yet another reason to insist that *b* should not be understood as 'background knowledge', when we consider actual cases. (See also the thought experiments, and more detailed discussion of different construals of knowledge and justification, in Rowbottom (2010).) In short, it suffices for justified true beliefs falling short of knowledge to sometimes be present.

4. 'Background Knowledge' without Truth?

Naturally, background assumptions used in actual science may also be false. Sometimes these assumptions may be rather simple, e.g. concern the number of planets in the solar system (in the case of the prediction of the paths of certain heavenly bodies). Other times they may be more complex, in so far as *b* may involve theories other than *h*; when the famous experiment in 1919 was conducted in order to examine whether light bends around the sun, for example, various theories other than general relativity were explicitly or implicitly taken for granted. The telescopes used were constructed on the basis of theories of optics. The results were taken to concern light, *qua* theoretical electromagnetic entity, in so far as Maxwell's equations were not questioned. Telescopically observed deviations in the apparent position of stars were not attributed to atmospheric effects. And this was ultimately due to theoretical notions such as index of refraction.

[16] Keep Duhem's thesis in mind. The law of cooling does not, taken alone, issue in any predictions. Hence one simple explanation of the Mpemba effect, for example, might have been evaporation.

However, one might argue that the *real* confirmation values are determined by the subset of scientific statements classified as true, at some time *t*, *which are true*. And once we have made confirmation an objective matter to the extent that it need not concern (individual or group) mental states at a particular time, as suggested above, this move appears far less unreasonable than it otherwise might. If one follows Floridi (2004, 2005) in thinking that all information is true, one might also maintain the view that 'background information' is an appropriate replacement for 'background knowledge'.

So should we hold that all information is true? I will not go into this issue in any serious depth here, for fear of straying too far from the central topic of this paper. Instead, suffice it to say that I am in two minds about whether a more lax notion of information would be appropriate. At first blush, "Darrell Rowbottom is six foot two inches tall" contains information about me despite its falsity (strictly speaking). But I suppose one might insist that this is an elliptical way of expressing "Darrell Rowbottom is approximately six foot two", in everyday discourse. So far so good, but matters become more complex and troublesome when we consider scientific models. Everyone now agrees that Rutherford's planetary model of the atom is false. Yet it may nevertheless hold valuable information about the atom that is not exhausted by the true observational consequences that can be derived from it (in combination with appropriate auxiliary hypotheses). In particular, some of the structure of the solar system may be present in the atom—the relation between the planets and the Sun may mirror that between the electrons and the nucleus in some respects, e.g. with respect to relative sizes, for instance—and thus we may think that we are informed about the atom to some extent in virtue of possessing such a model. It is also natural to think that science can progress through the development of models that bear better structural resemblances to their targets than their predecessors.[17]

Champions of the view that all information is true may respond by arguing that a number of true analogical statements can be derived from such models, e.g. 'Nuclei are bigger than electrons', and that *this* is the sense in which they are informative (and can be more or less so). Indeed, I discuss such analogical statements in Rowbottom

[17] One also makes progress in understanding science by changing the models in terms of which one thinks. See, for example, Collins and Gentner (1987).

(2011). However, it remains to be seen whether such a strategy will ultimately do the trick because it is unclear that all the information present in a model should be thought of as propositional.[18]

For present purposes, then, let me just say that even if one insists all information is true, it does not follow that *b* should be understood as background knowledge rather than background information. The result, rather, is that one route to arguing that *b* should not be understood as background knowledge—that often we are interested in how confirmed a theory is relative to propositions which are classified as true but are actually false[19]—is precluded. For the record, I should add that I currently prefer the view that information may be false, and that *b* should be allowed to include false statements.[20]

Acknowledgements

Thanks to Jon Williamson and two anonymous reviewers for comments on earlier versions of this paper. Special thanks to Patrick Allo for reading several versions of the paper, and for numerous suggestions about how to improve it. Finally, I am grateful to the British Academy for supporting some of my work on the paper.

References

Allo, P. 2010. 'Putting Information First: Luciano Floridi and the Philosophy of Information', *Metaphilosophy* 41, 247–254.

Allo, P. 2011. 'The Logic of "Being Informed" Revisited and Revised', *Philosophical Studies* 153, 417–434.

[18] This is one reason why I am more inclined toward a *representational* theory of information, where *x* contains information about *y* in virtue of bearing an appropriate representational relation to *y*, which I hope to have the opportunity to articulate and explore on another occasion.

[19] Naturally a defender of the view that *b* should be construed as background knowledge may suggest that we are interested, rather, in the counterfactual question of how well confirmed the theory would be *if* those propositions were really true. The only obvious criticism of this suggestion proceeds from general worries about the possible veracity of counterfactual claims. So if we can avoid counterfactual talk in our discourse concerning confirmation, so much the better.

[20] Williamson (2010, pp. 4–6) agrees with this, in advancing his view that our evidence is just what we take for granted.

Audi, R. 1994. 'Dispositional Beliefs and Dispositions to Believe', *Nous* 28, 419–434.

Bartley, W. W. 1984. *The Retreat to Commitment*. La Salle: Open Court.

Bauer, H. H. 2002. '"Pathological Science" is not Scientific Misconduct (Nor is it Pathological)', *Hyle – International Journal for Philosophy of Chemistry* 8, 5–20.

Bird, A. 2007. 'What is Scientific Progress?', *Noûs* 41, 64–89.

Bird, A. 2008. 'Scientific Progress as Accumulation of Knowledge: A Reply to Rowbottom', *Studies in History and Philosophy of Science* 39, 279–281.

Carnap, R. 1962 [1950]. *Logical Foundations of Probability*. Chicago: University of Chicago Press, 2nd Ed.

