Towards a Bayesian Account of Explanatory Power

February 1, 2010

Abstract

This paper defends an explication of explanatory power in Bayesian terms. It is argued that explanation and confirmation are intimately related, both on the qualitative and the quantitative level. Thus, the success of Bayesian confirmation theory can be used to argue for the transfer of Bayesian methods to an analysis of the quantitative, epistemic dimension of scientific explanations. The starting point is Hempel's (1965) observation that good explanations *rationalize* their explanandum, making the concept suitable for a Bayesian analysis. Thereby, the paper does not only reveal the deep structural similarities between confirmation and explanation, but also makes a contribution to better understanding explanatory reasoning in science, particularly statistics. We obtain a subjectivist theory of explanatory power which is better suited to describe statistical practice than several objectivist proposals that flourished in the past.

1. Introduction.

The search for an adequate explication of scientific explanation has been going on for the last sixty years, and has remained inconclusive. Starting with Hempel and Oppenheim's (1948) famous Deductive-Nomological (D-N) model, lots of sophisticated proposals have been made, but none of them succeeded at rebutting the flood of objections and counterexamples. There is no consensus about whether and how a scientific explanation can be characterized, using a set of precise and unambiguous conditions. Two ways out seem to suggest themselves: First, we could acknowledge that contextual and pragmatic elements in scientific explanations are irreducible (Van Fraassen 1980; Ruben 1990), standing in the way of any general explication or definition. Apparently, we have to be satisfied with a pluralist account where different characterizations of scientific explanations serve different purposes, while dismissing the search for an overall account. Second, we could claim that all good scientific explanations ultimately provide the *cause* of the explanandum. Then, we would presumably have to shift our focus from a conceptual analysis of explanation to an analysis of causation; such an analysis would, *a fortiori*, also deliver an account of successful scientific explanation (Salmon 1984; Woodward 2003).

I believe that there is a third option, too: to stick to the original explicative project, while setting different goals. When searching for a characterization of scientific explanation in terms of necessary and sufficient conditions, the *quantitative* aspect of explanation has been neglected. But it is all but natural to ask to what extent a phenomenon is explained by some hypothesis, how we can quantify its explanatory power.

Therefore I propose to focus on the salient quantitative aspects of explanation, without aiming at an exhaustive description of the concept itself. It is natural to adopt a probabilistic, Bayesian model for this purpose, so much the more as a steadily increasing number of natural and social sciences make use of probabilistic models and forecasts. This implies that a Bayesian model of explanatory power can naturally attach to the growing number of probabilistic inferences in science.

Such an approach is analogous to what has been done in confirmation theory over the last decades. And indeed, I claim that the success of the quantitative, Bayesian research program in the analysis of confirmation substantiates the hopes that a Bayesian account will fare equally well with respect to explanatory power. By the way, confirmation theory exemplifies, similar to the explanation debate, a history of failures and resistent problems in purely qualitative research programs aiming at a set of necessary and sufficient conditions.

To my mind, applying a Bayesian approach to *explanation* has the greatest potential to overcome the present deadlocks, to preserve the independence of explanation vis-à-vis causation, and to set new trends for the debate. It has to be argued, though, that the parallels between both concepts are close enough to warrant an argument by analogy. This similarity is elaborated and defended in the second section that briefly compares the qualitative aspects of both concepts. The third section devotes some detail to spelling out the merits of a quantitative analysis of explanation. The fourth section addresses some philosophical concerns, such as apparent subjectivism such of a Bayesian approach and the relation between explanation and causation. The final section concludes.

On a whole, I try to show that, although the subjectivist approach to explanations forces us to break with some traditions, it also opens the way to a fruitful research program about how to quantify explanatory power. A Bayesian account of explanatory power is very natural and not much bolder than a Bayesian analysis of confirmation which is an integral part of modern philosophy of science.

