May 18th, 9:00 AM - May 21st, 5:00 PM

Conclusions as Hedged Hypotheses

John R. Welch
Saint Louis University - Madrid Campus

Conclusions as Hedged Hypotheses

JOHN R. WELCH
Department of Philosophy
Saint Louis University – Madrid Campus
28003 Madrid
Spain
jwelch7@slu.edu

Abstract: How can the objectivity of an argument's conclusion be determined? To propose an answer, this paper builds on Betz's view of premises as hedged hypotheses. If an argument's premises are hedged, its conclusion must be hedged as well. But how can this be done reasonably? The paper first introduces a two-dimensional critical grid. This grid is then applied to three kinds of cases: those characterized by point-valued probabilities, by interval probabilities, and by non-numeric plausibilities.

Keywords: Bayes' theorem, cogency, credibility, inductive strength, Jeffrey's rule, plausibility, probability, soundness, truth, validity

1. Introduction

We attend to an argument's premises ultimately for the sake of its conclusion. We want to know whether to trust the conclusion and, if so, to what extent. How much trust we confer on a conclusion can be represented by assigning it a degree of credence. But how can we do this in a rational way? This paper addresses this question.

Our point of departure is Betz's view of premises as hedged hypotheses (Betz, pp. 212–214). Each premise is treated as a hypothesis, hedged by an estimate of its degree of credence. Think of each premise being paired with a tag that reflects its credence. The premises can then be conjoined, each still bearing its credence tag.

If an argument's premises are hedged, its conclusion must be hedged as well. We cannot do this reasonably without noting that degrees of credence can be represented with varying degrees of precision: by point-valued probabilities, interval probabilities, or non-numeric plausibilities such as 'very likely' or 'unconfirmed.' Conclusions drawn from premises with probabilistic hedges, for example, must be probabilistically hedged as well. Analogous points hold for conclusions drawn from premises hedged with probabilistic intervals or with non-numeric plausibilities.

The foregoing points will be developed in these pages as follows. A two-dimensional grid for evaluating arguments is introduced in Section 2. This grid is then applied to the probabilistic, probabilistic interval, and non-numeric plausibility cases in Sections 3, 4, and 5, respectively. Finally, this grid is linked to the argumentative standards of deductive soundness, inductive cogency, and plausibilistic credibility in the concluding Section 6.

2. A critical grid

This paper builds on the distinction between inductive and epistemic probability. Inductive probability is a property of an argument: the probability that a conclusion is true provided its premises are true. By contrast, epistemic probability is a property of a proposition: the probability that a proposition expressed by a premise, perhaps, or a conclusion is true.

Bondy, P., & Benacquista, L. (Eds.). *Argumentation, Objectivity, and Bias: Proceedings of the 11th International Conference of the Ontario Society for the Study of Argumentation (OSSA), 18-21 May 2016*. Windsor, ON: OSSA, pp. 1-11.

On the basis of this distinction, the paper proposes a two-dimensional grid for evaluating arguments. The grid's vertical dimension is inductive, reflecting the argument's downward flow from premises to conclusion. It specifies the inductive probability (or plausibility) of the conclusion given the premises. The grid's horizontal dimension is epistemic, focusing on the premises without dropping down to the conclusion. It evaluates the epistemic probability (or plausibility) of the premises when conjoined. Joint consideration of these dimensions yields a reasoned view of the degree of credence appropriate to the argument's conclusion. For example, argument A_1 in Figure 1 scores higher on the vertical (inductive) dimension than argument A_2 but lower on the horizontal (epistemic) dimension, while argument A_3 scores higher on both dimensions than either A_1 or A_2.

