This is the author's final version of the paper which was published in Synthese (2007) vol. 157, pp. 275-296. The original publication is available at
http://dx.doi.org/10.1007/s11229-006-9055-7

Coherence Measures and Inference to the Best Explanation

David H. Glass

School of Computing and Mathematics
University of Ulster
Newtownabbey
Co. Antrim
BT37 0QB
United Kingdom
E-mail: dh.glass@ulster.ac.uk

Abstract. This paper considers an application of work on probabilistic measures of coherence to inference to the best explanation (IBE). Rather than considering information reported from different sources, as is usually the case when discussing coherence measures, the approach adopted here is to use a coherence measure to rank competing explanations in terms of their coherence with a piece of evidence. By adopting such an approach IBE can be made more precise and so a major objection to this mode of reasoning can be addressed. Advantages of the coherence-based approach are pointed out by comparing it with several other ways to characterize 'best explanation' and showing that it takes into account their insights while overcoming some of their problems. The consequences of adopting this approach for IBE are discussed in the context of recent discussions about the relationship between IBE and Bayesianism.

1. Introduction

It seems obvious that science progresses by offering good explanations of features of the natural world. The success of science is not attributed merely to the discovery of previously unknown or unexpected phenomena, but rather to the fact that science can explain the phenomena in question. In many cases a number of competing explanations (or hypotheses) are proposed and it is the task of scientists to determine which offers the best explanation. Once this has been achieved, an inference can be made to the truth (or approximate truth) of the best explanation. This form of reasoning is referred to as inference to the best explanation (IBE) and is often identified with C. S. Peirce's concept of abduction[1]. Modern discussions of IBE date back to Harman (1965), while more recently Lipton (2004) has presented an exposition and defence of IBE. Adherents of IBE often point out that IBE, or something very like it, is used in many other areas of academic enquiry as well as in ordinary life. One such area is in debates about scientific realism where advocates of realism, such as Psillos (1999), often rely on IBE.

Although IBE seems extremely plausible, it has received serious criticism. One of the difficulties arises from the fact that IBE, while sounding good as a slogan, requires at least a reasonably precise formulation if it is to be taken seriously as a form of reasoning. This presents a substantial challenge to the defender of IBE since it is far from clear how the terms 'explanation' and 'best' (and indeed 'inference') can be made precise. The goal of this paper is to provide an account of 'best' in the sense that a measure will be proposed which can be used to compare and rank a number of competing explanations.

There has been considerable attention given to clarifying the notion of explanation (particularly scientific explanation) within the literature (see, for example, Salmon (1989, 1998) and references therein) and so it might be thought that any suitable account of explanation would automatically provide an account of how explanations can be compared. It would seem likely that this would be the case if a satisfactory account of probabilistic explanation could be provided in purely statistical terms. The reason for this is that whatever statistical measure is used to characterize explanation could also be used to compare competing explanations. Hempel's inductive-statistical account is a case in point (see section 2). Many counterexamples, however, have shown that such an account will not work and that more than probability is required to characterize explanation. Salmon (1998), for example, argues that an adequate account of explanation must take causality into account and, while statistical relationships may be important, this is because they provide evidence for an underlying causal relationship.

The above discussion suggests that two components are required for a full account of explanation:

a) an account of what *constitutes* an explanation, and
b) a suitable *methodology for comparing* competing explanations.

In this paper I propose to draw on recent work on probabilistic accounts of coherence in order to meet requirement (b) by providing a measure to rank explanations. In recent years a number of probabilistic accounts of coherence have been proposed (see section 3 for further details) and the implications for the coherence theory of justification investigated (Bovens and Olsson 2000; Olsson 2002, 2005; Bovens and Hartmann 2003a, b). Since there has been general agreement that coherence on its own does not result in a high likelihood of truth, the focus has been on the question of whether coherence is *truth conducive* so that more coherence gives rise to a higher likelihood of truth. Olsson (2005) presents an impossibility theorem to the effect that there is no truth conducive coherence measure. Bovens and Hartmann (2003b) also present an impossibility result, but argue that its impact on coherentism can be circumvented by adopting a partial ordering of information sets on the basis of coherence, i.e. in some cases one set can be identified as more coherent than another while in other cases no such comparison is possible. Clearly, these claims are extremely significant for any work on coherence and so they will be considered in section 3.

There are several ways in which the discussion of coherence in this paper differs from the work noted above. Much of the work on coherence has considered it in terms of reports, such as witness reports, which are taken to provide evidence for the facts they assert. The problem has then been to determine whether greater coherence of the facts asserted results in a higher posterior probability of those facts given the reports. Note that coherence is considered between the facts asserted, rather than *between* facts asserted and the evidence for them. By contrast, in this paper coherence is considered as a relation between an hypothesis and the evidence for it. Consequently, coherence relates the hypothesis to something which is already known to be true. In such a picture the coherence theory of justification is not in view. In fact, it might seem that this way of looking at things fits better with foundationalism since coherence is used to build on facts which are known in some other way. This relates to a further point. In the worked described above the main motivation was to find a coherence measure that would meet the demands of the coherentist programme. By contrast, the motivation here is to find a measure of coherence which matches our intuitive understanding of the concept and to investigate how such a conception might relate to explanation. In short, the connection between the notions of explanation and coherence is established by noting that a condition for a satisfactory account of the relation "… better explanation than …" proposed in section 2 turns out to be essentially the same as a plausible condition for the relation "… more coherent than …".

The idea being proposed is not that explanation can be understood in terms of probabilistic coherence, since this would be ignoring the literature which suggests that probabilistic relationships are inadequate in this respect. Rather, the proposal is that probabilistic coherence can be used to determine the goodness of explanations and hence to compare them. In many ways this work is similar to the approach advocated by Thagard (1989) where explanatory coherence is used in IBE. Thagard, however, uses a connectionist algorithm to determine coherence, whereas in this paper a probabilistic account is advocated. Nevertheless, this work agrees with Thagard in identifying the best explanation as the one which best coheres with the explanandum.

Even if this work presents an adequate account of 'best explanation', this in no way establishes the validity of IBE. Van Fraassen (1980, 1989) has set out a number of extremely significant criticisms of IBE. Psillos (1996, 1997) has criticized the arguments from 'the bad lot' and 'indifference' against IBE (see also the response from Ladyman et al (1997)), but in this paper another argument of van Fraassen's will be considered: the Dutch Book argument. It may be the case that providing a precise account of best explanation is enough to refute IBE since the resulting form of inference can be shown to be irrational. Okasha (2000) responds to van Fraassen by arguing that IBE and Bayesianism are compatible. This view is set out in a detailed and illuminating way by Lipton (2001, 2004) who argues that explanatory considerations might lubricate the Bayesian mechanism by helping to determine likelihoods and priors and in deciding what evidence is relevant in a given context. In response Salmon (2001) is somewhat sceptical about reconciling Bayesianism and IBE, but he does claim that a case can be made for the role of explanatory virtues in the assessment of priors. A significant area of agreement between Lipton and Salmon is that scientists are often more interested in fertile hypotheses with greater informational content even at the expense of high probability. I will return to this issue in section 5. Psillos (2004) adopts a rather different stance on the Bayesian-IBE issue since, although he wishes to defend IBE, he has reservations about trying to cast it within a Bayesian framework. His main points are that IBE is ampliative while Bayesian is not and that IBE is warrant-conferring rather than depending completely on Bayesianism for confirmation.

The structure of the paper is as follows. In section 2 I consider several straightforward ways to define 'best explanation' and point out problems with them. An analysis of these proposals gives rise to a necessary condition for any satisfactory account. Section 3 presents a discussion on coherence measures and the impossibility results, focussing particularly on information pairs. Section 4 provides a link between coherence and explanation, uses one of the coherence measures to rank explanations and compares this approach with those discussed in section 2. The implications for IBE are investigated in section 5 by considering van Fraassen's argument and two possible responses, one of which involves exploring the connection between IBE and Bayesianism.