2. Confirmation: From Successful Predictions to Inductive Arguments.

The oldest attempt to analyze explanation with elementary formal tools is Hempel and Oppenheim's (1948) D-N model – explanations are valid deductive arguments that contain at least one empirical law among their premises:

Deductive-Nomological Model of Explanation The pair $\langle H, A\rangle$ is an explanans for explanandum E if and only if

1. H is a true empirical law,[1]
2. E follows deductively from $H.A$ ($H.A \vdash E$),
3. E does not follow deductively from A alone ($A \nvdash E$)

In other words, an empirical law H provides, together with auxiliary hypotheses A, an explanation of E if and only if E can be derived from H and A, taken together, and H is crucially required for the derivation. Briefly, explananda are subsumed under covering laws. The structural similarity to the most venerable account of confirmation, the Hypothetico-Deductive (H-D) model of confirmation is obvious:

Hypothetico-Deductive Model of Confirmation A piece of evidence E confirms a hypothesis H relative to background assumptions K if and only if

1. $H.K$ is consistent,
2. E follows deductively from $H.K$ ($H.K \vdash E$),
3. E does not follow deductively from K alone ($K \nvdash E$)

The idea of H-D confirmation is thus that a piece of evidence E confirms a theoretical hypothesis H if it is a prediction of H, where some background knowledge K helped to create the deductive link between H and E. So in both cases, the evidence is predicted by a body of theoretical propositions, some of which have auxiliary character, and some of which stand central to

[1] That the explanans must be true is, for Hempel and Oppenheim (1948, 137), rather an empirical than a *logical* condition of adequacy, so that I will neglect this requirement in the remainder.

the explanation/confirmation relation. The "systematic power" of a theory (Hempel and Oppenheim 1948, 164), the power to derive as many data as possible from a small amount of initial information, is not only an asset for giving deductively-nomological explanations, but also for being confirmed in the hypothetico-deductive way.

Moreover, both models build, either implicitly or explicitly, on an account of lawlikeness. For the case of explanation, this requirement is motivated by observing that no arbitrary generalization which accidentally happens to be true, increases scientific understanding. Only empirical laws have the power to explain a phenomenon because they state the fundamental regularities that govern nature, and systematize our body of knowledge (Carnap 1966). Similarly, Goodman (1983) noticed that most natural accounts of confirmation, including the H-D-model, open the door to arbitrary confirmation, if not suitably qualified. They fail to single out a method of inductive projection: does the observation of a green emerald confirm that all emeralds are green or the rivalling hypothesis that all emeralds are "grue" (=green if examined in the past, blue if examined later)? Thus, the problem of lawlike vs. accidental generalizations equally affects the H-D model of confirmation and the D-N model of explanation. Both models share the same syntactic structure and the same problem of defining what kind of hypothesis counts as confirmable or explanatory.

Given all that, it is not surprising that both accounts resemble each other with respect to a further challenge as well: discerning relations of evidential/explanatory relevance. In both models, tacking irrelevant propositions to the confirmed hypotheses/the explanans preserves the confirmation/explanation relation, as we can easily see from the definitions. But irrelevance seems to be fatal to both explanation and confirmation (Salmon 1984; Ruben 1990).

At this point, one might conclude that the respective explications given by the H-D and the D-N model are too close to each other in order to neatly separate the concepts of explanation and confirmation. In the light of that concern, Hempel's (1945) decision to settle for a new account is noteworthy. In his "Studies in the Logic of Confirmation", he suggested that the view of confirmation as successful prediction was way to narrow; quite often deductive and quasi-inductive steps are both contained in the derivation of a prediction from some general theory. He proposed to view confirmation differently, as an argument from the evidence to (some specific implications of) the theory. Hempel's original characterization was the *satisfaction criterion*, a criterion stating that the restriction of the hypothesis to the domain of the evidence had to be satisfied in our evidential model. I am giving my own formulation of that criterion which preserves the Hempelian spirit, but points out the role of the background assumptions, too:[2]

Hempel's Satisfaction Criterion for Confirmation Evidence E (directly) confirms hypothesis H relative to background knowledge K if and only if

1. $E.K$ is consistent,

[2]This element is missing in Hempel's original account, where confirmation is treated as a two-place predicate.