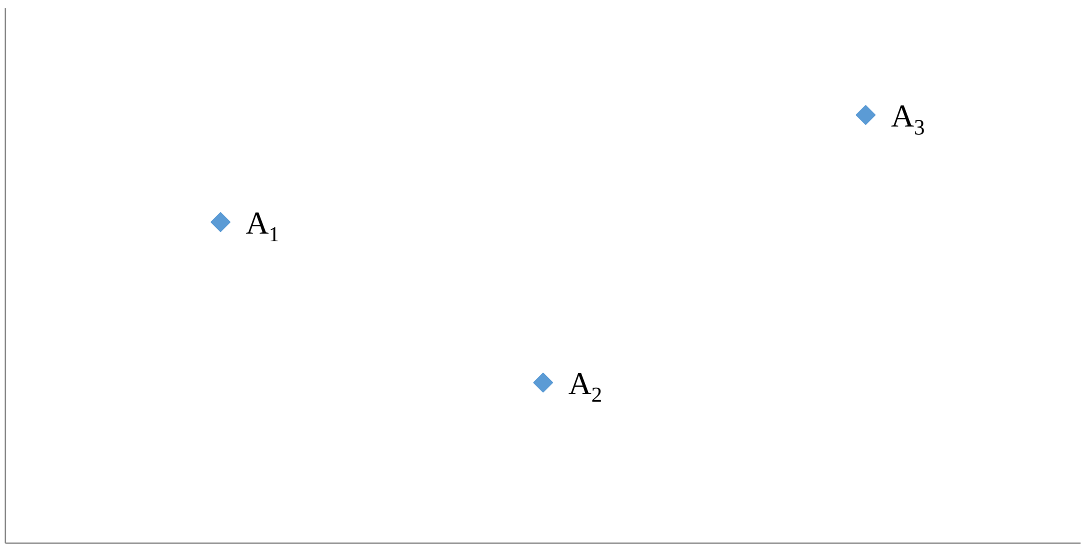

Figure 1. Joint Epistemic and Inductive Strength

This two-dimensional approach to evaluating arguments is applied below to three kinds of cases:

1) vertical and horizontal evaluations rely on point-valued probabilities;
2) vertical and horizontal evaluations rely on interval probabilities;
3) vertical and horizontal evaluations rely on non-numeric plausibilities (e.g., nearly certain, credible, doubtful, etc.).

The result is that, in each case, the argument's conclusion can be assigned a credence tag, as it were, that reflects a critical appraisal of its objectivity.

3. Conclusions with probabilistic hedges

Any conclusion derived from hedged hypotheses must itself be hedged. Obviously, the conclusion's hedge must be a function of the premises' hedges. But just how could the conclusion's hedge be reasonably determined?

I would like to propose an answer based on the following assumptions. I will assume, first of all, that the credences of premises can be represented by point-valued epistemic probabilities. While this assumption is not satisfied by most real-life argumentation, it offers the temporary advantage of concise illustration. Once the basic idea has been introduced, we can turn to more common, less precise representations of credence in Sections 4 and 5. Second, I will assume that the premises' conjunction E has a known epistemic probability $p(E)$. If the premises are independent, the probability of their conjunction falls out trivially from the probabilities of the premises. But if the premises are not independent, the conditional probability of one premise given the other premises and the joint probability of the other premises are also required. Third, I will assume that the inductive probabilities of the conclusion H given E and of H given E's negation are known. These are the conditional probabilities $p(H|E)$ and $p(H|\overline{E})$. Finally, I will make several trivial assumptions merely for the sake of illustration; these will be introduced as we proceed.

These assumptions suggest a typology of four cases:

I. Certain premises, valid argument: $p(E) = 1$, $p(H|E) = 1$.
II. Certain premises, invalid argument: $p(E) = 1$, $p(H|E) < 1$.
III. Uncertain premises, valid argument: $p(E) < 1$, $p(H|E) = 1$.
IV. Uncertain premises, invalid argument: $p(E) < 1$, $p(H|E) < 1$.

Drawing on the critical grid of Section 2, these four cases can be specified in terms of their horizontal (epistemic) and vertical (inductive) dimensions. They are represented schematically in Figure 2, where, for the sake of convenience, probabilities < 1 have been visualized temporarily as .5.