Chisholm, R. 1977. *Theory of Knowledge*. Englewood Cliffs: Prentice Hall, 2nd ed.

Collins, A. and Gentner, D. 1987. 'How People Construct Mental Models', in D. Holland and N. Quinn (eds), *Cultural Models in Thought and Language*, pp. 243–265. Cambridge: Cambridge University Press.

De Finetti, B. 1974. *Theory of Probability: A Critical Introductory Treatment, Volume 1*. London: John Wiley & Sons.

De Finetti, B. 2008. *Philosophical Lectures on Probability: Collected, Edited, and Annotated by Alberto Mura*.

Dunn, J. M. 2008. 'Information in Computer Science', in J. Van Benthem and P. Adriaans (Eds), *Handbook on the Philosophy of Information*, pp. 581–608. Amsterdam: Elsevier.

Eells, E. and Fitelson, B. 2000. 'Measuring Confirmation and Evidence', *Journal of Philosophy* 97, 663–672.

Esposito, S., De Risi, R., and Somma, L. 2008. 'Mpemba Effect and Phase Transitions in the Adiabatic Cooling of Water Before Freezing', *Physica A* 387, 757–763.

Floridi, L. 1996. *Scepticism and the Foundation of Epistemology: A Study in the Metalogical Fallacies*. Leiden: Brill.

Floridi, L. 2004. 'Outline of a Theory of Strongly Semantic Information', *Minds and Machines* 14, 197–222

Floridi, L. 2005. 'Is Information Meaningful Data?', *Philosophy and Phenomenological Research* 70, 351–370.

Gettier, E. L. 1963. 'Is Justified True Belief Knowledge?', *Analysis* 23, 121–123.

Gillies, D. A. 1991. 'Intersubjective Probability and Confirmation Theory', *British Journal for the Philosophy of Science* 42, 513–533.

Huber, F. 2007. 'Confirmation and Induction', *Internet Encyclopedia of Philosophy*, URL: http://www.iep.utm.edu/conf-ind/

Huber, F. 2008. 'Milne's Argument for the Log-Ratio Measure', *Philosophy of Science* 75, 413–420.

Keynes, J. M. 1921. *A Treatise on Probability*. London: Macmillan.

Jacoby, L. L., C. Kelley, J. Brown, and J. Jasechko. 1989. 'Becoming Famous Overnight: Limits on the Ability to Avoid Unconscious Influences of the Past', *Journal of Personality and Social Psychology* 56, 326–338.

Kelley, C. M. and D. S. Lindsay. 1993. 'Remembering Mistaken for Knowing: Ease of Retrieval as a Basis for Confidence in Answers to General Knowledge Questions', *Journal of Memory and Language* 32, 1–24.

Lagemann, R. T. 1977. 'New Light on Old Rays: N Rays', *American Journal of Physics* 45, 281–284.

Langmuir, I. 1985. 'Pathological Science: Scientific Studies Based on Non-existent Phenomena', *Speculations in Science and Technology*, 8, 77–94.

Milne, P. 1996. 'Log[P(h/eb)/P(h/b)] Is the One True Measure of Confirmation', *Philosophy of Science* 63, 21–26.

Mpemba, E. B. and Obsorne, D. G. 1969. 'Cool?', *Physics Education* 4, 172–175

Nye, M. J. 1980. 'N rays: An Episode in the History and Psychology of Science', *Historical Studies in the Physical Sciences*, 11, 125-56

Popper, K. R. 1968. 'Epistemology Without a Knowing Subject', in B. van Rotselaar and J. F. Staal (eds), *Logic, Methodology and Philosophy of Science III*, pp. 333–373. Amsterdam: North-Holland.

Popper, K. R. 1972. *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.

Popper, K. R. 1983. *Realism and the Aim of Science*. London: Routledge.

Price, D. de S. 1975. *Science Since Babylon*. New Haven: Yale University Press.

Roediger, H. L. and McDermott, K. B. 1995. 'Creating False Memories: Remembering Words Not Presented in Lists', *Journal of Experimental Psychology: Learning, Memory and Cognition* 21, 803–814.

Rowbottom, D. P. 2008a. 'Intersubjective Corroboration', *Studies in History and Philosophy of Science* 39, 124–132.

Rowbottom, D. P. 2008b. 'N-rays and the Semantic View of Scientific Progress', *Studies in History and Philosophy of Science* 39, 277–278.

Rowbottom, D. P. 2010. 'What Scientific Progress Is Not: Against Bird's Epistemic View', *International Studies in Philosophy of Science* 24, 241–255.

Rowbottom, D. P. 2011. 'The Instrumentalist's New Clothes', *Philosophy of Science* 78, 1200–1211.

Rowbottom, D. P. 2012. 'Objective Bayesianism Defended?', *Metascience* 21, 193–196.

Rowbottom, D. P. Forthcoming. 'Group Level Interpretations of Probability: New Directions', *Pacific Philosophical Quarterly*.

Salmon, W. C. 1990. 'Rationality and Objectivity in Science or Tom Kuhn Meets Tom Bayes', in C. W. Savage (ed.), *Scientific Theories*, pp. 175–204. Minneapolis: University of Minnesota Press.

Sartwell, C. 1992. 'Why Knowledge Is Merely True Belief', *Journal of Philosophy* 89, 167–180.

Schacter, D. L., J. L. Harbluk, and D. R. McLachlan. 1984. 'Retrieval without Recollection: An Experimental Analysis of Source Amnesia', *Journal of Verbal Learning and Verbal Behavior* 23, 593–611.

Steup, M. 2006. 'The Analysis of Knowledge', in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, URL: http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/

Williamson, J. O. D. 2010. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.