2. $E.K$ deductively entails the restriction of H to those objects that are (relevantly) mentioned in E,
3. K alone does not deductively entail the restriction of H to those objects that are (relevantly) mentioned in E,

While far from being the solution to all problems of qualitative confirmation theory (Fitelson 2008), Hempel's proposal at least circumvents the tacking paradoxes; it also clarifies the distinction between explanation and confirmation: While D-N explanation aims at providing arguments whose conclusions are observations in need of explanation, Hempel's satisfaction criterion explicates confirmation as a hybrid deductive-inductive argument from the evidence to a general hypothesis. While the H-D and the D-N model are linked by a direct structural *analogy*, we can perceive a *duality* between the satisfaction criterion and the D-N model: the role of the background knowledge K precisely corresponds to the role of the auxiliary assumptions A, and the respective deductive entailment relations between E and H are reversed. The following sections investigate whether such a duality can be found in quantitative (read: Bayesian) models of confirmation and explanation as well.

3. Towards a Bayesian Account of Explanatory Power.

Independent of the successes or failures of quantitative analyses of confirmation and explanation, it seems natural to aim at quantitative accounts of degree of confirmation, or explanatory power. First, not all explanations are equally convincing. It would be good to be able to say that hypothesis H explains evidence E better than hypothesis H', instead of just saying that both count as explanations of E. The solution of resilient puzzles (e.g. the "paradox of the ravens" in confirmation theory) often requires the ability to make such comparisons. Second, the history of confirmation theory has taught us that quantitative accounts are able to reveal subtle conceptual nuances, such as the distinction between the absolute and the incremental concept of confirmation, the various ways of understanding and measuring confirmation (e.g. increase in degree of belief vs. discriminative power of the evidence).

A natural candidate for a quantitative analysis is Bayesian modeling. Bayesianism measures how beliefs – or more neutrally: doxastic attitudes – are revised in the face of experience. Degrees of beliefs are modeled as conforming to the probability calculus so that we obtain a powerful tool for mathematical analyses. In other words, a Bayesian framework focuses on the epistemic dimension of explanation or confirmation and ties explanatory/confirmatory power to the extent that a person's convictions are changed by incoming evidence.

Of course, such a move subjectivizes the relationship between an explanans and an explanandum. We have seen that the Hempelian tradition conceived confirmation and explanation as relationships between first-order sentences. Since these sentences are meant to represent states of the material world, and not an agent's doxastic attitude, rational disagreement about whether something counts as confirming evidence, or as a good explanation, is impossible. By tying confirmation and explanation to subjective degrees of belief, Bayesianism breaks with this tradition.

And indeed, it might be argued that the Bayesian approach does not tell us anything about scientific explanation. The epistemic circumstances of giving explanations seem much less important than the exact way how science achieves progress and explains the world. Analyses of explanatory power that focus on a person's subjective degrees of belief are ostensibly missing the point: they are not about explanation in science, but about how scientists change their views. In other words, Bayesians describe *epiphenomena* of scientific explanation rather than explanation itself.

So apparently, confirmation seems to be more suitable for a Bayesian analysis than explanation. But the very same critique has been raised with respect to confirmation, too. Qualitative confirmation criteria did not – similar to Carnap's (1950) logical probability approach – refer to individual agent's psychological states. Instead, they analyzed the logical connections between theoretical and observational statements. This perspective has, with the advent of modern Bayesianism, gradually been replaced by conceiving confirmation as a subjective relationship about how a person's credences about a hypothesis are affected by some evidence. Increase in degree of belief became the central notion. But that move, although familiar to modern philosophers of science, has been far from uncontentious: in crucial episodes from the history of science where prominent theories are confirmed, reconstructions in terms of degrees of beliefs seem to miss the point. Strength of confirmation is, apparently, an *objective* relation between theory and evidence, and not dependent on the prior degrees of belief of any individual scientist, or of the community as a whole. Glymour (1980) gives a couple of incisive case studies, especially from physics.