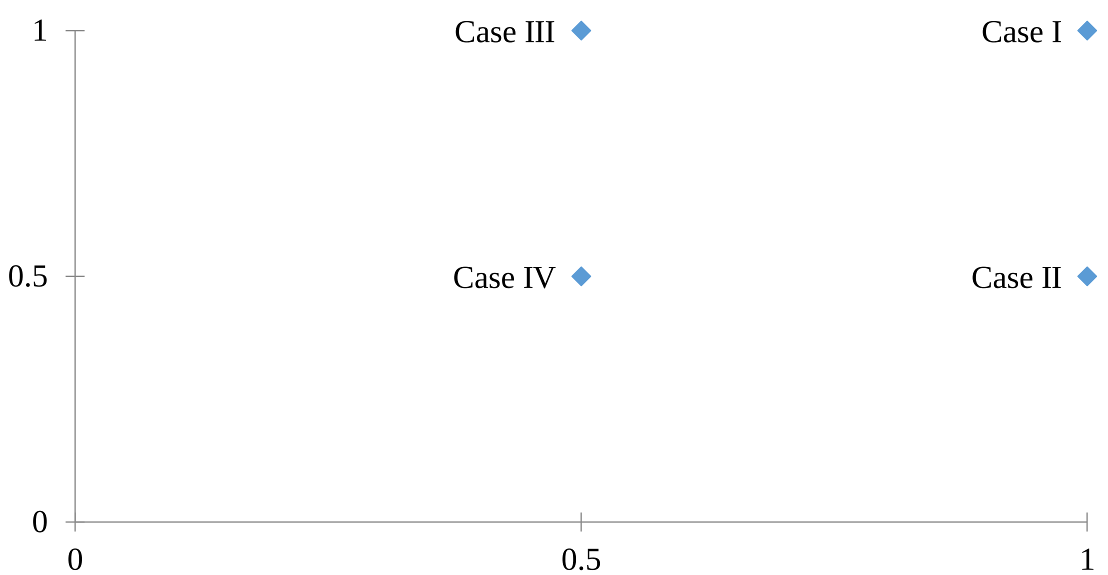

Figure 2. The Four Cases

The question is then whether the epistemic probability of the conclusion can be determined in each of these cases. I will suggest an answer in light of both Jeffrey's rule and Bayes' theorem.

3.1. Jeffrey's rule

The probability $p(H)$ of the conclusion must be considered in light of two possibilities: the premises' conjunction E is true, and the premises' conjunction $\overline{E}$ is false. Since H is equivalent to HE v $H\overline{E}$ and the clauses of this disjunction are mutually exclusive,

$$p(H) = p(HE \text{ v } H\overline{E}) = p(HE) + p(H\overline{E}). \qquad (1)$$

Applying the definition of conditional probability to each addend of the probability sum then yields

$$p(H) = p(E)p(H|E) + p(\overline{E})p(H|\overline{E}). \qquad (2)$$

(2) describes a sum of two *epistemic-inductive* (*EI) products*: the EI product for the evidence, which is $p(E)p(H|E)$, and the EI product for the negated evidence, which is $p(\overline{E})p(H|\overline{E})$.

I will refer to (2) as 'Jeffrey's rule', though Jeffrey proposed it for use in probability kinematics: changes from one probability assignment to another (Jeffrey, p. 169). Here the rule will be employed in what might be called probability statics: probability assignments that are fixed, at least for the moment.

Application of (2) to each of our four cases (with additional, purely illustrative assumptions for Cases II–IV) returns the following results:

I. $p(H) = 1$
II. $p(H) = .7$ where $p(H|E) = .7$
III. $p(H) = .76$ where $p(E) = .7$ and $p(H|\overline{E}) = .2$
IV. $p(H) = .55$ where $p(E) = .7$, $p(H|E) = .7$, and $p(H|\overline{E}) = .2$.

Note that in Cases I and II the certainty of the premises ensures that the epistemic probability of the conclusion equals the EI product for the evidence: $p(H) = p(E)p(H|E)$. The EI product for the evidence plays a key role in Cases III and IV as well, for it constitutes a lower bound for the epistemic probability of the conclusion. That is, $p(H)$ must be at least .7 in Case III and at least .49 in Case IV. In cases with uncertain premises, these lower bounds may be reached if, and only if, $p(H|\overline{E}) = 0$ and $p(H\overline{E}) = 0$.[1] For example, if E is the proposition P and H is the proposition PQ, logical considerations could determine that $p(H|E) = ½$ and that $p(H|\overline{E}) = p(H\overline{E}) = 0$. Here Jeffrey's rule collapses to the EI product for the evidence. Then as long as $p(E) = ½$, $p(H) = ¼$.