2. Problematic accounts of 'best' explanation

In attempting to provide a methodology for comparing competing explanations it is worth noting Hempel's distinction between *potential* and *actual* explanations (Hempel, 1965). An actual explanation is one which, as a matter of fact, explains the explanandum in question. A potential explanation is one which, if true, would be an actual explanation. This section compares three approaches for comparing potential explanations. In this context the goal of IBE can be understood as selecting the actual explanation from the potential explanations. Different forms of IBE arise from which of the three approaches is used to select the 'best' explanation from the potential explanations. [2]

2.1 Maximum likelihood (ML)

Hempel's (1965) account of inductive-statistical (I-S) explanation provides the natural place to start when looking for a way to compare explanations (or hypotheses) for a particular explanandum, E. The main point to note here is that for an I-S explanation the explanandum must be highly probable given the explanation. It is not appropriate here to reiterate the many criticisms of Hempel's account. However, even though these criticisms may be fatal to I-S explanation as an account of explanation perhaps the idea that a good explanation yields a highly probable explanandum is relevant here.

Suppose there are a number of competing hypotheses, H_i, for a piece of evidence, E. Furthermore, suppose that under some account of explanation (such as a causal account) all of the H_i are legitimate explanations even though $P(E|H_i)$ may be small for some of the H_i. It still seems reasonable to conclude that hypotheses with greater likelihood, $P(E|H_i)$, are *ipso facto* better explanations. These considerations yield a simple way to determine the best explanation. The best explanation is the hypothesis H_i with the maximum value of $P(E|H_i)$, i.e. the maximum likelihood.

An obvious problem with this approach is that nothing other than the likelihood is taken into account in assessing the goodness of an explanation and it is doubtful whether this is consistent either with scientific practice or considerations of explanation in ordinary life. To take an example from the history of science, consider the Copernican revolution. The Ptolemaic system was able to account for the same evidence as the Copernican system, but other considerations such as simplicity favoured the latter. To take a more mundane example, suppose that two different diseases would account equally well for a patient's symptoms. Does this mean that they are both equally good explanations? Suppose one disease is extremely rare, while the other is extremely common. Surely the more common disease would provide the better explanation. This example suggests that the prior probabilities of the hypotheses might also have a role to play as does the previous example if simplicity increases the prior probability. This point is closely related to the so-called base-rate fallacy. Psillos (2004), who discusses this issue in the context of IBE, points out "… the base-rate fallacy (no matter how one reads it) shows that it is incorrect to just equate the best explanation of the evidence with the hypothesis that has the highest likelihood."[3]

It is important to note that this approach would not seem to be a suitable one for a defender of IBE to adopt. The reason for this is that there is no good reason for thinking that an hypothesis with a high likelihood will also have a high posterior probability unless it also has a high prior probability. And of course a high posterior probability would seem to be required for IBE since the whole point of IBE is that the best explanation is likely to be true. Thus, we might expect that in many cases IBE, understood as inference to the hypothesis with the maximum likelihood, will not be a good approach for finding true (or highly probable) hypotheses.

2.2 The most probable explanation (MPE)

As noted above, one diagnosis of the problem with the ML approach is that it does not take the prior probability into account. Bayes' theorem provides a simple way to rectify the problem. If we now consider the posterior probability of an hypothesis, $P(H_i|E)$, which is the product of the likelihood, $P(E|H_i)$, and the prior probability, $P(H_i)$, of the hypothesis (divided by $P(E)$). Thus, in this case the best explanation can be identified as the hypothesis which is most probable given the evidence and so has the maximum value for $P(H_i|E)$. This is essentially the Bayesian approach to inference, but note that here it is being used to identify the best explanation. It is worth noting that this approach is the one usually adopted in artificial intelligence, particularly in Bayesian networks (see, for example, Pearl (1988) and Shimony (1994) where it is referred to as the most probable explanation (MPE) or maximum *a posteriori* (MAP).

A problem with this approach is that an hypothesis could turn out to be the best explanation even if the evidence is extremely unlikely given the hypothesis. For example, in many cases an hypothesis which lowers the probability of the evidence relative to the unconditional case, ie $P(E|H_i) < P(E)$, will turn out to be the best explanation according to the MPE approach because it has a high prior probability. While it may not be desirable to rule out the possibility that the evidence could be negatively dependent on the best explanation[4], it still seems reasonable to say that such dependence should count against the hypothesis more than it does in the MPE approach. The main drawback with the ML approach was that it ignored prior probabilities, but it seems that the MPE approach overcompensates.

The MPE approach seems like the ideal account of best explanation for a defender of IBE. By selecting the best explanation, the most probable explanation is automatically selected. Thus, the connection between IBE and truth (or probable truth) is guaranteed. However, this success comes at a high price indeed since the 'best explanation' has been defined as the 'most probable explanation' and this seems to make IBE trivial. This point has been emphasized by Lipton (2004) who draws the distinction between the loveliest explanation and likeliest explanation. To quote Lipton, "We want a model of inductive inference to describe what principles we use to judge one inference more likely than another, so to say that we infer the likeliest explanation is not helpful" (Lipton 2004, 60). It seems that the goal for defenders of IBE is to give an account of 'best explanation' in terms of loveliness and show that a feature of such an explanation will be its likeliness, i.e. high posterior probability.

These considerations along with the discussion in section 2.1 highlight a fundamental problem for IBE. Once the link between the best explanation and the most probable explanation has been severed there is no guarantee that it can be re-established. The ML approach in section 2.1 provided an obvious way to account for the loveliness of an explanation, but as we have seen it does not follow that lovely explanations will also be likely explanations.

2.3 A conservative Bayesian approach (CB)

Despite the criticisms that have be levelled against the ML and MPE approaches they do nevertheless capture some of the features which we might expect good explanations to have. Furthermore, in many cases they yield the same conclusion as to which explanation is best. It is worth looking at Bayes' theorem in order to compare the approaches and see if they can be brought together in a more satisfactory way. The MPE and ML measures of goodness are related by

$$P(H_i \mid E) \quad = \quad \frac{P(E \mid H_i)}{P(E)} \quad \times \quad P(H_i) \tag{1}$$

and so MPE offers a particular strategy for incorporating likelihood and prior probability. One strategy that a defender of IBE might attempt is presented by Okasha (2000). Bayesian inference could be adopted so that expression (1) is used to infer the most probable explanation, but this need not be equated with the best explanation.[5] Okasha's point is that explanatory considerations come into play when making the inference since the likelihood, $P(E|H_i)$, and prior probability, $P(H_i)$, are taken into account.

Why does Okasha stop short of adopting the MPE definition of 'best'? Perhaps this is because objections can be made along the lines indicated in section 2.2. Okasha's point seems to be that a more precise account of 'best' is not required since a defence of IBE can remain neutral on questions such as the ranking of explanations. It is rather that qualitative features of explanation can come into play in making inferences and, since this can be accounted for in Bayesian terms, IBE is not defeated. However, this does not seem quite right. Surely IBE is committed to the idea that an inference should be made to the *best* explanation and this seems rather hollow unless there is a way to rank explanations.

Perhaps another strategy could be adopted for making 'best explanation' more precise without adopting the problematic MPE approach. Okasha certainly seems to accept that some explanations are better than others. Consider Okasha's example of the child taken to the doctor by his mother. The doctor forms two competing hypotheses on the basis of the mother's information: the child has pulled a muscle (hypothesis, H_1), and he has torn a ligament (hypothesis, H_2). Using IBE the doctor tentatively accepts H_2 on the grounds that the prior probability of a torn ligament is greater than that of a pulled muscle ($P(H_2) > P(H_1)$) and the symptoms E are exactly what would be expected if the child had a torn ligament, but not if he had pulled a muscle ($P(E|H_2) > P(E|H_1)$).