Nevertheless, the subjectivist research program in confirmation theory has been extremely successful. The paradox of the ravens could be reconstructed and solved by means of a Bayesian account (Fitelson and Hawthorne 2010), the Duhem-Quine problem lends itself to a Bayesian analysis (Earman 1992), and so on. Even more, confirmation theory has connected to the foundations of statistics and interacted with research in statistics that centers around the question about how to measure probabilistic evidence, or goodness of fit between theory and data (Good 1983; Howson and Urbach 1993; Royall 1997). Given that statistics is taking more and more ground in the empirical sciences, the rewards of these advances can hardly be overestimated. Summing up, the power of the Bayesian approach and its interaction with statistical research have justified *post hoc* the identification of confirmation with increase in degree of belief and more than compensated for the loss of accuracy in describing some historical cases. There is certainly more to scientific confirmation than a Bayesian analysis, but this does not imply that the latter is futile.

Hence, not only struggled early Bayesian confirmation theory with similar objections that dawn on the horizon for Bayesian accounts of explanations: the benefits are similar, too. So I suggest to focus our attention on the quantitative dimension of explanation, and the epistemic impact of explanatory considerations. This field is not covered by classical conceptual analysis, but extremely important for science. For instance, if all candidate models of a certain phenomenon are very improbable and data are messy, it might be more interesting to ask for an analysis of the models' explanatory power than to ask which model is confirmed best.[3] Providing a quantitative account of explanatory power thus closes a crucial gap at the intersection of science and philosophy, allows us to compare different cases of explanation, and, by searching for an adequate measure of explanatory power, to get deeper insight into the mechanics of how the concept of explanation works. Since the parallels – or better: dualities – between explanation and confirmation on the qualitative level are very strong, we have reason to suspect that a Bayesian analysis of explanation will have similar success. That does not render the search for a qualitative account of scientific explanation obsolete; e.g. the question whether all explanations are causal or mechanistic can still be meaningfully asked.

Still, the initial question has remained open so far: why should the *epistemic* dimension of scientific explanation, the change in degree of belief, stand central to our analysis? A classical answer, already anticipated in some of Charles S. Peirce's writings, has been articulated by Carl G. Hempel:

> "the explanatory information must provide good grounds for believing that X [the explanandum] did in fact occur; otherwise, that information would give us no adequate reason for saying: 'That explains it – that does show why X occurred.' " (Hempel 1965, 368)

In a more modern formulation, Hempel's idea can be captured in the principle that the explanatory power of an hypothesis relative to some evidence gives us *reason* to believe that the explanandum is true. Imagine someone saying that "X explains Y, but given X, Y is not more expected than before". This would strike us as plainly absurd, and we would refuse to count such an X as an explanation of Y. All explanatory considerations have the normative power to reinforce the belief that the explained proposition is true, and explanatory power measures the degree to which a potential explanans *rationalizes* the explanandum. And it is hard to see how the often-quoted "scientific understanding" that good explanations provide could be different from a rationalization of the explanandum in the light of the explanans. Explanations in statistical sciences, with their use of various goodness-of-fit measures, have already learned that lesson. Moreover, this approach is tailor-made to analyze the role of explanations in raising the posterior probability of a theory, or in appraising one theory over another.

One misunderstanding should be addressed directly: since the idea of Bayesian accounts of explanation consists in measuring explanatory relevance by means of probabilistic relevance, which is a symmetric relations, Bayesian approaches do not seem to get off the ground: the asymmetrical character of explanation, e.g. in the famous flagpole/shadow case, gets lost. On a Bayesian reading, the length of a flagpole's shadow would, presumably, always be explanatorily relevant to the height of the flagpole. Isn't that an absurd outcome and a conclusive counterexample to all Bayesian analyses?

[3] Of course, this claim depends on the specific confirmation measure used.

Those who believe that there is a problem assume, however, the wrong premise that the Bayesian analysis is a bidirectional conceptual analysis of the meaning of explanatory power, in the sense of specifying a set of necessary and sufficient conditions. However, it is meant as a *unidirectional* analysis of the epistemic impacts of explanation on the degrees of belief of a rational agent. We are only committing ourselves to the fact that cases of genuine explanation have to yield a positive degree of explanatory power, but not to the contrary (e.g. when reserving the flagpole/shadow case). Hence, this objection does not prevent us from pursuing our explicative project in a Bayesian framework. Considering how fruitful a probabilistic analysis of confirmation has proven, and how similarly the both concepts have been analyzed on a qualitative level, a Bayesian explication of explanatory power is overdue, at least for the rapidly growing number of statistical explanations in science.