Cases with uncertain premises have an additional lower bound for $p(H)$: the EI product for the negated evidence. Consistent with this additional lower bound, $p(H)$ must be at least .06 in Cases III and IV. This lower bound may be reached in exceptional circumstances when $p(E) = 0$ or $p(H|E) = 0$.

3.2. Bayes' theorem

Bayes' theorem in its simplest form can be stated as:

---

[1] The definition of conditional probability ensures that $p(H|\overline{E}) = 0$ if, and only if, $p(H\overline{E}) = 0$.

$$p(H|E) = \frac{p(E|H)}{p(E)} \times p(H). \tag{3}$$

Reconfigured to reflect our focus on the epistemic probabilities of conclusions, the theorem specifies $p(H)$ as follows:

$$p(H) = \frac{p(H|E) \times p(E)}{p(E|H)}. \tag{4}$$

Applying the theorem to our four cases yields these values:

I. $p(H) = 1$
II. $p(H) = .7$ where $p(H|E) = .7$
III. $p(H) = .7/p(E|H)$ where $p(E) = .7$
IV. $p(H) = .49/p(E|H)$ where $p(E) = .7$ and $p(H|E) = .7$.

Necessarily, the results for Cases I and II coincide with those obtained from applying Jeffrey's rule. The results for Cases III and IV coincide also in the rather unhelpful sense that the Jeffrey value of .76 for Case III must be equal to the Bayes value of $.7/p(E|H)$, and the Jeffrey value of .55 for Case IV must be equal to the Bayes value of $.49/p(E|H)$. But EI products play different roles in Jeffrey and Bayes calculations. As we have seen, EI products in Jeffrey values for Cases III and IV set lower bounds for $p(H)$. But EI products in Bayes values for the same cases mark lower bounds for $p(E|H)$. For example, $p(E|H)$ cannot be less than .7 in Case III and less than .49 in Case IV on pain of violating the axioms of probability.

As a result, Jeffrey's rule has a marked advantage over Bayes' theorem in treating the problem at hand. In Cases III and IV, the lower bound identified by Jeffrey's rule is directly relevant to determining $p(H)$, whereas that marked by Bayes' theorem is not. In addition, since division by 0 is undefined, Bayes' theorem cannot be applied on the anomalous occasions when $p(E|H) = 0$. Finally, even though Bayes' theorem is just as useful as Jeffrey's rule for Cases I and II, it is not as useful for Cases III and IV. Hence the Jeffrey principle offers the advantages of one-stop shopping: one rule can be applied to all four cases. In ascertaining the value of $p(H)$, therefore, we would be better off to rely on Jeffrey's rule.

4. Conclusions with interval hedges

Section 1 proposed that premises are hedged hypotheses, each bearing a tag, as it were, to indicate its degree of credence. Section 3 showed how conclusions could also be assigned credence tags, provided the relevant credences could be expressed with point-valued probabilities. Unfortunately, as we noted, credences cannot usually be expressed with this degree of precision. But they can sometimes be expressed less precisely as probabilistic intervals. This might occur with Cases III and IV of Section 3, for example. Could we hedge a conclusion appropriately in these less favorable circumstances?

The answer is affirmative, and the procedure is relatively straightforward. We simply redouble our probabilistic efforts. For the uncertain premises of Cases III and IV, the probability $p(E)$ of the conjoined premises is calculated twice: once with the lower bound of the interval for each premise and once with the upper bound of the interval for each premise. The resulting probabilities delimit a range of epistemic probabilities for $p(E)$. Then, assuming the inductive probabilities $p(H|E)$ and $p(H|\overline{E})$ are also intervals, Jeffrey's rule can be applied twice: once with the lower values and once with the higher values. The result is a range of values for $p(H)$.

As a trivial example, consider this invalid argument:

$$P \rightarrow Q$$
$$Q$$
$$\text{Therefore } P.$$

Here $p(H) = p(P)$ and $p(E) = p((P \rightarrow Q) \wedge Q)$. Suppose that $p(P \rightarrow Q) = .75 – .80$, that $p(Q) = 2/3 – 3/4$, and that the premises are independent. Hence $p(E) = .5 – .6$. Suppose also that $p(H|E) = .5 – .6$ and that $p(H|\overline{E}) = .4 – .5$. Then we can apply Jeffrey's rule twice to find $p(H)$. We apply it once for the lower values of $p(E)$, $p(H|E)$, and $p(H|\overline{E})$:

$$(.5 \times .5) + (.5 \times .4),$$

and once for the higher values:

$$(.6 \times .6) + (.4 \times .5).$$

Consequently, $p(H)$ must lie in the interval .45 – .56.