Chapjewska and Halpern[6] (1997) provide a way of making this insight precise (see also van Fraassen 1980, chap. 2). With each explanation H_i of evidence E associate the pair of numbers $\{P(E|H_i), P(H_i)\}$. A partial ordering can now be defined on the explanations of E by taking H_1 to be a better explanation than H_2 if and only if,

$$P(E \mid H_1) \quad > \quad P(E \mid H_2) \qquad and \qquad P(H_1) \quad > \quad P(H_2)\,. \tag{2}$$

This seems to provide a good solution since now some explanations can be ordered and, among those that can, the better the explanation the higher its posterior probability. And this can be achieved in such a way that no commitment as to which is the better explanation need be made in cases where H_1 has a greater likelihood, but a lower prior probability than H_2 which gave rise to the problems in section 2.2. This will mean, of course, that there are cases where explanatory considerations will not tell us which is the best explanation, but this does not seem problematic since it is not a requirement of IBE that it can account for all inferences. Let us call this approach the conservative Bayesian approach (CB) since it does not provide a total ordering.

Unfortunately there is a major problem with this approach. As noted earlier there are many cases where the ML and MPE approaches agree as to which is the best explanation. For example, in the case of two competing explanations, it is easy to see from Bayes' theorem in equation (1) that they will agree in all cases where both conditions in (2) are satisfied. However, the ML and MPE approaches will also agree in all cases where the priors of the competing explanations are equal and in many cases where the explanation with the greater likelihood has a lower prior. In such cases the CB approach will fail to rank the explanations and so it appears to be too conservative. I therefore propose a minimal requirement that should be a feature of any account of 'best explanation' if it is to replace ML or MPE:

> Any account of best explanation should give the same result as the ML and MPE approaches whenever these approaches give the same result as each other.

This requirement can also be formulated as a condition for the relation "… better explanation than …", which will play a central role in connecting coherence and explanation in section 3. The condition can be set out as follows,

Explanation Ranking condition:

> For two explanations, H_1 and H_2, of a piece of evidence E, if $P(E|H_1) > P(E|H_2)$ and $P(H_1|E) > P(H_2|E)$ then H_1 is a better explanation of E than H_2.

It is easy to see that the CB approach does not satisfy this condition. Consider again the example of the child being brought to the doctor. Let us suppose that the probabilities are as follows: $P(H_1) = 0.6$, $P(H_2) = 0.4$, $P(E|H_1) = 0.5$, $P(E|H_2) = 1$. Bayes' theorem gives $P(H_1|E) = 0.43$ and $P(H_2|E) = 0.57$. Thus, the ML and MPE approaches agree that H_2 is the better explanation, but because H_2 has a lower prior probability than H_1 the CB approach does not provide an ordering. A simple way to overcome this problem would be to provide a partial ordering based on when ML and MPE agree, but a more unified way to integrate the approaches is considered in Sect. 4.

3. Coherence measures

Intuitively it seems that the concept of coherence might be related to IBE. As Psillos (2002, 619) points out, "this coherence-enhancing role of IBE … is ultimately the warrant-conferring element of IBE." Roughly speaking the coherence of a set of beliefs describes how well they fit together or to what extent they support each other and this seems to be what is required for IBE. The scientist, the doctor and the detective could all be described as trying to find the hypothesis which best fits the evidence available. The difficulty here is that we are trying to give a reasonably precise account of IBE and to appeal to coherence seems fraught with difficulties unless a precise account of coherence can be given. This section explores some recent work on coherence, including the impossibility results, but the main focus is on aspects of the debate that will turn out to be relevant for ranking explanations.

3.1 Defining coherence

Suppose that a number of independent witnesses give testimonies which are in agreement with each other. Even if the witnesses are not completely reliable, the fact that their testimonies agree or cohere with each other may well lead us to accept their claims. For Lewis "congruence [coherence] of the reports establishes a high probability of what they agree upon" (Lewis 1946, 346) even if the antecedent probability was low. Bonjour goes further by stating that "no antecedent degree of warrant or credibility is required" and claims that this will be the case even if the individual reports "are much more likely to be false than true" (Bonjour 1985, 148).

A serious problem for coherentism is that it is remarkably difficult to give a satisfactory account of the notion of coherence. However, since it is generally agreed that coherence comes in degrees and is concerned, at least partially, with the extent to which beliefs agree with each other it is tempting to think that a probabilistic account can be given. In the case of two beliefs Bovens and Olsson (2000) set out in probabilistic terms what they describe as a minimal sufficient condition for the relation "… more coherent than …",

Bovens-Olsson condition:

> For an information pair {A,B} and probability distributions P and P′, if $P(A|B) > P'(A|B)$ and $P(B|A) > P'(B|A)$, then {A,B} is more coherent on probability distribution P than on probability distribution P′ (2000, 688).

This condition brings out the idea of agreement between beliefs in a very clear manner since it is the probability of each belief given the other that is all important. It also captures the extreme cases well since A and B will be more coherent on a probability distribution P for which $P(A|B)= P(B|A) = 1$ than for a distribution where this is not the case and, similarly, they will be less coherent on a probability distribution P′ for which $P'(A|B)= P'(B|A) = 0$ than for a distribution where this is not the case. Despite the plausibility of the condition, however, Bovens and Olsson also point out that,

> Even this minimal condition is not entirely innocent. Let there be a roulette wheel with one hundred numbers and an equal chance for each number to be the winning number (WN). The wheel is spun. On scenario a, Joe says that WN is 49 or 50 and Amy says that WN is 50 or 51. On scenario b, Joe says that WN is 1, 2, …, or 70 and Amy says that WN is 31, 32, …, 100. Then, on our minimal condition, the information pair S ={[Joe is correct], [Amy is correct]} is more coherent on P^b than on P^a, which seems counterintuitive (2000, 688-9).

Bovens and Olsson go on to point out that one response to this is to say that coherence can be thought of as either agreement or as striking agreement. Thus, in the example the information pair S is more coherent (in the sense of agreement) on P^b than on P^a, but is more coherent (in the sense of striking agreement) on P^a than on P^b. According to this way of looking at coherence the Bovens-Olsson condition applies to *coherence as agreement.*

One probabilistic measure of coherence appearing in the literature (Shogenji 1999) gives the coherence of two beliefs A and B to be $P(A \wedge B)/P(A)P(B)$. This measure yields a value of 0 in the case where A and B are logically inconsistent, a value of 1 when they are probabilistically independent and greater than 1 when there is a positive probabilistic dependence between them. However, suppose that A and B are logically equivalent and that $P(A) = P(B) = p$. In this case Shogenji's measure yields a value of $1/p$ which seems counterintuitive since logical equivalence seems to be a case where the coherence should be maximal. An alternative probabilistic coherence measure proposed by Fitelson (2003) is based on a measure of support and yields values in the interval [-1,1] with the value of -1 obtained if A and B are logically inconsistent, 0 if they are probabilistically independent and 1 if they are logically equivalent. Thus Fitelson's measure has a number of characteristics which make it a plausible measure of coherence. A rather different approach is adopted by Bovens and Hartmann (2003), who propose a coherence measure which depends on the reliability of the sources of the information. Although this seems counterintuitive, their coherence measure induces a partial ordering of information sets which does not depend on the reliability of the sources.