4. Measuring Explanatory Power, and Concerns Addressed.

It might be questioned whether our Bayesian approach has substantial advantages over, say, Hempel's (Hempel 1965) inductive-statistical (I-S) or Salmon's (1971/1984) statistical-relevance (S-R) approach to explanatory power. Hempel's quantification of the strength of an I-S argument (Hempel 1965) seems to anticipate a Bayesian measurement of explanatory power, and Salmon, on the other hand, takes probabilistic (=statistical) relevance as the main idea of his proposal, just as modern Bayesians do.

There is, however, a crucial difference: both Hempel and Salmon lean on relative frequencies in relevant reference classes as the measures of explanatory strength. In the case of Salmon, this is motivated by the desire to connect an account of statistical explanations to relative frequencies in nature, whereas Hempel and his co-author Oppenheim aim at connecting their account of explanatory power to the Carnapian framework of logical probability. Thus, they are not giving a subjective Bayesian account, as this paper proposes.

Hempel and Oppenheim even explicate the "systematic (=explanatory, predictive) power" S of a theory with respect to some evidence as a notion that is *dual* to the theory's logical probability P: theorems about S can be translated into theorems about P by swapping the logical connectives '$\wedge$' and '$\vee$' in the arguments of those functions (Hempel and Oppenheim 1948, 173). This might be reconstructed as a duality between "confirmedness" of a theory (=the absolute concept of confirmation), and the predictive power of a theory. Now, it is natural to suspect that a Bayesian measure of explanatory power could be dual to the concept of incremental confirmation, i.e. increase in degree of belief. For this purpose, we have to define the way we would like to measure explanatory power.

A Bayesian measure of explanatory power $\mathcal{E}$ is based on a probability space $(\Omega, \mathcal{A}, p(\cdot))$ and is defined as a measurable function from three propositions $E, H, K \in \mathcal{A}$ to a real number $\mathcal{E}(E, H, K)$. Conventionally, one assumes that if hypothesis H is positively explanatory relevant to explanandum E, then $\mathcal{E}(E, H, K) > 0$, and vice versa for negative explanatory relevance. If H is said to be explanatorily irrelevant to E, relative to K, then $\mathcal{E}(E,H,K)=0$.

Keeping in mind the basic idea that the explanans rationalizes the explanandum, there are a couple of measures that have or could been suggested to measure explanatory power, often starting from the likelihood of the explanandum under the explanans $p(E|H.K)$. Then, one compares this term to the explanandum's prior expectedness $p(E|K)$ $(\mathfrak{d}, \mathfrak{r})$, or to its likelihood under alternative potential explanations $p(E|H.K)$ $(\mathfrak{s})$, and so on (Eells 1991; McGrew 2003; Schupbach and Sprenger 2010):

$$\mathfrak{d}(E,H,K) = p(E|H.K) - p(E.K) \qquad \mathfrak{s}(E,H,K) = p(E|H.K) - p(E|\neg H.K)$$
$$\mathfrak{r}(E,H,K) = \log\frac{p(E|H.K)}{p(E.K)} \qquad \mathfrak{k}(E,H,K) = \frac{p(H|E.K) - p(H|\neg E.K)}{p(H|E.K) + p(H|\neg E.K)}$$

These measures have well-known analogues in confirmation theory (Fitelson 2001), obtained by swapping the arguments:

$$d(E,H,K) = p(H|E.K) - p(H.K) \qquad s(E,H,K) = p(H|E.K) - p(H|\neg E.K)$$
$$r(E,H,K) = \log\frac{p(E|H.K)}{p(E.K)} \qquad k(E,H,K) = \frac{p(E|H.K) - p(E|\neg H.K)}{p(E|H.K) + p(E|\neg H.K)},$$

so that there is also a duality on the level of modern Bayesian measures of degree of confirmation and explanatory power. Given the previous discussion, this result may not be unexpected, but it is quite important as it vindicates the structural analogy between explanation and confirmation on the quantitative level, too.