5. Conclusions with plausibilistic hedges

The situation can get worse, of course. In dealing with point-valued and interval probabilities, we have assumed that $p(E)$ is known. But what if the credence of the premises can be represented only with qualitative terms such as 'nearly certain' or 'doubtful'? All is not lost, for plausibility measures can map propositions to members of any partially ordered set (Friedman and Halpern, p. 176). But, as one might expect, qualitative plausibilities complicate matters considerably.

To investigate the possibility of a plausibilistic analogue of $p(E)$, let us assume that we are equipped with a plausibility function π that maps propositions to plausibility values. For the sake of simplicity, let there be a minimal set of these values: superior (S), average (A), and inferior (I), ordered in the obvious way. The plausibility of the premises' conjunction might then be determined by appeal to a variant of John Pollock's weakest link principle: "The degree of support of the conclusion … is the minimum of the degrees of support of its premises" (Pollock, p. 99).[2] For example, if the evidence E for an argument's conclusion consists of three premises with plausibilities S, A, and I, then the weakest link principle stipulates that $\pi(E) = I$. The weakest link serves as an upper bound for the plausibility of the conjoined premises.

In working with Jeffrey's rule in Sections 3 and 4, the epistemic probabilities $p(E)$ and $p(\overline{E})$ were required, but the general negation rule permitted $p(\overline{E})$ to be inferred from $p(E)$. Similarly, might we infer $\pi(\overline{E})$ from $\pi(E)$? Not with quantitative precision, of course, but there is a reciprocal

[2] This variant of the weakest link principle differs from Pollock's principle in two ways. Whereas Pollock's principle governs the degree of support for the conclusion, the variant regulates the credence of the conjoined premises. In addition, though Pollock limited his principle to "deductive arguments," I propose to employ it for arguments with uncertain premises, whether valid or not.

relation between the plausibility of a proposition and that of its negation. Given our sample scale of plausibilities, if $\pi(E) = S$, then $\pi(\overline{E}) = I$; and if $\pi(E) = A$, then $\pi(\overline{E}) = A$.

Jeffrey's rule also relies on the inductive probabilities $p(H|E)$ and $p(H|\overline{E})$. These probabilities could conceivably be combined with the epistemic plausibilities $\pi(E)$ and $\pi(\overline{E})$. The result would be the hybrid hedges described in Section 5.1. But if these probabilities are unavailable, we might nonetheless be able to specify the inductive plausibilities $\pi(H|E)$ and $\pi(H|\overline{E})$. The combination of non-numeric epistemic and inductive plausibilities would produce the pure hedges treated in Section 5.2.

5.1. Hybrid hedges

In working with probabilistic credences, we assumed that the inductive probabilities $p(H|E)$ and $p(H|\overline{E})$ were at hand. The same assumption may not be unrealistic here, for inductive probabilities may be known on empirical or logical grounds. For example, if the propositions expressing H and E are unquantified, Carnap's λ-continuum (Carnap 1952) and its successor Basic System (Carnap 1971, 1980) permit inductive probabilities to be determined on the basis of an empirical factor and a logical factor. If there are no empirical data to draw on, Carnapian systems can derive inductive probabilities from logical considerations alone. Although these systems have a well-known weakness in dealing with analogy, the problem can be addressed straightforwardly by supplementing the empirical and logical factors with an analogy factor (Kuipers, p. 69; Welch, pp. 241–242). On the other hand, if the propositions expressing H and E are quantified, Hintikka's α-λ continuum (Hintikka), which extends the λ-continuum to improve the handling of inductive generalization, and Hintikka and Niiniluoto's *K*-dimensional system (Hintikka and Niiniluoto), which axiomatizes much of the α-λ continuum, can be used to evaluate $p(H|E)$ and $p(H|\overline{E})$.

