Ampliative Reasoning: Induction or Abduction?

Stathis Psillos*

Induction and abduction are modes of ampliative reasoning. Ampliative, or content-increasing, is the type of reasoning in which the conclusion of the inference exceeds what is already stated in the premises. As such, the truth of the conclusion is not guaranteed whenever the premises are true. It is perfectly consistent with the truth of the premises that the conclusion be false. So far, philosophers have mostly concentrated on the problem of justification (Hume's problem): given that ampliative reasoning is not necessarily truth-preserving, how can we justify it? That is, how can we show that it is rational to employ ampliative reasoning?

Although in the philosophical literature this is known as the problem of induction, Hume's challenge concerns ampliative reasoning in general. His point was that there is no non-circular way to show that ampliative reasoning is truth-preserving. Since the premises of an ampliative argument do not logically entail the conclusion, there are possible worlds in which the premises are true and the conclusion false. How, then, can we show that the actual world is one of the possible worlds in which when the premises of an ampliative argument are true, its conclusion is also true? Hume pointed out that any attempt to establish the former is bound to be the conclusion of an ampliative argument, and hence it's bound to beg the question. Take for instance the predictive inference from the premise that 'all observed Fs have been G' to the conclusion that 'the next F is going to be G'. Any attempt to show that this inference is truth-preserving would rest on the premise that so far whenever all observed Fs turned out to be G, so did the next F we observed. But, short of an ampliative argument, we have no way to infer that next time we'll observe a number of Fs being G, the next F is also going to be G.

However, there is another problem that should be dealt with, which is independent of the issue of justification. The descriptive problem: what is the structure of ampliative reasoning? In this piece, I am going to focus on the following aspect of the descriptive problem: what is the relation, if any, between induction and abduction? In particular, when we analyse and describe ampliative reasoning, do we need to posit two (or more) different modes of such reasoning and lay down different rules of warrant? Or, can we manage our ampliative inferential economy by positing one all-encompassing mode? Notice that we may use the term 'induction' very broadly to cover any attempt to go beyond a finite set of data to a hypothesis that is well-supported—but not entailed—by them. Then, 'induction' would be a short for ampliative reasoning in general. This is a purely terminological matter, I think. The real issue is what mode of inference best represents ampliative (and, in the above broad sense, inductive) reasoning. When we deal with this issue we'd better examine specific models of inductive reasoning.

The rule associated with enumerative induction is that if one has observed n As being B and no As being not-B, and if the evidence is enough, then one should infer that with high probability 'All As are B'. The general version of the rule is this: if one has found that r out of n observed As being B ($r \leq n$), then one should infer that with high probability r/n of all As are B'. The essence of enumerative induction is that ampliative reasoning proceeds via generalisation. We observe a pattern among the data (or, among the instances of two attributes) and then generalise it so that it covers all relevant data (or all instances of the two attributes). But can this offer a general model of ampliative reasoning?

I think there are four desiderata that a general model of ampliative reasoning should satisfy. a) It must be non-monotonic, that is, it must allow that a certain conclusion be defeated by the inclusion of extra information in the premises. b) It must deal with the 'cut-off point' problem, i.e., it must show in virtue of what generalisations from samples to populations as a whole are warranted. c) It must allow for 'vertical' extrapolation, that is

* Dept of Philosophy, Logic and Scientific Method, London School of Economics, Houghton Street, London WC2A 2AE, UK. e-mail: psillos@lse.ac.uk

it must support conclusions that involve reference to types of entity (or, more generally, vocabulary) that are not already referred to in the premises, e.g., positing scientific unobservables. d) It must accommodate the eliminative dimension of ampliative reasoning, that is, the fact that in typical cases of ampliative reasoning more than one hypotheses consistent with the premises are considered and attempt is being made to find grounds to eliminate (ideally) all but one of them.