While much more could be said about the advantages and disadvantages of the various coherence measures discussed above, an important feature shared by all of these measures is that they fail to satisfy the Bovens-Olsson condition.[7] In the case of Shogenji's and Fitelson's measures this has been considered in another paper (Glass 2005)[8], where the relationship between the Bovens-Olsson condition, dependence on priors and inclusion of probabilistic (in)dependence is also considered. The fact that Shogenji's and Fitelson's measures have a neutral point (1 in Shogenji's measure and 0 in Fitelson's measure) where the beliefs in question are probabilistically independent results in the failure to satisfy the Bovens-Olsson condition. By contrast, another measure of coherence, which has appeared in the literature (Olsson 2002; Glass 2002), does satisfy the Bovens-Olsson condition. According to this measure the coherence of two beliefs A and B is given by,

$$C(A,B) = \frac{P(A \wedge B)}{P(A \vee B)}, \tag{3}$$

provided $P(A \vee B) \neq 0$. The idea that coherence is given by the probability of the conjunction divided by the probability of the disjunction captures our intuition that coherence is concerned with the degree of agreement between the two beliefs. Note that the coherence is a measure on the interval [0,1] with $C(A,B) = 0$ occurring when $P(A|B) = P(B|A) = 0$ and $C(A,B) = 1$ when $P(A|B) = P(B|A) = 1$. This means that logically equivalent beliefs are maximally coherent (eg $C(A,A) = 1$), while logically inconsistent beliefs are incoherent (eg $C(A,\neg A) = 0$) and so the extreme deductive cases are treated appropriately. In contrast to the other probabilistic measures of coherence expression (3) depends only on the extent of agreement between two beliefs rather than how probable those beliefs are in the fist place. More precisely, for fixed values of the conditional probability the coherence measure is independent of prior probabilities (see Glass (2005) for further discussion of this point). Consequently, it satisfies the Bovens-Olsson condition.

The main point of this section has been to outline some important differences between several accounts of coherence, rather than to argue that one gives the correct account while the remaining accounts fail. Rather, even if the different accounts do justice to different intuitions regarding the nature of coherence, they are not equally suitable for ranking explanations. The Bovens-Olsson condition is crucial in this respect as will be explored further in Sect. 4.

3.2 The Impossibility Results

Since the goal of this paper is not to assess the coherence theory of justification, the impossibility results do not impinge directly on the argument. Thus, even if the coherence measure C defined in (3) is not truth conducive in the required sense, it might still be useful for ranking explanations. However, some points are noted here which might provide some further motivation for the use of C as a coherence measure.

A central question in assessing the truth conduciveness of coherence measures concerns what should be kept fixed between information sets that are being compared, i.e. what should be included in the *ceteris paribus* clause. Bovens and Hartmann think that the prior joint probability should be kept fixed, while Olsson thinks it should not. The argument presented by Olsson is based on the idea of independent variation so that a quantity should be kept fixed if it can be varied independently of the quantity under consideration. While this argument is convincing when applied to strength in the case of Shogenji's measure, it is much less plausible for the prior probability in the case of the C measure. The reason for this is that the prior probability and coherence are independent in a well-defined sense for this measure (see Glass (2005) for discussion) so that fixing the prior does not place any restriction on coherence. Thus it seems that in the case of the C measure the prior should be included in the *ceteris paribus* clause.

There is, however, a further problem since Bovens and Hartmann as well as Olsson think that the size of the information sets under consideration should not be kept fixed. It is certainly plausible to think that comparisons should be permitted across sets of different size. However, as Olsson points out, fixing the set size does not restrict the values of coherence, although he does note that size and coherence are not completely separable. This is a particularly important issue for the C measure since the main objection to C as a coherence measure is due to its dependence on the size of the information set being considered. The obvious extension of C to the case were more than two beliefs are being considered is to take the ratio of the probability of the conjunction divided by the probability of the disjunction. This means that adding an extra belief can never increase the value of C, whereas it seems clear that it can increase the coherence (see Bovens and Hartmann (2003b) for an example). There are other ways to generalize C, however. For example, by taking into account overlap between subsets of beliefs in the information set it might be possible to avoid this problem. Even if such a solution is possible, however, the Bovens and Hartmann result seems to establish that no extension of a measure such as C will be truth conducive for sets with more than two members.

In the case of information pairs, however, everything is different. Perhaps there are domains where it is of interest to restrict comparisons to information pairs (this indeed is the focus of this paper). Furthermore, it is possible to think of building up a set of beliefs incrementally so that coherence between two beliefs is all that needs to be considered: coherence of the new belief with the conjunction of the current beliefs. It is of particular interest to note that the C measure is truth conducive for information pairs (see Appendix for a proof).

Even if the arguments in this section are not sufficiently convincing to avoid the implications of the impossibility results, they do at least provide a motivation for taking the C measure as a good account of coherence in the case of information pairs. Advantages of this measure will now be pursued further.

3.3 The Bovens-Hartmann Account for Information Pairs

As noted earlier, the Bovens-Hartmann account provides a partial ordering of information sets. In the case of information pairs it turns out that the ordering resulting from this approach is similar to that resulting from use of the C measure. For this reason its merits as an account of coherence are considered here before going on to look at how coherence can be linked with explanation.

Suppose we have information sets $S = \{A,B\}$ and $S'=\{A',B'\}$, with $a_0 = P(A,B)$ and $a_1 = P(A,\neg B) + P(\neg A,B)$ and similarly for a_0' and a_1'. Bovens and Hartmann (2003b, p. 37) provide conditions for the information set S being more coherent than (or equally coherent to) the information set S' (i.e. $S \geq S'$). These conditions are both necessary and sufficient whenever there are two elements (n=2) and can be interpreted as follows.

There are precisely two ways to decrease the coherence in moving from information set S to S':

1. Shrinking the overlapping area ($a_0' \leq a_0$) and expanding the non-overlapping area ($a_1' \geq a_1$);

2. Expanding the overlapping area ($a_0' \geq a_0$) while expanding the non-overlapping area to a greater degree ($a_1' / a_1 \geq a_0' / a_0$).

There appears to be an asymmetry between 1 and 2, which can be seen by contrasting condition 1 with

1'. Shrinking the overlapping area ($a_0' \leq a_0$) and shrinking the non-overlapping area to a lesser degree ($a_1' / a_1 \geq a_0' / a_0$).

Note that 1′ and 2 can be replaced with the simple requirement that ($a_1' / a_1 \geq a_0' / a_0$). This is precisely the requirement for decreasing coherence in moving from S to S' which arises from the C measure discussed earlier. Thus, the orderings obtained by the Bovens-Hartmann account are a proper subset of those obtained by using the C measure. Although the Bovens-Hartmann account does not satisfy the Bovens-Olsson condition, this is because it is more conservative than the C measure in terms of which sets should be ordered.

There is a symmetry between conditions $1'$ and 2 which is lacking between 1 and 2. Since condition $1'$ is very similar to 2, it is not obvious how the latter could be justified and yet the former ruled out. It seems that a more defensible approach would be to replace condition 1 with $1'$ or else to leave condition 1 as it stands but drop condition 2.

A related problem with the asymmetry is that it leads either to
a) cases which intuitively should be ordered, but are not; or
b) cases which intuitively should not be ordered, but are.
Whether one views a) or b) as the main problem will depend on intuitions about coherence orderings. A problem with condition 1 can be illustrated by the following scenarios for the 'Corpse in Tokyo' example (Bovens and Hartmann 2003b, p. 39):

Scenario λ, S^{λ} - witness 1 reports squares 1-20 and witness 2 reports squares 11-20. This gives $a_0^{\lambda} = 0.1$ and $a_1^{\lambda} = 0.1$.

Scenario μ, S^{μ} - witness 1 reports squares 1-10 and witness 2 reports square 10. This gives $a_0^{\mu} = 0.01$ and $a_1^{\mu} = 0.09$.