The plurality of candidate measures introduces the problem of *measure sensitivity* – some statements that we would like to make about the concept analyzed (here: explanatory power) depend on the choice of our measure. In this paper, I cannot go into the details of the arguments which measure should be preferred: suffice to say that I have elaborated an argument for $\mathfrak{k}$ in another place [reference omitted for blind review].[4] But be this as it may, all measures are based on probabilistic relevance, and as such, they all entail that sound explanations give us reason to believe that the explanans is true, too. Successful explanations confirm the explanans in the Bayesian sense. This natural consequence of a Bayesian model of explanatory reasoning explains why we like to infer to explanatory successful hypotheses, and affects all accounts of theory choice and appraisal where the explanatory power of a theory is an important criterion. In

[4]This measure gives rise to several instructive theorems, and it is also derived from a parsimonious set of adequacy conditions: sensible technical demands, uncontentious judgments about explanatory irrelevance, maximal explanatory power, and the relationship between positive to negative explanatory power. Moreover, it is dual to the Kemeny-Oppenheim measure k which has been defended by Fitelson (2001) and which is ordinally equivalent to the likelihood ratio measure that plays an outstanding role in statistical inference.

other words, we can give a Bayesian reconstruction of how explanatory successes increase the credibility of a hypothesis, and for all measures that are not fully symmetric, such as $\mathfrak{r}$, we can also state the extent to which explanatory power and degree of confirmation differ from each other. Note that this property does *not* commit us to a theory of abductive inference, such as Inference to the Best Explanation, only to acknowledging the normative pull of good explanations.

Finally, we must defend our parallels between confirmation and explanation against the criticism that only *causal* relations truly explain. Probabilifying explanation in terms of subjective degrees of belief apparently neglects that genuine scientific explanations reveal causal mechanisms whereas, clearly, not all probabilistic correlations count as explanatory. My reply to this objection is threefold.

First, the unidirectional nature of our Bayesian analysis does not force us to accept that all correlations are explanatory – rather, explanatory relations will be mirrored in positive correlations. Hence, there is no logical tension. Second, there is no need to adopt an overly narrow concept of what counts as a sensible explanation. As has been pointed out in the past, this will be highly context-sensitive anyway (Van Fraassen 1980), and even without causal relations, theories can explain the data by subsuming them in an illuminating way. Newton's law of gravitation with its instantaneous "action at a distance" is a prime example of an explanatory highly successful covering law. Yet, it is hard to interpret causally as long as causal processes are believed to operate *locally*. Explanation has, much more than causation, a connotation of rationalizing evidence by means of positing a hypothesis, and this justifies a "subjective" understanding in terms of degrees of belief vis-à-vis an "objective" understanding in terms of causal relations. Subjectivism with respect to scientific explanation – as opposed to objectivism about causal relations – is not a threat, but rather a natural way of keeping the different functions of both concepts apart.

Third and last, there is a natural way to integrate a Bayesian analysis of explanatory power into a particular way of thinking of causation, namely into a probabilistic analysis. For instance, statisticians infer from probabilistic relevance to a cause-effect-relation when all possible further causes of the effect are under control. Conversely, the Causal Markov Condition explicates causal screening-off by means of probabilistic independence conditional on all common causes. Finally, probabilistic relevance under a varying set of conditions forms the core of several philosophical analyses of causation, e.g. Eells (1991). On such an account, positive probabilistic relevance under all relevant boundary conditions – causation – would give rise to an "ideal explanation", an explanation independent of the contingent configuration of things in the world and our limited knowledge thereof. Whereas for normal explanations, it would be sufficient that the explanandum is rationalized by the explanans given our *actual* background knowledge. Understanding causation as an ideal form of explanation is, to my mind, a much more fruitful perspective than letting them coincide. Then, it remains an interesting open question how measures of causal strength (Eells 1991; Fitelson and Hitchcock 2010) can be related to the starting debate about measures of explanatory power (McGrew 2003; Schupbach and Sprenger 2010).

5. Conclusions.

The project of explicating explanation qualitatively has suffered from many setbacks, caused by a continued strain of objections. To resolve this deadlock, I have proposed to shift attention to the quantitative, epistemic dimension of scientific explanations, in the same way that it has been done for confirmation. A Bayesian approach is the natural way of doing so. This move offers a much better and fine-grained understanding of the grammar of explanation, and has the ability to connect to the growing statistical literature on measures of confirmation and goodness-of-fit.

To motivate the project further, I argued at length for the structural duality of explanation and confirmation, where Bayesianism has been a very fruitful research program. This duality can be sustained and qualified when passing from a qualitative to a quantitative, Bayesian framework. Shifting the focus to the epistemic dimension of successful explanations, to the rationalization of the explanandum by the explanans, the paper does not only illuminate the relationship between confirmation and explanation, but also makes a contribution to better understanding explanatory reasoning in science, and its role in theory choice and appraisal. A subjectivist, Bayesian analysis of explanatory power in terms of rational degrees of belief is not far-fetched and futile, but fruitful, feasible, and reveals exciting structural connections between causation, explanation and confirmation.

References

Carnap, Rudolf (1950): *Logical Foundations of Probability.* Chicago: University of Chicago Press.

Carnap, Rudolf (1966): "Explanation in Science and History", in: R. G. Colodny (ed.), *Frontiers of Science and Philosophy*, Pittsburgh: University of Pittsburgh Press.

Earman, John (1992): *Bayes or Bust?*. Cambridge/MA: The MIT Press.

Eells, Ellery (1991): *Probabilistic Causality.* Cambridge: Cambridge University Press.

Fitelson, Branden (2001): *Studies in Bayesian Confirmation Theory.* Ph.D. thesis, University of Wisconsin-Madison. Available at `www.fitelson.org`.

Fitelson, Branden (2008): "Goodman's 'New Riddle'", *Journal of Philosophical Logic* 37, 613-643.

Fitelson, Branden, and James Hawthorne (2010): "How Bayesian Confirmation Theory Handles the Paradox of the Ravens", in: Ellery Eells and James Fetzer (ed.), *Probability in Science*, Chicago: Open Court.

Fitelson, Branden, and Christopher Hitchcock (2010): "Probabilistic Measures of Causal Strength", in: F. Russo and J. Williamson (ed.), *Causality in the Sciences*, Oxford: Oxford University Press. Available at `www.fitelson.org`.

Glymour, Clark (1980): *Theory and Evidence.* Princeton: Princeton University Press.

Good, I. J. (1983): *Good Thinking: The Foundations of Probability and Applications.* Minneapolis: The University of Minnesota Press.

Goodman, Nelson (1983): *Fact, Fiction and Forecast.* Fourth Edition. Harvard: Harvard University Press.

Hempel, Carl G. (1945): "Studies in the Logic of Confirmation", originally appeared in *Mind*, reprinted in: *Aspects of Scientific Explanation*, 3-46. New York: Free Press.

Hempel, Carl G. (1965): *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science.* New York: Free Press.

Hempel, Carl G., and Paul Oppenheim (1948): "Studies in the Logic of Explanation", *Philosophy of Science* 15, 135-175.

Howson, Colin and Peter Urbach (1993): *Scientific Reasoning: The Bayesian Approach.* Second Edition. La Salle: Open Court.

McGrew, Timothy (2003): "Confirmation, Heuristics, and Explanatory Reasoning", *British Journal for the Philosophy of Science* 54, 553-567.

Royall, Richard (1997): *Statistical Evidence: A Likelihood Paradigm*. London: Chapman & Hall.

Ruben, David-Hillel (1990): *Explaining Explanation*. London: Routledge.

Salmon, Wesley (1971/1984): "Statistical Explanation", reprinted in Salmon (1984), 29-87.

Salmon, Wesley (1984): *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

Schupbach, Jonah, and Jan Sprenger (2010): "The Logic of Explanatory Power", manuscript under review, available at `www.laeuferpaar.de`.

Van Fraassen, Bas (1980): *The Scientific Image*. Oxford: Clarendon Press.

Woodward, James (2003): *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.