Enumerative induction can accommodate only the first of the above: it is clearly non-monotonic. For instance, if added in the premises, the information that the sample was biased undercuts the previous conclusion. But it fails on the rest. It cannot solve the 'cut-off point' problem. It does not tell us when the evidence is enough to warrant the generalisation. Without extra information about the quality of the data, we cannot tell when we have seen enough instances to warrant the generalisation. It cannot accommodate 'vertical' extrapolation. In a certain sense, enumerative induction, pretty much like deduction, operates with the principle 'garbage in, garbage out': the descriptive vocabulary of the conclusion cannot be different from that of the premises. Hence with enumerative induction, although we may arguably gain knowledge of hitherto unobserved correlations between instances of the attributes involved, we cannot gain 'novel' knowledge, i.e., knowledge of entities and causes that operate behind the phenomena. Finally, enumerative induction cannot accommodate the eliminative dimension of ampliative reasoning: unless we consider alternative generalisations of the observed correlations, we cannot attempt to eliminate them. But enumerative induction does not have the resources to generate alternative generalisations.

So enumerative induction cannot adequately model ampliative reasoning in general. But this is hardly news, one may think. Enumerative induction is too crude and too rough a first approximation to ampliative reasoning. We definitely need a better model. But still one may argue that enumerative induction can keep its place in our economy of ampliative reasoning as a distinct mode. My own view is that it need not be kept as such. I will suggest that abduction, suitably understood as Inference to the Best Explanation, can adequately cover ampliative reasoning.

What is abduction? In its original Peircean conception, abduction (also called retroduction) was mostly seen as a heuristic to generate hypotheses which are worth critical attention. These would then be put to the test by other inductive means. The abduced hypotheses were such that could explain the evidence that probed their generation. Seen in this light, abduction has no probative force. It's not a mode of reasoning, strictly speaking, but rather a way to generate the material (hypotheses) that we reason about. Induction also, seen in this light, is no longer crudely enumerative. But it's its close relative: instance confirmation. Hypotheses are put to the test by examining their positive instances. The more positive instances we observe, coupled with the absence of negative instances, the more the hypothesis gets confirmed. Instance-confirmation can be generalised (and liberalised) in the form of hypothetico-deductivism. According to this, hypotheses are tested by examining the observational consequences which they entail together with suitable initial conditions and auxiliary assumptions. In the more general form, the observational predictions need not be instances of the hypotheses—since the latter may involve reference to unobservable entities. But since they are entailed by these hypotheses, if they turn out to be right, the hypotheses get inductive credit.

One may naturally think that Peircean abduction (as a mode of generating hypotheses) together with hypothetico-deductivism (as a mode of testing them) overcome the limitations of enumerative induction and can capture the most general mode of ampliative reasoning. Surely, such a combination avoids some of the problems of the naive enumerative model. For instance, it bypasses the 'vertical extrapolation' problem, since abduction can help us generate hypotheses whose vocabulary is distinct from that of the evidence-statements and whose content can be about unobservable entities and causes. Besides, since hypotheses are not mere generalisations of instances, we no longer face the 'cut-off point' problem. That is, we no longer worry about when the evidence (number of instances) is enough to license generalisation. But, abduction together with hypothetico-deductivism fail the other two desiderata. Hypothetico-deductivism is not non-monotonic. Given that the relation between a hypothesis H and the evidence is

deductive, the addition of any extra information H' to the hypothesis H will entail all that H already entailed. Therefore, it cannot undermine the initial hypothesis H. One may object here that H may be defeated by finding another hypothesis H'' which is more confirmed by the same evidence. This would be a natural 'way out', if it weren't for the fact that hypothetico-deductivism also fails to accommodate the 'eliminative' dimension of ampliative reasoning. For, if all that matters for inductive confirmation is the entailment of the evidence, then it's not clear at all that there are grounds to eliminate alternative hypotheses that entail the same evidence. In fact, whenever an observational result of an H-D test confirms a hypothesis, it also confirms a multitude of other hypotheses that are incompatible with the first but entail the same observational consequence(s). Elimination can only be effected by disconfirmation.

But H-D faces another family of problems. These, too, relate to the deductive part of the mode. Here again, the problems have been extensively discussed in the literature concerning the deductive-nomological (D-N) model of explanation.[1] In essence, the model suggests that to explain a singular event is to construct a deductive argument whose conclusion is the explanandum and whose premises essentially contain a law-like statement (expressing a law of nature). In other words, explanation amounts to nomic-expectability: an event is explained when it is predicted, that is when it is shown that one should have expected this event to happen, had one taken account of the laws that govern the specific situation together with the specific initial conditions.

More generally, it turns out that the D-N model is both too weak and too strong. It's too strong in that there are bona fide explanations which are not deducible from law-like statements. For instance, a doctor can explain the development of paresis by suggesting that the patient suffered from latent untreated syphilis, although the probability that latent untreated syphilis will lead to paresis is only 25%. But also, the D-N model is too weak in that there are perfectly legitimate deductions from law-like hypotheses that do not count as bona fide explanations. For instance, one can deduce, but clearly not explain, the occurrence of a regularity R_1 by from the conjunction of R_1 together with an unrelated regularity R_2. In any case, deduction is neither necessary nor sufficient for explanation, nor for confirmation. (This , however, is not meant to suggest that deduction plays no role in explanation.) Most of the literature has focused on the existence of statistical explanations that are clearly non-deductive. But a central related problem that Hempel [2] brought home is that, strictly speaking, all hypotheses and generalisations that feature in the hypothetico-deductive method as well as the deductive-nomological model of explanation are ceteris paribus. They don't simply state that 'All Fs are G', but rather that 'Other things being equal, all Fs are G'. But when are other things equal? And what happens, if they are not? Ceteris paribus clauses call for the absence of certain disturbing factors. Strictly speaking, they cannot be replaced by a conjunction of all these possible factors, for on most occasions we don't even know what they are, nor can they be described in the language of the hypothesis under test. The point then is that the seemingly innocuous deductive part of the H-D method cannot operate as simply as we think, unless we add as an extra premise that cetera are paria (but how can we affirm that?), or unless we resolve that the deduction concerns systems in which, by definition, cetera are paria (but then how do we connect these systems with the real ones, in which cetera aren't paria?). I don't want to suggest that these problems cannot be resolved. (On the contrary, I think a lot can be gained here if we apply rules of default reasoning.) Rather, I want to point out something that will also be useful later on in the discussion of inference to the best explanation: deduction plays a much less straightforward role in the derivation of consequences (predictions) by which we test hypotheses.

In sum, ampliative reasoning cannot be adequately described as Peircean abduction plus hypothetico-deductivism. But there is no reason to despair. For more recently (cf. Harman, [2]; Lipton, [5]) Peircean abduction has been placed in a broader framework: inference to the best explanation (IBE). In a sense, IBE is a liberalised version of Peircean abduction together with hypothetico-deductivism. IBE can be defended as the mode of ampliative reasoning that can encompass hypothetico-deductivism and enumerative induction as special cases, while satisfying all four features of ampliative

[1] This literature is voluminous. But for an excellent recent account, see Salmon [6]

reasoning: non-monotonicity; the 'cut-off' point; 'vertical' extrapolation and the eliminative dimension. Although, we are still behind in adequately explaining 'explanation', we can nonetheless outline the broad structure of abductive reasoning (in the foregoing broad sense). IBE can be described as explanation-based reasoning: one infers to the best of a number of mutually inconsistent explanations of the evidence (known as 'abducibles'). What makes the inference warranted is that the conclusion is arrived at by virtue of being the best explanation of the evidence. The premises of the inference include more than the evidence (the explanandum). They include relevant background knowledge as well as what computer scientists have called 'integrity constraints'. The inference is licensed only in light of background knowledge and the integrity constraints.

As Lipton [5] has stressed, the best explanation of the evidence is that which offers the most understanding of the explanandum. But how are we to understand 'understanding'? Understanding is an elusive notion. However, one can argue that understanding is connected with expanding /modifying the existing background knowledge so that it encapsulates the explanandum by offering more complete information about it. Yet, not every expansion/modification consistent with the background knowledge improves understanding equally well. Such expansions/modifications should be regulated by constraints, some of which may be fairly general and intuitive marks of good understanding, some other being domain-specific. So, for instance, it adds little to our understanding of an explanandum *e*, if we just consider *e* itself as one of the abducibles and annex it in the background. I will say more on these constraints later, but the basic idea is that telling an informative narrative is not enough for understanding. Rather this narrative should satisfy extra constraints such as avoidance of ad hoc accounts, unification etc.

Abductive reasoning should then be embedded in abductive problem-situations. Building on Kowalski et al. [4], an abductive problem-situation can be characterised as an ordered quadruple $<e, B, A, I>$, where *e* is the explanandum, *B* the background knowledge (or theory), *A* a set of abducibles, and *I* a set of integrity constrains that, generally, consists of two sub-sets: i) a set of more general constraints on explanation and ii) some more domain-specific requirements. In such a situation, we infer to $A_i \in A$ that best explains *e* in the light of *B* and *I*. More specifically, one would expect that *B* narrows down the set of abducibles by licensing some of them while excluding others as uninformative, and *I* selects among the remaining abducibles that which satisfies the extra constraints.