According to the conditions 1 and 2 (or equivalently condition 2.12 on p.37), S^{λ} and S^{μ} cannot be ordered since $a_1^{\mu} / a_1^{\lambda} = 0.9$ is not greater than or equal to $1 = max(1, a_0^{\mu} / a_0^{\lambda})$ and $a_1^{\lambda} / a_1^{\mu} = 1.11$ is not greater than or equal to $10 = max(1, a_0^{\lambda} / a_0^{\mu})$. However, intuitively it seems that S^{λ} is more coherent than S^{μ} since, in moving from S^{λ} to S^{μ}, the overlapping area has been reduced very substantially (0.1 to 0.01) while the non-overlapping area has only been reduced by a very small amount (0.1 to 0.09). Of course, if condition $1'$ were adopted, this intuition would be taken into account.

Perhaps the problem with $1'$ is that it would give a complete ordering and so would permit information sets to be ordered that intuitively should not be, eg it would result in S^{ε} being more coherent than S^{δ} in the examples on p.40 of Bovens and Hartmann (2003b). However, while condition 1 avoids these problems, condition 2 seems to run into similar problems. For example, consider again

Scenario α - witness 1 reports squares 50-60 and witness 2 reports squares 51-61. This gives $a_0^{\alpha} = 0.1$ and $a_1^{\alpha} = 0.02$.

But now consider a new scenario:
Scenario ϕ - witness 1 reports squares 50-64 and witness 2 reports squares 51-66. This gives $a_0^{\phi} = 0.14$ and $a_1^{\phi} = 0.03$.

According to condition 2 above S^{α} is more coherent than S^{ϕ}, although it does not seem any more intuitively obvious that there should be an ordering here than in the comparison of S^{δ} with S^{ε}.

These examples highlight problems for the Bovens-Hartmann account[9] for information pairs and suggest that the C measure might be more appropriate in this case since it avoids the asymmetry involved in conditions 1 and 2. Of course, as Bovens and Hartmann point out the most obvious way of extending the C measure to larger information sets shows that it has serious problems. However, in this paper we are interested in information pairs and here it seems like a good coherence measure since it is prior independent, satisfies the Bovens-Olsson condition, is truth conducive in this case and avoids the problem of asymmetry noted above. These characteristics not only make the C measure a good measure of coherence for information pairs, but make it suitable for ranking explanations as will be explored in the next section.

4. Coherence and the best explanation

We have looked at ways of ordering explanations in section 2 and ordering information sets in terms of coherence in section 3. In this section we now consider whether a suitable coherence measure can be used to order explanations. After establishing a link between coherence and explanation and demonstrating the suitability of the C measure, we discuss an example to show how it compares with the other approaches for ordering explanations discussed in section 2.

4.1 Making the link

The approach to be adopted is to use a coherence measure to rank explanations so that the best explanation is the one which best coheres with the evidence. However, this leaves the question as to which coherence measure should be used. Recall that the *Explanation Ranking condition* in section 2 was proposed as a condition for a satisfactory account of 'best explanation'. Closer inspection reveals that this is nothing other than a special instance of the *Bovens-Olsson condition* for coherence. The Explanation Ranking condition can be recast in terms of two probability distributions P and P′ for an hypothesis H and evidence E, with the additional requirement that $P(E) = P'(E)$. Thus, a strong link has been made between the notions of coherence and explanation (or at least between a particular conception of coherence and the proposal for ranking explanations). Since the coherence measure defined in expression (3) is the only one of those considered to satisfy the Bovens-Olsson condition, it also seems the most appropriate measure to use in ranking explanations[10].

The suitability of this measure for selecting the best explanation can be seen more clearly by noting that provided the joint probability $P(H_i \wedge E) \neq 0$ expression (3) can be rewritten (replacing A with H_i and B with E) as,

$$C(H_i, E) \quad = \quad \left[\frac{1}{P(H_i \mid E)} + \frac{1}{P(E \mid H_i)} - 1 \right]^{-1}. \qquad (4)$$

The best explanation can now be defined as the hypothesis which is most coherent with the evidence and so has the maximum value for $C(H_i,E)$.[11] It is straightforward to show from expression (4) that this coherence measure satisfies the Bovens-Olsson condition and also that it satisfies the Explanation Ranking condition. Thus this coherence approach subsumes the ML and MPE approaches since in all cases where ML and MPE agree as to which explanation is best, the coherence approach will yield the same answer (unlike the CB approach). Returning to the example discussed in section 2.3 it is found that $C(H_1,E) = 0.3$ and $C(H_2,E) = 0.57$ and so H_2 (torn ligaments) is found to be the best explanation in agreement with ML and MPE. The suggestion here is that in the more difficult cases where ML and MPE disagree the coherence approach will mediate between them and yield the answer that is more intuitively correct.

As noted earlier, the idea in this work is similar in many respects to that of Thagard (1989) who uses the concept of explanatory coherence in determining the best explanation. In Thagard's case, explanatory coherence occurs when there is an explanatory relation between two (or more) beliefs. The most obvious difference is that Thagard adopts a connectionist as opposed to a probabilistic approach, although a comparison with Bayesian networks has been carried out (Thagard 2000). The approach proposed here would need to be extended in a number of directions in order to make a general comparison with Thagard's work possible.

It is worth noting that the use of coherence advocated here in defining the best explanation differs from another way in which coherence may be considered to be relevant. It could be claimed that the coherence of an hypothesis is concerned with its internal consistency rather than with the relationship between the hypothesis and the evidence.[12] According to such a view coherence along with other factors such as simplicity and fruitfulness may play a role in theory choice, but could not be used on its own to determine the best explanation. Note that in the account of coherence used in this paper the main focus is on the relationship between the hypothesis and the evidence, although this is not to say that internal consistency is ignored. This can be seen from the fact that the problem with using maximum likelihood to rank explanations was that it ignored prior probabilities completely, whereas the coherence approach overcomes this problem to some extent. Thus, if internal consistency contributes to the prior probability, its influence can be taken into account.

4.2 A comparison of approaches

In this section the ML, MPE, CB and coherence approaches are compared by applying them to several simple scenarios. Suppose that people with a range of medical conditions are known to be more susceptible to brain tumours. Here we consider two of these medical conditions which are believed to be causally related to brain tumours. Let us call these conditions H_1 and H_2 and try to determine which provides a better explanation of the occurrence of tumours (which we will represent as E). Let us also assume that these conditions are mutually exclusive, but not exhaustive (ie some people who are susceptible to brain tumours do not have either condition H_1 or H_2). Consider now four fictitious scenarios which specify prior probabilities of H_1 and H_2 within the group of susceptible people and the probabilities for developing a brain tumour given H_1 and given H_2. In all cases the prior probability of a tumour for the group, P(E), is 1/10. The results for each scenario are presented in table 1.

Scenario 1:
$P(H_1) = 1/25$, $P(H_2) = 3/50$, $P(E|H_1) = 1/2$ and $P(E|H_2) = 1$.

Scenario 2:
$P(H_1) = 2/25$, $P(H_2) = 1/25$, $P(E|H_1) = 1/4$ and $P(E|H_2) = 1$.

Scenario 3:
$P(H_1) = 9/25$, $P(H_2) = 1/50$, $P(E|H_1) = 1/6$ and $P(E|H_2) = 1$.

Scenario 4:
$P(H_1) = 2/25$, $P(H_2) = 1/50$, $P(E|H_1) = 3/4$ and $P(E|H_2) = 1$.

Table 1. A summary of the various measures obtained using the ML, MPE and coherence approaches for various scenarios discussed in the text. CB does not provide a measure, but where it provides an ordering the result is displayed.