Sections 3.1–3.2 noted the role of EI products in determining the point-valued epistemic probabilities of conclusions. Here there is no question of an EI product, for we are assuming that the credence of the evidence is represented by a non-numeric plausibility while the credences of the conclusion given the evidence and the conclusion given the negated evidence are represented by probabilities. But we could rely on *EI indicators*, which represent credence in hybrid fashion by combining an epistemic plausibility like *S*, *A*, or *I* with an inductive probability. Although the operation of multiplication is undefined for such disparate domains, EI indicators could be formed by simple juxtaposition of epistemic and inductive terms. *A*.5 is a straightforward example.

EI indicators may or may not present problems of comparability. The indicator *A*.5, for instance, is plainly superior to *I*.4. But we could not compare the indicators *S*.5 and *A*.6, since the epistemic plausibility of the first is superior to that of the second while the inductive probability of the second is superior to that of the first.

Jeffrey's rule determines the probability of *H* as a sum of the EI product for the evidence and the EI product for the negated evidence. Extending this logic to plausibilities in general, we might say that the plausibility of *H* depends on both the EI indicator for the evidence and the EI indicator for the negated evidence. We will not be able to add these EI indicators, of course, but we can consider them jointly. For example, *A*.5, *A*.3 could represent the plausibility of *H* based on the EI indicator for the evidence (*A*.5) and the EI indicator for the negated evidence (*A*.3).

Since some EI indicators are comparable and others are not, the plausibilities of conclusions determined by joint consideration of EI indicators may be comparable or not. A conclusion characterized by the indicators *S*.6, *I*.3 is more plausible than a conclusion with the indicators *S*.5, *I*.1. But the plausibility of a conclusion with indicators *S*.6, *I*.3 cannot be compared to the plausibility of a conclusion with indicators *S*.5, *I*.5, for the first conclusion fares better on the indicators for the evidence while the second conclusion does better on the indicators for the negated evidence.

Let us beware of exaggerating the problem, however. The principal practical problem in argument evaluation is to determine whether a conclusion *H* or its contradictory –*H* is a better inference from a given set of premises. In making these judgments, the problem of incomparability is ameliorated, though not eliminated, by two considerations. The first is that the epistemic plausibilities of the evidence and the negated evidence are the same for *H* as they are for –*H*, since both conclusions are based on the same premises. Consequently, the problem can be addressed entirely on the basis of the inductive probabilities of the respective EI indicators. The second consideration is that the inductive probabilities for –*H* may be inferred from those for *H*, for $p(H|E) + p(\overline{H}|E) = 1$. For example, if the plausibility of *H* is represented as:

$$\pi(E)p(H|E),\ \pi(\overline{E})p(H|\overline{E}) = S.7,\ I.6,$$

then the plausibility of –*H* must be:

$$\pi(E)p(\overline{H}|E),\ \pi(\overline{E})p(\overline{H}|\overline{E}) = S.3,\ I.4.$$

As a result, the crucial inductive probabilities are all comparable, and *H* in this instance is the better inference. But this does not prevent the indicators for the evidence from supporting a different verdict than the indicators for the negated evidence. If the plausibility of *H* is represented as:

$$\pi(E)p(H|E),\ \pi(\overline{E})p(H|\overline{E}) = S.4,\ I.7,$$

and the plausibility of –*H* as:

$$\pi(E)p(\overline{H}|E),\ \pi(\overline{E})p(\overline{H}|\overline{E}) = S.6,\ I.3,$$

the indicators for the evidence favor –*H* while those for the negated evidence support *H*.

5.2. Pure hedges

Even if it is possible in principle to determine the inductive probability of the conclusion on the premises, we may have neither the time nor the expertise to do so. In such cases, the best we could do would be to represent the credence of the conclusion based on the premises in terms of non-numeric plausibilities. Our EI indicators would no longer be hybrids of epistemic plausibilities and inductive probabilities; they would be pure compounds of non-numeric plausibilities.

There are two scenarios to consider. In the first, the same set of values represents both epistemic and inductive plausibility. In the second, one set of values expresses epistemic plausibility while another expresses inductive plausibility. The differences are relatively unimportant. The first scenario admits comparable EI indicators like *AS* and *IS*, but it also allows incomparable indicators such as *AS* and *SA*. The second scenario does much the same thing. If *S*, *A*, and *I* constitute the epistemic vocabulary and *s*, *a*, and *i* the inductive one, we will have comparable EI indicators like *Sa* and *Ai* as well as incomparable ones such as *Sa* and *As*.