As I noted above, IBE fulfils all four desiderata of ampliative reasoning. It is non-monotonic. The 'best explanation' can be defeated when an even better explanation is found, even though the evidence does not change. The nice feature here is that occasionally, a better explanation can be found only if we substantially revise elements of the background. Hence abductive reasoning can nicely cater for conceptual innovation.

Abduction can solve the 'cut-off point' problem. The mere enumeration of joint instances of two attributes cannot tell when we are warranted to generalise. What licenses generalisation is explanatory considerations, vis., that in a certain abductive problem-situation, a certain generalisation would best explain the evidence. So for instance, we move from an observed correlation between two attributes in a certain sample to an extrapolation to the population as a whole only when such an extrapolation best explains why the frequency in the sample had the observed value. In so doing, we try to exclude spurious correlations, bias in the data selection etc. If there is a better explanation as to why the sample exhibited a certain pattern, e.g., a certain bias on behalf of the selector, then generalisation is not legitimate. So, the cut-off point is when the generalisation emerges as the best explanation of the evidence.

Let's, however, be careful here. I think it's wrong to say that, for instance, the generalisation 'All ravens are black' is licensed because it best explains why all observed ravens have been black. The generalisation 'All ravens are black' offers no such explanation. It can be paraphrased as the conjunction: "'All observed ravens are black' and 'All unobserved ravens are black'" and, hence, it can 'explain' why 'All observed ravens are black' only by affirming that unobserved ravens are black. But in what sense is the claim that 'All unobserved ravens are black' an explanation of the evidence that 'All observed ravens are black'? It clearly offers nothing by way of understanding why

'All observed ravens are black'. It seems that the best explanation here would be one that accounts for the fact that the property of being a raven and the property of blackness are invariably associated, thereby explaining why all ravens (both observed and hitherto unobserved) are black. My claim is that generalisations themselves are not explanatory of the instantial evidence that supports them. In particular, as is well known, the claim that 'All Fs are G' does not explain why each particular F is G.

However, if we view amplification-by-generalisation as a special case of inference to the best explanation the situation becomes quite different. A generalisation is licensed because (and when) it best explains why the pattern of association among the data is what it is. Generally, the licensed inference would be along the following lines: the pattern of association between instances of the two attributes F and G is best explained by the claim that the same pattern holds in the population as a whole (e.g. 'All Fs are G') in the light of the fact that the sampling methods have been random, that instances were drawn from a variety of environments etc. So, in a certain sense, although the generalisation 'All Fs are G' cannot emerge as the best explanation of the evidence that 'All observed Fs are G', it can certainly emerge as the best explanation of the evidence that 'All observed Fs have been found to be G and not not-G'. A similar point has been recently made by Josephson in [3], but I think it deserves more attention.

Abduction can clearly cater for 'vertical' extrapolation. Frequently, the best explanation employs hypotheses that involve reference to unobservables and need vocabulary that goes beyond the conceptual resources of the evidence. Most importantly though, abduction can accommodate the eliminative dimension of ampliative reasoning. When it comes to IBE, the explanation that is being inferred to is to be chosen from a pool of competing potential explanations of the evidence according to whether it offers the most understanding.

So, IBE can present itself as the general mode of ampliative reasoning. But still two central descriptive issues need to be dealt with: a) the explication of 'explanation'; and b) what makes a potential explanation 'best' among its competitors. We know for instance that it won't do to plug into the model the D-N account of explanation. Apart from all the problems that the D-N model faces, IBE would then end up being too close to hypothetico-deductivism for comfort. On the other hand, our understanding of causation is still not good enough to spell out the slogan that explanation amounts to assigning of causal responsibility. There are, of course, clear-cut cases in which we explain why an event occurred (e.g., why the grass is wet) by looking into its causal history and specifying the causal factor that brought it about (e.g., that it rained in the night). But this cannot be the whole story. For one, causal responsibility seems to be bound up with the presence of causal laws and we still have no adequate understanding of them. For another, there are clear-cut cases of explanations which are not causal but rather point to a deductive (or better, deductive by default) connection between the explanandum and the explanans (e.g., when we explain the law of the pendulum by showing how it follows from Newton's laws). Perhaps, causal explanations and nomological explanations are, as Salmon argues, just the two sides of the same coin. At any rate though, there is still more work to be done here.