	CB	ML		MPE		Coherence	
	Best	$P(E\|H_1)$	$P(E\|H_2)$	$P(H_1\|E)$	$P(H_2\|E)$	$C(H_1,E)$	$C(H_2,E)$
Scenario 1	H_2	1/2	1	1/5	3/5	1/6	3/5
Scenario 2	-	1/4	1	1/5	2/5	1/8	2/5
Scenario 3	-	1/6	1	3/5	1/5	3/20	1/5
Scenario 4	-	3/4	1	3/5	1/5	½	1/5

The first scenario is very straightforward since H_2 has greater likelihood and higher prior probability and so all the approaches agree that H_2 is the better explanation. Scenario 2 is not so clear cut since H_2 has a much greater likelihood, but a lower prior probability. Consequently the CB approach fails to rank H_1 and H_2 even though the ML and MPE approaches agree that H_2 is better. As expected, the coherence approach also yields H_2 as the better of the two explanations. The difference in likelihood is more significant in this case than the difference in prior probabilities.

In scenario 3 it is much more difficult to get an intuitive picture as to what the correct answer should be. H_2 has a much greater likelihood, but now a much lower prior probability. In this case ML and MPE differ, with ML very strongly in favour of H_2 and MPE strongly favouring H_1. In this case coherence agrees with ML in supporting H_2 but note that it reflects the tension by ranking H_2 only slightly ahead of H_1. Intuitively, this seems correct since the only reason H_1 has a higher posterior probability is because of a very high prior probability. Consequently, there will be many patients who have H_1 but do not develop brain tumours.

The prior probability and likelihood also point in different directions in scenario 4 and, although ML favours H_2 and MPE favours H_1 as in scenario 3, there has been a sufficiently large change so that coherence now agrees with MPE. The reason for this is clear: although there will be some patients who have H_1 but do not develop brain tumours, most people who develop brain tumours have H_1 (since $P(H_1|E) = 3/5$); on the other hand, very few people who develop brain tumours have H_2 (since $P(H_2|E) = 1/5$).

These examples are intended to illustrate advantages that the coherence approach has over other ways of comparing and ranking competing explanations. It is claimed that these are general features of the coherence approach, which arise because of the way in which it is able to avoid the extremes of the ML and MPE approaches and yet takes into account the insights of both.

Another feature of the coherence approach which distinguishes it from the ML and MPE approaches concerns the ranking of the 'catch-all' hypothesis, H_C, which in the examples given would be the hypothesis that neither condition H_1 nor H_2 explains the tumour. It will usually be the case that the specified hypotheses, H_1 and H_2, will have reasonably high likelihoods and will not be extremely improbable, otherwise they would not have been considered in the first place. For similar reasons H_C will typically have a low likelihood, as is the case for all four scenarios considered (eg $P(E|H_C) = 1/45$ in scenario 1), and so according to ML it will not be as good an explanation as H_1 or H_2. By contrast, H_C will often have a high prior probability, as is the case for all four scenarios (eg $P(H_C) = 9/10$ in scenario 1), and in some cases this will result in a high posterior probability. For example, in scenario 2 MPE ranks H_C equal to H_2 as the best explanation, while in all the other scenarios it is ranked equal to the weaker explanation. Thus, according to MPE the catch-all hypothesis, H_C, tends to do quite well in competition with the specified hypotheses. This is despite the fact that it is a very vague hypothesis covering all possible ways in which H_1 and H_2 could be false and so intuitively it does not seem that it should be a serious competitor.

It is no surprise that the coherence approach mediates between ML and MPE in this respect. In all four scenarios it agrees with ML that H_C is by far the weakest explanation. For example, in scenario 2, for which MPE ranks H_C joint best, the coherence of H_C with the evidence is 2/47. Thus, the coherence approach seems to capture our intuitions better than MPE. Furthermore, unlike the ML approach, coherence will rank H_C higher than a specified hypothesis if the latter has a sufficiently small prior probability. Thus, it could also be argued that coherence has an advantage over ML in this respect since in such a scenario we might be tempted to say that the specified hypothesis should not have been proposed as a serious competitor.

5. Inference to the most coherent explanation (IMCE)

So far coherence has been proposed as a measure which can be used to compare and rank competing explanations, but can it also be used within the context of IBE? The claim that 'best explanation' can be understood in terms of coherence is independent of whether it is appropriate to make an inference to the truth (or probable truth) of the most coherent explanation. However, it is interesting to pursue the possibility that coherence might provide a satisfactory way to make IBE more precise. The major objection to such a strategy comes from van Fraassen's attack on IBE (van Fraassen 1980, 1989). As noted in section 1, van Fraassen presents a number of arguments against IBE, but here it is his Dutch book argument that is particularly relevant. The idea is that IBE can be conceived of as a rule for updating probabilities, but a consequence of this is that it necessarily leads to a diachronic Dutch book argument and hence IBE can be shown to be irrational. Van Fraassen describes a Bayesian called Peter who is trying to determine the bias of a die and as a good Bayesian he conditionalizes on evidence as it becomes available. But Peter is also a believer in IBE and so increases the posterior probability of good explanations after conditionalization. However, this departure from Bayesian conditionalization leaves Peter susceptible to a diachronic Dutch book.

It is important to note that van Fraassen's argument does not depend on a particular account of IBE, but is intended to apply to any probabilistic version of IBE. Van Fraassen's conclusion can be resisted, however, by arguing that IBE need not involve any departure from Bayesian probabilities and so the two approaches are in fact compatible. In this vein, Okasha (2000) and Lipton (2001, 2004) have adopted similar approaches, arguing that explanatory considerations play an important role in implementing Bayesian reasoning. Such considerations are brought to bear in the determination of priors and likelihoods as well as in judgements of which pieces of evidence are relevant.

A possible difficulty with this reconciliation concerns what exactly constitutes the best explanation. Within this framework, it is difficult to see how the best explanation could be anything other than the hypothesis with the highest posterior probability. However, it is important to point out that Lipton is certainly not *defining* 'best explanation' as the most probable explanation as pointed out in section 2.2, but in reconciling IBE with Bayesianism he does seem to be committed to the idea that the best explanation will turn out to be the most probable explanation. This would certainly be the ideal conclusion for the defender of IBE, but can such a close association be established? Lipton has certainly made a strong case for the idea that explanatory considerations such as simplicity, scope, etc. play an important role in determining priors and likelihoods. But how can it be established that a proper consideration of all the relevant explanatory factors will always yield the most probable explanation? Is it possible that the best explanation in terms of all the relevant explanatory factors could fail to be the most probable explanation as judged by Bayes' theorem? Clearly not, if Bayesianism and IBE are to be reconciled, although section 2.2 above suggests otherwise.

If the best explanation is selected in terms of its coherence with the evidence, it is clear that reconciliation between Bayesianism and IBE is impossible. Although no rule for updating probabilities in terms of coherence has been proposed in the foregoing sections, it is undoubtedly true that if such a rule were to be devised van Fraassen's argument would be applicable. Thus, if inference to the most coherent explanation (IMCE) is conceived of as a rule for updating probabilities it seems doomed to fail. Furthermore, any conception of IMCE would seem to involve attributing belief to the most coherent explanation and so, since in general this is in conflict with believing the most probable explanation, van Fraassen's argument seems to apply. If this is the correct way to understand IMCE, and if IMCE is the correct way to understand IBE, then the discussion in the earlier parts of this paper actually supports van Fraassen and shows that IBE is untenable. Two options open to the defender of IMCE are discussed below.

5.1 IMCE as descriptive, Bayesianism as normative

In outlining a number of approaches to relating IBE and Bayesianism, Lipton (2004) discusses the possibility that IBE gives a correct *description* of our inductive practices, while Bayesianism gives the correct *normative* account. While this claim is debatable it does raise the question as to whether the reasoner using IBE would tend to make correct inferences. The hope would be that someone who makes inferences to the best explanation will tend to arrive at the same conclusion as they would have done had they adopted a Bayesian approach.