In working with the hybrid hedges of Section 5.1, we extended the logic of Jeffrey's rule. Rather than treat the probability of *H* as a sum of EI products, we assessed the plausibility of *H* by jointly considering hybrid EI indicators. An analogous approach can be adopted here. Although we lack inductive probabilities, we might still assess the plausibility of *H* by jointly considering pure EI indicators: one for the evidence, and one for the negated evidence. As in the hybrid case, we cannot add pure EI indicators, but we can juxtapose them to form an estimate of *H*'s plausibility. For example, *Sa*, *Ii* might represent the plausibility of *H* based on the EI indicator for the evidence (*Sa*) and the EI indicator for the negated evidence (*Ii*).

Pure hedges have the same spotty record of comparability that we noticed in connection with hybrid hedges: the plausibilities of conclusions determined by joint consideration of EI indicators may be comparable or not. A conclusion with the indicators *Sa*, *Ia* is more plausible than a conclusion with the indicators *Si*, *Ii*. But a conclusion with the indicators *Sa*, *Ia* cannot be compared to a conclusion with the indicators *Si*, *Is*, since the first conclusion's indicator for the evidence is superior to that of the second, while the second conclusion's indicator for the negated evidence is superior to that of the first.

Nevertheless, as noted for the hybrid case, we need not exaggerate the problem. Section 5.1 invoked two considerations to ameliorate the practical problem of inferring *H* or –*H* from a given set of premises. Analogous considerations apply here. First, the epistemic plausibilities of the evidence and the negated evidence are the same for *H* as they are for –*H*. As a result, the problem can be addressed on the basis of the inductive plausibilities in the respective EI indicators. Second, we observed in the hybrid case that the inductive probability of –*H* could be inferred from that of *H*, since $p(H|E) + p(\overline{H}|E) = 1$. This degree of quantitative precision is unattainable with pure indicators, but there is nonetheless a reciprocal relation between the inductive plausibilities for *H* and for –*H*. To illustrate with our sample scale of plausibilities: if $\pi(H|E) = s$, then $\pi(\overline{H}|E) = i$; and if $\pi(H|E) = a$, then $\pi(\overline{H}|E) = a$. Hence if the plausibility of *H* is represented as:

$$\pi(E)p(H|E),\ \pi(\overline{E})p(H|\overline{E}) = Ss,\ Ia,$$

then the plausibility of –*H* must be:

$$\pi(E)p(\overline{H}|E),\ \pi(\overline{E})p(\overline{H}|\overline{E}) = Si,\ Ia.$$

Here *H* is the better inference. But, as in the hybrid case, incomparability can still raise its enigmatic head. If the plausibility of *H* is represented as:

$$\pi(E)p(H|E),\ \pi(\overline{E})p(H|\overline{E}) = Si,\ Is,$$

and the plausibility of –*H* as:

$$\pi(E)p(\overline{H}|E),\ \pi(\overline{E})p(\overline{H}|\overline{E}) = Ss,\ Ii,$$

the indicators for the evidence favor –*H* while those for the negated evidence support *H*.

6. Argumentative standards with hedges?

Section 2 of this paper proposed a two-dimensional grid for evaluating arguments. The horizontal dimension of the grid represents the epistemic properties of the premises, while the vertical dimension represents the inductive flow from premises down to conclusion. An argument's representation on this grid is immediately relevant to our usual standards of argumentation. The horizontal dimension responds to the conditions these standards impose on argument content, and the vertical dimension to the conditions they impose on argument form. Let us review three of these standards.

Deductive soundness
Content: The premises must be true.
Form: The form must be valid, i.e., the conditional probability of the conclusion on the premises must be 1.

Inductive cogency
Content: The premises must be true.
Form: The form must be inductively strong, i.e., the conditional probability of the conclusion on the premises must be greater than or equal to that of any rival conclusion based on the same premises.

Plausibilistic credibility[3]
Content: The premises must be true.
Form: The form must be plausibilistically strong, i.e., the plausibility of the conclusion given the premises must be greater than or equal to that of any rival conclusion based on the same premises.