Similarly, an account of what it is for a potential explanation to be the 'best' should be given. As I noted above, explanation relates to understanding and understanding is related to background knowledge (even when understanding can only be achieved by modifying the background). So, a natural suggestion has been that the best potential explanation is the one that offers the best understanding of the explanandum in the light of background knowledge. But doesn't this trivialise the idea of IBE? For then isn't the best potential explanation just the likeliest explanation given the background? Hence, how can explanatory considerations (as conceptually distinct from considerations of probability) guide ampliative reasoning? The special feature of abductive reasoning is that inference is guided by explanatory considerations. Yet, if in the end of the day we infer to whatever potential explanation is most likely in the light of the background, this special feature seems to be lost. This problem has been nicely put by Lipton [5] when he distinguishes between the loveliest explanation (the one that, if true, would provide the most understanding) and the likeliest explanation (the one that is most probable). Lipton has observed that an adequate model of abductive reasoning should show how explanatory considerations, i.e., considerations of

loveliness, although conceptually distinct from considerations of likeliness, can nonetheless confer likeliness to the conclusion of an ampliative argument.

Here again, we have a challenge that an adequate model of abductive reasoning should meet. But by way of sketching an answer we should note the following. If abductive reasoning is the general mode of ampliative reasoning, then the background itself has been the product of such reasoning and hence its acceptance has also been based on explanatory considerations. So we start off, as it were, from a situation in which explanatory considerations and considerations of likeliness mesh rather than diverge. But, normally, the background only delimits the range of potential explanations of a certain explanandum. It rarely dictates just one potential explanation. More frequently, it licenses more than one potential explanation of the explanandum. There is the point where the considerations encapsulated in the integrity constraints are helpful. They can guide the selection of the abducible that enhances understanding the most. Such considerations include: **-** completeness, i.e., whether the potential explanation accounts for all available evidence, or only parts of it; **-** parsimony: whether or not the potential explanation employs unnecessary assumptions that cannot be independently defended; **-** lack of ad hocness: whether or not the potential explanation is such that the explanandum has been essentially used in its construction; **-** predictive power: whether or not the potential explanation yields further predictions; **-** unifying power: whether or not the potential explanation unifies independently accepted parts of the background as well as explaining the explanandum. These considerations form a by no means complete or uncontroversial list of constraints. Moreover, even if this list were complete and uncontroversial—which it is not, given that we should make space for domain-specific principles, too— we would require a metric to combine and suitably weigh all these features in order to apply them in each case. But at least we seem to have a rough guide of what contributes to 'bestness' which is quite close to features we actually look for when we exercise our intuitive judgement to separate good from bad potential explanations.

To sum up, then, I have argued that abduction can be a good candidate for the manager of our ampliative inferential economy. Induction can just be seen as a special case, and even so the role of explanatory considerations in it should be brought to the fore.[2]

REFERENCES

[1] G. Harman, 'Inference to the Best Explanation', Philosophical Review, **74**: 88-95 (1965).

[2] C.G. Hempel, 'Provisos: A Problem Concerning the Inferential Function of Scientific Theories', in A. Grünbaum & W. Salmon (eds.) The Limitations of Deductivism, University of California Press, (1988).

[3] J.R. Josephson, 'Conceptual Analysis of Abduction' in J.R. Josephson & S.G. Josephson (eds.) Adductive Inference, Cambridge University Press (1994).

[4] A.C. Kakas, R.A. Kowalski, & F. Toni, 'Abductive Logic Programming', J. Logic and Computation, **2**: 719-770 (1992).

[5] P. Lipton, Inference to the Best Explanation, Routledge (1991).

[6] W. Salmon, Four Decades of Scientific Explanation, University of Minnesota Press (1989).

[2] What preceded owes a great deal to the work of Gilbert Harman [1] and Peter Lipton [5]. I want to thank especially Peter for constant inspiration on matters abductive and for writing a fine book that has guided most of the thoughts that appear in this piece. Many thanks are also due to Bob Kowalski. Lots of insightful conversations with him had made me clarify and improve my ideas about abduction.