To see how IMCE might be relevant in this respect consider again the definition of coherence in equation (3). If this coherence measure provides an adequate account of 'best explanation' it goes some way to satisfying Lipton's requirement for explanatory loveliness, "After all, if Inference to the Loveliest Explanation is a reasonable account, loveliness and likeliness will tend to go together, and indeed loveliness will be a guide to likeliness" (Lipton 2004, 61). It might seem that the link between coherence and likeliness is not as strong as it needs to be since a likely explanation (ie an explanation with a high posterior probability) will have a low value of coherence if it has a small likelihood. This is not a problem, however, since it is not a requirement of IBE that it account for all inferences. It is a requirement that good explanations will tend to have high posterior probabilities and this requirement is satisfied. To see that this is the case note that

$$\forall x \in [0,1] \quad C(H,E) \geq x \quad \Rightarrow \quad P(H \mid E) \geq x. \tag{5}$$

Thus, it seems reasonable to say that coherence and likeliness 'tend to go together' at least in the sense that is relevant for IBE and, furthermore, that loveliness defined in terms of probabilistic coherence is a guide to likeliness.

Even if this is correct, IMCE seems to face an insurmountable problem: determining the most coherent explanation as specified in equation (3) requires knowledge of the relevant probabilities, but this makes IMCE superfluous as an approximation to Bayesian reasoning. However, if this definition of coherence is not intended to be used as a formula to calculate coherence, but rather a formal explication of a concept that is used informally then perhaps the problem can be resolved. If humans are better at assessing and ranking explanations in terms of how well they cohere with the evidence than they are at making informal judgements about likelihoods and priors (as Bayesianism would require) then the account of 'best explanation' in this paper suggests that they will tend to come up with explanations that are also highly probable.

If this is correct the defender of IMCE cannot claim that the most probable explanation and the best explanation are coextensive. In fact, by definition they are not and so the relationship between the two must be somewhat weaker. It could be claimed that good explanations will tend to have high posterior probabilities and perhaps this is good enough. From this perspective IMCE could be seen as a kind of surrogate for Bayesian inference. This would be relevant if IMCE (or something like it) is, in fact, how humans reason rather than by using Bayesian inference. Consequently the status of IMCE is one for psychologists to address, although there is an important theoretical difference between Bayesianism and IMCE that might be relevant. Bayesianism relies on judgements concerning prior probabilities, while IMCE is independent of the priors provided the conditional probabilities are fixed. Thus judgements about coherence are only concerned with how well the hypothesis and the evidence fit together rather than how probable they are in the first place. If humans are worse at making judgements about prior probabilities than conditional probabilities this might favour IMCE.

5.2 IMCE instead of Bayesianism

So far we have assumed that the goal of inference is to select the hypothesis that is most probable given the evidence, but perhaps this is not quite right. That high probability is not the only factor (and perhaps not the most important) has been pointed out by many authors such as Popper (1959, section 83), Lipton (2001, 2004), Salmon (2001) and Psillos (2004, 2006). As Lipton points out, "… high probability is not the only aim of inference. Scientists also have a preference for theories with great content, even though that is in tension with high probability, since the more one says the more likely it is that what one says is false." (Lipton 2004, 116). To illustrate the point consider the following hypotheses: H_1: the die is biased to an even number; H_2: the die is biased to 2. No matter what evidence is obtained from rolling the die the posterior probability of H_1 will always be higher since H_1 is entailed by H_2, yet we might be more interested in making an inference to the more specific hypothesis H_2. This of course is particularly relevant to science since we are interested in theories which are as precise as possible and testable rather than being vague and compatible with both a piece of evidence and its negation.

In many ways the ideal type of hypothesis adopted by scientists is one which entails the evidence ($P(E|H_1) = 1$) in question and is also highly confirmed by the evidence ($P(H_1|E)$ is high). This kind of scenario often occurs when the evidence cannot be accounted for on the basis of other theories. Note that these are exactly the requirements for a high degree of coherence and so perhaps inferences *should* be made to the most coherent explanation *rather* than the most probable explanation in some cases.

Despite these considerations, however, it is still clear that in many cases we do wish to find the most probable explanation. For example, if we wish to make an inference to the explanation for why a particular patient suffered a brain tumour it will presumably be the most probable explanation that is required. Thus, the response outlined in this section would seem to be inadequate on its own.

6. Conclusions

The first goal in this paper has been to provide an adequate account of 'best explanation'. By using a particular definition of probabilistic coherence I have proposed that 'best explanation' should be understood as 'most coherent explanation' in the sense that the best explanation should have a higher degree of coherence with the explanandum than the competing explanations. Advantages of this approach over obvious rivals have been discussed. The main advantage is that it takes into account the insights of its rivals, while avoiding their drawbacks.

A second goal in the paper has been to determine whether the account of 'best explanation' supports or undermines IBE. If the account is used as a rule for updating probabilities it clearly undermines IBE since there will be cases where the best explanation is not the most probable one. Nevertheless, it does suggest that good explanations will be probable explanations and so someone who reasons using IBE will tend to make probable inferences. Furthermore, it was pointed out that the kind of inference often required in science is not inference to the most probable hypothesis, but inference to an hypothesis with various features which might be better characterized in terms of probabilistic coherence.

Acknowledgements
I would like to thank Alexander Bird, Donald Gillies, Erik Olsson and Stathis Psillos for helpful comments on earlier drafts of this paper.

Appendix

Following the notation of Bovens and Hartmann (2003b), consider an information set of the form $\{R_1, R_2\}$, where R_i are referred to as fact variables, and corresponding report variables $REPR_1$, $REPR_2$, where $REPR_i$ is a report that R_i is true. For all i, let $p = P(REPR_i \mid R_i)$ and $q = P(REPR_i \mid \neg R_i)$ and define $r=q/p$. If we assume that $p > q > 0$ and that each report variable is conditionally independent of other report variables and fact variables given the fact variable on which it reports, Bovens and Hartmann have shown that,

$$P(R_1, R_2 \mid REPR_1, REPR_2) = \frac{a_0}{a_0 + a_1 r + a_2 r^2},$$

where $a_0 = P(R_1, R_2)$ and $a_1 = P(R_1, \neg R_2) + P(\neg R_1, R_2)$ and $a_2 = 1 - a_0 - a_1 = P(\neg R_1, \neg R_2)$.

Definition. Consider a pair of information sets $S=\{R_1, R_2\}$ and $S'=\{R_1', R_2'\}$. A coherence measure *Coh* is truth conducive *ceteris paribus* if and only if: if $Coh(R_1, R_2) > Coh(R_1', R_2')$, then $P(R_1, R_2 \mid REPR_1, REPR_2) > P(R_1', R_2' \mid REPR_1', REPR_2')$. In the ceteris paribus clause we keep the prior probability fixed, i.e. $P(R_1, R_2) = P(R_1', R_2')$ and r fixed in the interval $(0,1)$.

Proposition. The C measure defined in equation (3) in section 3.1 is truth conducive for information pairs.

Proof

Let us assume that $C(R_1, R_2) > C(R_1', R_2')$. Note that this can be written as,

$$\frac{a_0}{a_0 + a_1} > \frac{a_0}{a_0 + a_1'},$$

since we assume that $a_0 = a_0'$. It follows that $a_1 < a_1'$.

Since $r \in (0,1)$, it follows that,

$$(a_1' - a_1)r - (a_1' - a_1)r^2 > 0$$

and hence,

$$a_0 + a_1 r + (1 - a_0 - a_1)r^2 < a_0 + a_1' r + (1 - a_0 - a_1')r^2.$$

From this it immediately follows that,

$$P(R_1, R_2 \mid REPR_1, REPR_2) > P(R_1', R_2' \mid REPR_1', REPR_2').$$

Thus, C is truth conducive for information pairs.

Notes

[1] In fact, Peirce identifies abduction with the process of generating hypotheses rather than comparing them, but the terms IBE and abduction are often used interchangeably.
[2] In the rest of the paper the term 'explanation' should be understood as 'potential explanation'.
[3] This is not to say that base-rate information should always be taken into account. Psillos (2004) also distinguishes between correct probabilistic reasoning on the one hand and getting at the truth on the other. To get closer to the truth it may be appropriate in certain cases to ignore base-rates because of particular features of the case at hand.
[4] I am not denying that such an hypothesis could count as an explanation. In fact, Salmon (1989) argues if there are explanations of particular facts that are genuinely statistical in character then it follows that there will be explanations which lower the probability of the explanandum. Nevertheless, it still seems reasonable to say that the probability of the explanandum given the hypothesis should be taken into account when considering the goodness of an explanation.

[5] To note that this is the case consider Okasha's remark that "I have … offered no criteria for ranking explanations in terms of their goodness" which would be false if 'best' were equated with 'most probable'.
[6] This paper provides a good overview of explanation in Bayesian networks and addresses a number of issues closely related to the current work.
[7] This is not to say that they are therefore inadequate measures of coherence. Perhaps they capture the intuition of coherence as striking agreement, for example. Furthermore, an alternative response to the roulette example is also given by Bovens and Olsson (2000) where they suggest that their minimal condition could be weakened by requiring that $P(A,B) = P'(A,B)$. This weakened version of the condition would be satisfied by the Bovens and Hartmann (2000) account of coherence.
[8] For Shogenji's and Fitelson's accounts counterexamples to the Bovens-Olsson condition can be found where the probabilistic relations are as specified in the condition, but the information set $\{A,B\}$ is more coherent on distribution P' than on P. This is not the case for the Bovens-Hartmann account, although there are counterexamples in this case where the probabilistic relations are as specified in the condition, but their account fails to order the set $\{A,B\}$ as more coherent on one distribution than the other.
[9] Similar objections have been raised by Meijs and Douven (2005).
[10] Since the Explanation Ranking condition is a special case of the Bovens-Olsson condition, it is not an absolute requirement that a coherence measure satisfy the latter in order to satisfy the former. For example, the coherence measure of Shogenji, which in the case of an hypothesis H and evidence E can be written as $P(E|H)/P(E)$, falls into this category (since $P(E) = P'(E)$). However, it only satisfies the Explanation Ranking condition in the trivial sense that it is in effect identical to the likelihood and so will always agree with the ML approach in its selection of the best explanation.
[11] The use of coherence here does not necessarily require a commitment to coherentism. A foundationalist will need some way to justify non-foundational beliefs on the basis of foundational beliefs and perhaps the coherence measure could be used in this way.
[12] The view of coherence as internal consistency seems to be what Bovens and Hartmann (2003a) have in mind. They explore the role that their account of coherence would play in scientific theory choice and find that a more coherent theory could result in a higher degree of confidence in certain circumstances.

References

Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press

Bovens, L. and Olsson, E. J. (2000). Coherentism, reliability and Bayesian networks. *Mind*, *109*, 685-719

Bovens, L and Hartmann, S. (2003a). Solving the riddle of coherence. *Mind, 112*, 601-633

Bovens, L and Hartmann, S. (2003b). *Bayesian Epistemology.* Oxford: Oxford University Press

Chajewska, U. and Halpern, J. Y. (1997). Defining explanation in probabilistic systems. *Proceedings of the 13th Conference on Uncertainty in AI*, 62-71

Fitelson, B. (2003). A probabilistic theory of coherence. *Analysis*, *63*, 194-199

Glass, D. H. (2002). Coherence, explanation and Bayesian networks. In M. O'Neill, R.F.E Sutcliffe et al (Eds.), *Artificial Intelligence and Cognitive Science*. Berlin Heidelberg New York: Springer-Verlag, *Lecture Notes in Artificial Intelligence*, *2464*, 177-182

Glass, D. H. (2005). Problems with priors in probabilistic measures of coherence. *Erkenntnis*, *63*, 375-385

Harman, G. (1965) The inference to the best explanation. *The Philosophical Review*, *74*, 88-95

Hempel, C. G. (1965). *Aspects of Scientific Explanation*. New York: Free Press

Ladyman, J., Douven, I., Horsten, L. and van Fraassen, B. (1997). A defence of van Fraassen's critique of abductive reasoning: reply to Psillos. *The Philosophical Quarterly*, *47*, 305-321

Lewis, C. I. (1946). *An Analysis of Knowledge and Valuation*. LaSalle: Open Court

Lipton, P. (2001). Is explanation a guide to inference? A reply to Wesley C. Salmon. In G. Hon and S. S. Rakover (Eds.), *Explanation: Theoretical Approaches and Applications* (pp. 93-120). Dordrecht: Kluwer

Lipton, P. (2004). *Inference to the Best Explanation*. London: Routledge. 2nd Edition.

Meijs, W. and Douven, I. (2005). Bovens and Hartmann on Coherence. *Mind*, *114*, 355-363

Okasha, S. (2000). Van Fraassen's critique of inference to the best explanation. *Studies in the History and Philosophy of Modern Science*, *31(4)*, 691-710

Olsson, E .J. (2002). What is the problem of coherence and truth? *The Journal of Philosophy*, *99 (5)*, 246-272

Olsson, E. J. (2005). *Against Coherence*. Oxford: Oxford: Oxford University Press

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems*. San Mateo: Morgan Kaufman

Popper, K. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

Psillos, S. (1996). On van Fraassen's critique of abductive reasoning. *The Philosophical Quarterly*, *46*, 31-47

Psillos, S. (1997). How not to defend constructive empiricism: a rejoinder. *The Philosophical Quarterly*, *47*, 369-372

Psillos, S. (1999). *Scientific Realism: how science tracks truth*. London: Routledge
Psillos, S. (2002). Simply the best: a case for abduction. In F. Sadri and A. Kakas (Eds.), *Computational Logic: From Logic Programming into the Future*. Berlin Heidelberg New York: Springer-Verlag. *Lecture Notes in Artificial Intelligence*, *2408*, 605-625
Psillos, S. (2004). Inference to the best explanation and Bayesianism. In F. Stadler (ed.), *Induction and deduction in the sciences*, (pp. 83-91). Dorderecht: Kluwer.
Psillos, S. (2006). Thinking about the ultimate argument for realism. In C. Cheyne and J. Worrall (Eds.), *Rationality and Reality: Conversations with Alan Musgrave*. Dorderecht: Kluwer.
Salmon, W. C. (1989). *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press
Salmon, W. C. (1998). *Causality and Explanation*. Oxford: Oxford University Press
Salmon, W. C. (2001). Reflections of a bashful Bayesian: a reply to Peter Lipton. In G. Hon and S. S. Rakover (Eds.), *Explanation: Theoretical Approaches and Applications* (pp. 121-136). Dordrecht: Kluwer
Shimony, S. (1994). Finding maps for belief networks is NP-hard. *Artificial Intelligence*, *68*, 399-410
Shogenji, T. (1999). Is coherence truth-conducive? *Analysis*, *59*, 338-345
Thagard, P. (1989). Explanatory Coherence. *Behavioral and Brain Sciences*, *12*, 435-467
Thagard, P. (2000). Probabilistic networks and explanatory coherence. *Cognitive Science Quarterly*, *1*, 93-116
Van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Clarendon Press
Van Fraassen, B. C. (1989) *Laws and Symmetry*. Oxford: Clarendon Press