Note that the content conditions of the three standards are identical, but the formal conditions grow progressively weaker as we move from soundness to cogency to credibility. The formal conditions are intimately related, however. As Wittgenstein pointed out, "The certainty of logical inference is a limiting case of probability" (Wittgenstein, 5.152); hence validity is a special case of inductive strength. In addition, since the standard probability axioms are special cases of plausibility axioms (Chu and Halpern, pp. 209–210), inductive strength is a special case of plausibilistic strength.

In evaluating an argument relative to any of these standards, epistemic meta-questions can (and should) be asked: 'I think the premises are true', for example, 'but how sure am I?' One response to these questions would be to augment the argumentative standard by embedding an epistemic meta-condition within it. For example, we might judge that an argument is sound and that we have reason to be certain of both the truth of the premises and the validity of the argument's form. Such an argument would approximate an Aristotelian demonstration.

Though some may find this maneuver attractive, an alternative approach seems preferable to me. As customarily stated, the standards are objective. If an argument meets the standard of deductive soundness, for example, its premises *are* true and its form *is* valid. Incorporating an epistemic meta-dimension within an argumentative standard would create an uneasy mix of objective and subjective considerations. Better, I think, to handle our epistemic meta-concerns as

[3] To my knowledge, the term 'plausibilistic credibility' is not as widely accepted as 'deductive soundness' and 'inductive cogency,' but there is a clear need for some such term.

questions of credence, pure and simple. After all, the payoff in evaluating an argument is the ability to specify the credence of its conclusion: Based on these premises, how sure can I be of this conclusion? As this paper has shown, a degree of credence can frequently be assigned to a conclusion on the basis of the epistemic credences of its premises and the inductive credence of the conclusion given the premises. In such cases, then, we can obtain our critical payoff, deciding in a reasonable fashion how much weight to place on a conclusion, without diluting the objectivity of our usual standards of argumentation. The conclusion ends up with a credence tag, but not the argumentative standard.

Acknowledgements: I am grateful to Michael Shenefelt and Heidi White for helpful comments on an earlier version of this paper.

References

Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science* 3, 207–220.

Carnap, R. (1952). *The Continuum of Inductive Methods*. Chicago: University of Chicago Press.

Carnap, R. (1971). Inductive logic and rational decisions. A basic system of inductive logic, part 1. In: R. Carnap and R. C. Jeffrey (Eds.), *Studies in Inductive Logic and Probability* (vol. I, pp. 5–31, 33–165), Berkeley: University of California Press.

Carnap, R. (1980). A basic system of inductive logic, part 2. In: R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability* (vol. II, pp. 7–155), Berkeley: University of California Press.

Chu, F. C. and J. Y. Halpern. (2004). Great expectations. Part II: Generalized expected utility as a universal decision rule. *Artificial Intelligence* 159, 207–229.

Friedman, N. and J. Y. Halpern. (1995). Plausibility measures: A user's guide. In: P. Besnard and S. Hanks (Eds.), *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* (*UAI '95*) (pp. 175–184), San Francisco: Morgan Kaufmann.

Hintikka, J. (1966). A two-dimensional continuum of inductive methods. In: J. Hintikka and P. Suppes (Eds.), *Aspects of Inductive Logic* (pp. 113–132), Amsterdam: North-Holland.

Hintikka, J. and I. Niiniluoto. (1976). An axiomatic foundation for the logic of inductive generalization. In: M. Przełęcki, K. Szaniawski, and R. Wójcicki (Eds.), *Formal Methods in the Methodology of Empirical Sciences* (pp. 57–81), Dordrecht: D. Reidel and Wrocław: Ossolineum.

Jeffrey, R. C. (1983). *The Logic of Decision*, 2nd ed. Chicago: University of Chicago Press.

Kuipers, T. A. F. (1984). Two types of inductive analogy by similarity. *Erkenntnis* 21, 63–87.

Pollock, J. L. (2003). *Cognitive Carpentry: A Blueprint for How to Build a Person*. Cambridge, MA: MIT Press.

Welch, J. R. (1999). Singular analogy and quantitative inductive logics. *Theoria* 14, 207–247.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul.