Script for "Logical Foundations of Evidential Support"
(Do not quote or cite without permission.)

Preamble, before slides...

The handouts contain the main slides that I will be projecting up on the screen. **The slides provide an outline the main points I'll be making, but my script is slightly more verbose, and adds various additional remarks**.

Slide 0 (not on handout, only on screen):

This presentation is called "Logical Foundations of Evidential Support", which is something of a tribute to Rudolf Carnap, whose book "Logical Foundations of *Probability*" was the motivating force behind this project. While I will deviate from the *letter* of Carnap's proposals in several important ways, I will try to recapture what I think is the *underlying spirit* of Carnap's project. I think Carnap's project was abandoned too soon, and for the wrong reasons.

Slide 1:

In the first part of the talk, I will discuss the contemporary project of Bayesian confirmation theory, which is a distant (I should say, *very* distant) descendant of Carnap's own confirmation project. I will briefly provide some background material on Bayesian confirmation, and then I will describe a serious problem (called "the measure sensitivity problem") faced by the modern Bayesian account of confirmation or evidential support.

In the second part of the talk, I will backtrack, chronologically. I will provide a historical and conceptual trace of the development of inductive logic. This will involve identifying three historical desiderata for the inductive logic project, and then examining various existing attempts to produce an adequate explication of a *logical* concept of confirmation that satisfies these desiderata. In the end, I will argue that the historical proposals — in particular, Carnap's — are lacking in at least two crucial ways, and I will suggest a new, "Neo-Carnapian" approach to inductive logic. Then, I will show how this approach provides both a solution to the Bayesian problem of measure-sensitivity, and a general, logical foundation for Bayesian and non-Bayesian probabilistic (or statistical) notions of evidential support.

In closing, I will try to place this work in a broader philosophical context, mainly by discussing its epistemological significance.

(I)

Slide 2:

Bayesianism (which has recently been called "Bayesian epistemology" – although I think that's a misnomer, as I'll explain later) assumes that the degrees of belief (or credence) of rational agents are *probabilities*, with certainty of truth corresponding to probability 1, certainty of falsehood corresponding to probability 0, and uncertainty being gauged by intermediate values of a probabilistic credence function.

Let Pr(H) denote the (unconditional) degree of belief some rational agent *S* assigns to proposition H (at some time *t*). I will call this S's *prior* probability of H (at *t*).

Let Pr(H | E) be the degree of belief S assigns to H (at *t*), *on the supposition that E is true* (or, if you prefer, the degree of belief S *would* assign to H upon learning that E is true with certainty). This is called S's *posterior* probability of H (on E, at *t*). For simplicity, I will hereafter omit explicit reference to S and *t*, but it is to be kept in mind that Bayesian probabilities are always indexed to specific rational agents S at specific times *t*. We will return to this subjective aspect of contemporary Bayesianism later on.

Here is a simple and intuitive toy example (just to fix our ideas, and to get our probabilistic juices flowing): Let H be the hypothesis that a card (to be drawn at random from a standard 52 card deck) is a spade, and let E be the proposition that the card is the ace of spades. In this case (assuming the standard probability model of random card draws, which assigns equiprobability to each card, etc.), we have that the prior probability of H is 1/4, since 1/4 of the cards are spades, but the posterior probability of H, given E is 1, since the card *must* be a spade on the supposition that it is the ace of spades. So, intuitively, learning E raises the probability of H here. Indeed, E *verifies* H in this case.

Slide 3:

According to contemporary Bayesianism, evidence E confirms (or supports) a hypothesis H if learning E raises the probability of H — that is, if the *posterior* probability of H (given E) is greater than the *prior* probability of H.

If learning E lowers the probability of H, then E disconfirms (or counter-supports) H, and if learning E has no effect on the probability of H, then E is confirmationally irrelevant to (or neutral regarding) H.

For modern Bayesians, then, confirmation involves *relevance*, and – in particular – *probabilistic relevance*. I will argue later on that this is a *virtue* of modern Bayesian confirmation theory.

In probability theory, there are many logically equivalent ways of saying that E raises the probability of H. As such, there are many logically equivalent ways of says that E confirms H in the modern, Bayesian sense. Here are three such ways:

…explain these…mention likelihoods…compute numbers…
– E confirms H if $Pr(H \mid E) > Pr(H)$ [$1 > 1/4$ in our toy example]
– E confirms H if $Pr(E \mid H) > Pr(E \mid {\sim}H)$ [$1/13 > 0$ in our toy example]
– E confirms H if $Pr(H \mid E) > Pr(H \mid {\sim}E)$ [$1 > 12/51$ in our toy example]

Intuitively, if one wanted to measure the *degree* to which E confirms H, one could take *differences* (or ratios, etc.) of the left and right hand sides of any of these inequalities. For instance, the difference between the posterior $Pr(H \mid E)$ and prior $Pr(H)$ is one intuitive measure of the degree to which E confirms H, since it is positive if E confirms H, negative if E disconfirms H and zero if E is confirmationally irrelevant to H. Such measures are called *relevance* measures. In this way, *many* syntactically different quantitative Bayesian relevance measures of degree of confirmation can be formed.

Slide 4:

Dozens of (syntactically distinct) Bayesian relevance measures of confirmation have been proposed and defended in the philosophical literature over the past century or so (dating back at least to Keynes' teacher W.E. Johnson around the turn of the 20th century). Here are the four most popular (and representative) contemporary Bayesian measures of confirmation:

… go over the measures here … mention *names*: d, r, l, s … decreasing order of popularity in the literature

Here, I take logarithms of the ratios, just to ensure that the measures yield positive numbers for confirmation, negative numbers for disconfirmation, and zero for neutrality. That is, to ensure that the ratio measures are easily identifiable as *relevance measures*. This is just a useful convention, which does not affect the ordinal or comparative structure of the measures (since Logs are strictly increasing functions).

The first part of my presentation concerns the nature of the disagreement exhibited by these measures, and its ramifications for Bayesian confirmation theory …

Slide 5:

All Bayesian relevance measures of confirmation (including our four) agree on *qualitative* judgments of the form "E confirms, disconfirms, or is irrelevant to H." This is because all such measures are positive in cases of confirmation, negative in cases of disconfirmation, and zero in cases of irrelevance.

Surprisingly, not only do these (and many other) relevance measures disagree *syntactically* or *numerically* (that is to say, *conventionally*), but they also disagree *radically* on *comparative* or *ordinal* judgments of the form "E1 confirms H1 more strongly than E2 confirms H2."

This ordinal or comparative disagreement has serious consequences for a wide variety of arguments in the literature (many of which have previously gone unnoticed and/or unappreciated).

For instance, most famously, it is part of Bayesian lore that the observation of a black raven (E1) confirms the hypothesis that all ravens are black (H) more strongly than the observation of a red herring (E2) does. But, this conclusion *depends* *sensitively on one's choice of confirmation measure*. That is, one can render Bayesian arguments for this conclusion about Hempel's Ravens Paradox invalid, merely by changing the confirmation measure used in the argument — while holding all other assumptions fixed. When an argument can be made invalid *simply* by switching confirmation measures, we say that the argument is *sensitive to choice of measure of confirmation*. Sadly, *almost all* comparative arguments in the Bayesian literature are sensitive to choice of measure! In my dissertation, I provide a historical survey of dozens of measure-sensitive arguments. And, since then, many more measure-sensitive arguments have been discovered.

Slide 6:

This table enumerates 9 of my favorite measure-sensitive arguments from the Bayesian literature (this is sort of my "top 9"). These arguments encompass many of the central problems and applications of confirmation theory (including Bayesian treatments of Hempel's Ravens paradox and Goodman's New Riddle of Induction, or "Grue"). There are two important things to notice about this table. First, no argument has "Yes"s all the way across its row. This means that no argument is valid for *all* choices of *d*, *r*, *l*, or *s*, as measure of confirmation. *That, in a nutshell,* is the problem of measure-sensitivity. Second, note that both arguments *within* the Bayesian framework (7 on top) as well as *critiques of* Bayesian confirmation theory (2 on bottom) fall prey to the measure-sensitivity problem. Critiques of Bayesianism due to Popper-Miller and Earman are two well-known examples. Time limitations prevent me from going over the details of these measure-sensitive arguments, so I will now bring this first part of my presentation to a close.

--- SHOW OVERVIEW AGAIN --- Click on "Bayesianism" at bottom of Table

Let's Take Stock: So far, we have had a brief introduction to Bayesian confirmation theory, and we have also seen a serious problem with that theory. There are many, radically different Bayesian measures of degree of confirmation. And, almost every comparative argument in the literature on Bayesian confirmation theory depends for its cogency on the preference of certain measures over others. Somewhat embarrassingly, there have not been many attempts in the literature to provide a compelling argument in favor of one of these measures over the others. And, the few existing attempts to resolve this problem tend to be highly technical in nature, and not very satisfying or compelling. This was the subject of my dissertation.

In the second part of the presentation, I will take a look back at the historical project of probabilistic inductive logic. Here, I will try to follow the *spirit* of the Carnapian inductive-logic programme. But, my discussion will diverge from Carnap's approach in several important respects. In the end, I will propose a new, "Neo-Carnapian" account of inductive logic. This will allow us to resolve the Bayesian measure-sensitivity problem, and to provide a general logical foundation for both Bayesian and non-Bayesian probabilistic accounts of evidential support.

… **water**…

Click on "(II)" on overview slide to jump to part II

(II)

Slide 7:

Now, I am going to shift gears and backtrack, historically. I am going to outline a new way to think about inductive logic. But, before I get to my own account, I want to look at some historical proposals and conceptions. I begin with some quotes from Carnap's LFP. The first quote gives us a sense of the basic idea behind inductive logic, as a quantitative analogue (or generalization) of deductive logic:

``Deductive logic may be regarded as the theory of the relation of logical consequence, and inductive logic as the theory of another concept [c] which is likewise objective and logical … degree of confirmation.''

The next two quotes from Carnap provide intuitive, pre-theoretical characterizations of what it means for inductive logic to be *logical* and *objective*, respectively.

``The principal common characteristic of the statements in both fields is their independence of the contingency of facts. This characteristic justifies the application of the common term `logic' to both fields.''

``That c is an objective concept means this: if a certain c value holds for a certain hypothesis with respect to a certain evidence, then this value is entirely independent of what any person may happen to think about these sentences."

In what follows, I will try to capture the *spirit* of the project described by these intuitive remarks of Carnap's.

Slide 8:

Abstracting away from the idiosyncrasies of the vast array of specific historical attempts to explicate a *logical* concept of confirmation (especially, Carnap's), we can extract the following three (nearly universal) central desiderata from the historical literature on inductive logic:

[Draw diagram on board for D1 -- then complete it later]

D1: Inductive logic should provide a quantitative generalization of deductive logic. That is, the relations of deductive entailment and deductive refutation should be captured as limiting (extreme) cases of confirmation (disconfirmation) with cases of `partial entailment' and `partial refutation' lying somewhere on a c-continuum (or range) between these extremes.

D2: Inductive logic should use probability (in its technical, modern sense) as its central conceptual building block. [This is mainly a *historical* rather than a *material* desideratum.]

D3: Inductive logic (i.e., the non-deductive relations between propositions that are characterized by inductive logic) should be *objective* and *logical*.

Here, I will follow Carnap's *intuitive, pre-theoretical* characterizations of the terms `logical' and `objective' (given above in the quotes from LFP). As we will see, I will later *diverge* from the *letter* of Carnap's precise, *formal* explications of these concepts.

Slide 9:

Here's a Naive attempt to provide a generalization of deductive entailment using Probability:

According to classical deductive logic (here, I have in mind propositional logic, but this can all be generalized to richer logical languages), premises P1, …, Pn entail a conclusion C if the material conditional P hook C is (logically) necessarily true (I will use P to abbreviate the conjunction of the premises P1, … , Pn).

Naively, then, one might try to quantitatively generalize deductive entailment by saying that the argument from P to C is inductively strong if the material conditional P hook C is *probably* (but not *necessarily*) true. This suggests the following proposal:

(NIL) $c(C, P) = Pr(P \supset C) = Pr(\text{either not-P or C})$ … explain …

As Brian Skyrms explains in his book *Choice and Chance*, (NIL) will not do. Consider the following argument of Skyrms':

> (P) There is a man in Cleveland who is 1999 years and 11-months-old and
> in good health. Therefore, (C) No man will live to be 2000 years old.

Intuitively, this argument is *not* strong, since (P) seems to *dis*confirm (C). That is, the existence of a 1999+ year old healthy man seems to provide evidence that the conclusion (which says that no man will live to be 2000) is *false*. But, intuitively, the probability of the material conditional is *high* — *merely because* the conclusion (C) *itself* is highly probable *unconditionally* (basically, we're just making a disjunction probable here by probabilifying one of its disjuncts). The crucial point here is that the corresponding material conditional is highly probable, but *not* because of any confirmation (or evidential) *relation* between the premises and the conclusion. Hence, as Skyrms rightly points out, this is a counterexample to (NIL).

Footnote: Skyrms doesn't mention this, but (NIL) doesn't even satisfy desideratum (D1). If the premises refute the conclusion, then according to (NIL), the degree of inductive strength of the argument is the probability that the premise is false (which is not minimal and in fact can be quite high). So, (NIL) will *not* serve as an adequate explication of confirmation.

Slide 10:

According to Skyrms, (NIL) conflates the probability of the material conditional P hook C with the conditional probability of C, given P; and, it is the latter (not the former), according to Skyrms, which captures the confirmation *relation*. This is "The Received View" of inductive logic, which has been *very* widely accepted by inductive logicians since the inception of the discipline:

(TRV) $\mathrm{c}(\mathrm{C}, \mathrm{P}) = \Pr(\mathrm{C} \mid \mathrm{P}) = \dfrac{\Pr(P \,\&\, C)}{\Pr(P)}$ … explain …

It is clear that (TRV) handles Skyrms' example better than (NIL) does. Intuitively, given that there is a 1999 year 11-month-old healthy man in Cleveland, the probability that no man will live to be 2000 is *not* high. And, so, according to (TRV) this argument is *not* a strong one, which seems to be the correct verdict.

It is also clear that (TRV) satisfies the first two of our historical desiderata:

D1: If P entails C, then $\Pr(\mathrm{C} \mid \mathrm{P}) = 1$, which is maximal, and if P refutes C (i.e., if P entails ~C), then $\Pr(\mathrm{C} \mid \mathrm{P}) = 0$, which is minimal (and there is a range of values in between). So, (TRV) does define ***a*** quantitative generalization of deductive entailment and refutation in the sense of D1.

D2: (TRV) *identifies* the confirmation function with the conditional probability function, and so (of course) it makes probability a central concept in its explication. [To wit: the title of Carnap's book LF***P***]

What about desideratum D3? Does (TRV) provide us with a conception of confirmation that is objective and logical? That seems to depend on how (TRV) *interprets* the conditional probability function Pr.

Slide 11: **… don't speed up…relax…**

If one accepts (TRV), then it seems that one must provide a "logical" or (*a priori*) conditional probability function Pr in order to satisfy desideratum D3.

Many philosophers — including Leibniz, Wittgenstein, Keynes, Waismann, Carnap, and many others — have tried to explicate something like the following "logical" concept of conditional probability. Very roughly and heuristically, we might try something like the following idea:

$$\Pr(\mathrm{C} \mid \mathrm{P}) = \frac{\text{the proportion of logically possible worlds in which both P and C are true}}{\text{the proportion of logically possible worlds in which P is true}}$$

It is *notoriously* tricky to make this vague and informal notion precise. It is *so* tricky, in fact, that no satisfying account of "logical" or *a priori* probability has ever been found (the consensus now seems to be that no such account is forthcoming). The main problem with this sort of project has been trying to justify a choice of *measure* of the "proportion of logically possible worlds in which a proposition is true".

Early attempts to explicate logical probability tried to apply the principle of indifference to justify assigning *equal a priori* probability to various partitions of logical space. The basic idea behind this principle is (roughly) that if a collection of evidence does not favor one element of a partition of logical space over any other, then – relative to that collection – each cell in the partition is equiprobable. As stated, the principle seems to be an epistemic one (as it is relative to a collection of evidence), which makes it somewhat surprising that the early Carnap endorsed it as a *logical* principle. Early on, Carnap used to it motivate assigning equiprobability *a priori* to the *state descriptions* in certain logical languages [we can discuss Carnap's theory of "logical probability" further in the question period if people are interested – I have included some auxiliary slides on that]. Carnap ultimately rejected this assignment of equal *a priori* probability to state descriptions, because it led (in his systems) to probability functions that could not reflect any *correlations* between the attributes of different objects in the universe of discourse [Carnap and others have characterized that as not supporting "learning from experience"]. And, he realized that such correlations are essential for inductive and statistical inference. So, Carnap went on to apply the principle of indifference to a *different* partition: the *structure* descriptions of his languages. But, there were still various kinds of correlations (analogical ones) that these probability functions could not model. Similar problems kept reappearing in more and more subtle ways as Carnap made his syntactical constructions more and more complex in an attempt to overcome this fundamental limitation of his systems. But, in the end, none of his accounts of "logical" probability exhibited the requisite flexibility to undergird all the kinds of correlations we think are essential to inductive inference. This was rather frustrating for Carnap. And, Carnap was not alone in his frustration. The search for "logical" or "a priori" probability has eluded many great philosophers and probabilists over a span of hundreds of years, and nobody even seems to be getting warm…

Alas, it seems that there is no "logical aether" with respect to which propositions receive their "absolute probabilities". Just as there is no "physical aether" with respect to which physical objects obtain their "absolute velocities". Probability seems to be an *inherently relational* property of propositions — a property they only seem to have, *relative* to specified *probability models*. Just as velocity seems to be an inherently relational property of physical objects — a property they only seem to have relative to specified *frames of reference* [that's a hokey analogy — but, I see that Nozick makes this same hokey analogy in his most recent book *Invariances*]. This "relativity of probability" is the main reason that contemporary Bayesians have moved toward *subjective* or *epistemic* accounts of probability which are always *indexed* to rational-agents-at-times or to collections of background evidence, *etc*.. Almost all Bayesians have given up hope of finding "the" *a priori* credence function. However, in doing so, they have moved toward *epistemically* or *pragmatically relative* probabilities, which for all practical purposes constitutes an abandonment of desideratum D3. In other words, modern probabilists have all but given up on the traditional project of Inductive Logic.

Since my goal today is precisely *not* to give up on the traditional inductive logic project, I will need to do something different. If, as it appears, probability judgments *are* inherently relational, and that there is no adequate theory of "logical" or "a priori" probability forthcoming, then we will need an alternative way to satisfy desideratum D3.

I propose a "Neo-Carnapian" Solution…

… **water**…

Slide 12:

When it came to deductive logic, Carnap (the early Carnap, at least) adopted a very liberal *Principle of Tolerance*, which said (roughly) that the deductive logician is not in the business of telling people which deductive-logical frameworks (or logical languages) they should adopt. For Carnap, the question of which deductive logical framework one should adopt is a *pragmatic* and *not* a *logical* question – a question which is "external" to logic proper.

On this view, the deductive logician's job is simply to tell us what the deductive-logical relations are between statements *within* specified logical frameworks. For Carnap, this *is* a logical question, which is "internal" to logic proper.

I suggest that we make an analogous sort of move in the case of *in*ductive logic. That is, I submit that the inductive logician is not in the business of telling people which inductive-logical frameworks (i.e., which *probability models*) they should adopt or use. *That*, it seems to me, is a *pragmatic* (or perhaps *epistemological*) question, but *not* a logical one.

The inductive logician's job (I submit) is simply to tell us what the inductive-logical relations (*i.e*., the confirmation relations) are between statements *within* (*arbitrary!*) specified probability models.

NOTE: This is just an *analogy*. We can't push it *too* hard. I am *not* endorsing *logical pluralism* here. I am *not* suggesting that there is a plurality of acceptable *measures* of confirmation. On the contrary, I will argue below that there is *one true measure* of confirmation, and it is *the* job of the inductive logician to tell is what that measure is. What I'm arguing for is tolerance with respect to choices of *probability models*, *not* choices of measures of confirmation.

Carnap's inductive-logical tolerance was similar. He described various systems of inductive logic, and he thought that choices among them were pragmatic (and not purely a matter of logic). And, Carnap's inductive tolerance (like mine) included *only* the probability *models* and not the fundamental confirmation measure itself (which he took to be just the conditional probability function). However, in the ends, Carnap was not tolerant *enough*, because the class of probability models Carnap was able to emulate with his "logical constructions" was not rich enough to cover many of the models which appear to be essential to inductive and statistical inference (even by Carnap's own lights). There seem to be two reasons for this. First, Carnap believed that in order for judgments of *confirmation* to be logical, the underlying probability function in terms of which the confirmation relation is defined must *also* be "logical" (hence, *analytic* in Carnap's technical sense). And, second, Carnap's *formalistic* views about *logic* and *analyticity* placed very strong *syntactical* constraints on his *accounts* of "logical probability". There is much more that can (and should) be said about this historical tension between Carnap's views about the indispensability and analyticity of "logical probability," and his desire for enough flexibility and tolerance in his systems of inductive logic to meet the demands of our inductive and statistical inferential practices. But, in the interest of time, I must now move on to discuss the details of my own "Neo-Carnapian" proposal.

Slide 13:

Here's my proposal: Confirmation is a three-place relation between premises, conclusion, and *a probability model M* (that is, an algebra of propositions – *with real numbers assigned to them* – where this assignment satisfies your favorite axiomatization of probability). This proposal has several advantages, including the following:

• It cleanly separates the *logical confirmation* relation from *epistemic* concepts like "a priori" credence, which often crop-up in so-called "logical" accounts of probability — *via* the Principle of Indifference. [In particular, Carnap's perplexing remarks on the principle of indifference (to which I alluded above) show just how difficult it can be to keep these logical and epistemic concepts straight.]

• It guarantees the *transparent* satisfaction of all three historical desiderata (especially, the "logical" desideratum D3) — *without* requiring a "logical interpretation" of probability.

This approach is "Neo-Carnapian" since it is a *logical* conception of confirmation (in Carnap's sense), but the confirmation function is 3-place (not 2-place), and it is not constructed syntactically out of the logical languages that underlie our formalized descriptions of the probability models.

Note that there are *two* senses in which my account leads to a "logical" confirmation function. First, my account is logical in the *intuitive* sense Carnap discusses in LFP: on my account, degrees of confirmation are determined *mathematically* by the specification of the probability model. Once the model is specified, the degrees of confirmation are just *mathematical properties* of the given model. This seems as objective and non-contingent as one could want. Of course, this approach may require us to broaden our "logical ontology". I am suggesting that we should expropriate the mathematical theory of probability into our inductive logical framework. This is probably not a move all logicians will welcome with open arms. However, philosophers of mathematics with "logicist" leanings – like Carnap himself – should not be too alarmed by this stratagem (especially since almost all useful parts of probability theory can be expressed in very weak fragments of second order logic). The implications of this move for the philosophy of logic and mathematics merit more careful scrutiny. In any case, I don't see how we can *avoid* this move, if what we want is a probabilistic, quantitative ***generalization*** of deductive logic. It seems to me that if we want a more general, probabilistic, logical relation to come *out* of our framework, then we will need to put a sufficient amount of probabilistic structure and theory *into* our framework. Once we have seen the details of my proposal laid out, below, it will become clearer that there is a *second* sense in which my account can be considered "logical" — it provides a proper quantitative *generalization* of *deductive entailment* (*i.e.,* it satisfies desideratum D1).

Next, I will propose a fourth and final material desideratum for the three-place confirmation function c(C, P, M). This will be a *relevance* desideratum, borrowed (in a formal sense) from contemporary Bayesian confirmation theory. This final desideratum will lead to a precise and more-or-less complete rendition of my proposal for inductive logic.

Slide 14:

Consider the following argument:

> (P) Dennis Rodman has been taking birth control pills for the past year.
> Therefore, (C) Dennis Rodman is not pregnant.

Intuitively, the conditional probability of C, given P is quite high (assuming a probability model M which reflects our knowledge that Dennis Rodman is a male, *and* our knowledge of human biology, etc.). It is highly probable that Dennis Rodman is not pregnant, given that he has been taking the pill for the past year. So, according to (TRV), this is a *strong* inductive argument (in M).

This just seems wrong. Intuitively, P is *irrelevant* to C (in M). That is, Rodman's taking the pill does not affect the probability that he is pregnant one way or the other. In this case, it appears that the conditional probability of the conclusion given the premises is high *solely because* the unconditional probability of the conclusion is high in the first place, and *not* because of any confirmation or evidential *relation* between the premises and the conclusion.

This is exactly the same sort of criticism Skyrms levels against the (NIL) proposal. And, it is just as compelling here. This suggests a fourth and final desideratum, which has already been widely adopted in contemporary Bayesian circles:

(D4) The confirmation function should be a *probabilistic relevance measure*.

That is, the confirmation function should be *positive* if P raises the probability of C (in M), negative if P lowers the probability of C (in M) and zero if P has no effect on the probability of C (in M). This immediately *rules-out* (TRV) as an inadequate account of confirmation.

Slide 15:

We can summarize all four of our desiderata for logical confirmation as follows:

(D) … briefly explain its meaning … by completing diagram from D1 on board …

Of the dozens and dozens of relevance measures of confirmation or evidential support that have been proposed and defended in the literature, *only one* (up to ordinal equivalence) satisfies all four of our desiderata, and that is the log-likelihood-ratio measure *l*, which we saw above.

This provides a very elegant, inductive-logical solution to the (historical) problem of measure-sensitivity in Bayesian confirmation theory. Just a few simple and intuitive inductive-logical desiderata pick-out a unique confirmation measure, among all of the historical proposals. This seems to me a better resolution of the measure sensitivity problem than the few others I have seen. Existing resolutions typically require rather complicated mathematical assumptions that are not grounded in any intuitive characteristics of the explicanda: confirmation or evidential support.

Moreover, our approach provides a general, logical foundation for both Bayesian and non-Bayesian probabilistic accounts of evidential support — *via* likelihood ratios. I will conclude with a brief discussion concerning the generality and flexibility of the present framework.

Slide 16:

Personalistic Bayesians often make claims like ``E1 is better evidence for H1 than E2 is for H2''. For instance, they often say that the observation of a black raven is better evidence for "all ravens are black" than the observation of a red herring is. Such Bayesian claims are mere *biographical remarks* concerning relations among some rational agent *S*'s degrees of belief at some time *t*. As such, these comments have little *epistemological significance*, since the only constraint on a rational *Bayesian* agent's degrees of belief is that they satisfy the axioms of probability theory — *any* probability function is as "rational" as any other, from a personalistic Bayesian point of view [DeFinetti, for example, makes this very clear in his writings.].

Non-Bayesians (or more objective Bayesians) in statistics use *objective* statistical models (*e.g.*, models *of* causal-statistical regularities in experimental set-ups). Claims about *statistical* evidential support are thus meant to be *objective* and not merely personalistic.

In either case, our framework provides a *logical* foundation for cashing out these claims as expressions of *likelihood-ratio* comparisons in the salient models — rather than – say – probability comparisons, or comparisons involving other *relevance* measures. That confirmation *within* probability models should be gauged by likelihood-ratios (rather than any of the other proposed measures) does not depend on what these models are models *of* or how they are *applied* — this is just the *nature* of quantitative logical support.

In general, you can plug *any* probability model you want into my framework (even Carnapian "logical" ones!), and what you'll get out are the logical relations of support between the propositions (or statements) in that model. Whether these relations will be of any epistemological (or, more generally, *philosophical*) significance will not be something that inductive logic *qua logic* can decide. To answer *that* question, we'll need to do some additional philosophical analysis. In particular, we'll need to delve more deeply into the epistemology and metaphysics of the *probability models* in question.

… **water**…

--- SHOW OVERVIEW AGAIN --- Now for some concluding remarks ---

…show overview slide again… Click "-models" on the bottom of the present slide…

(III)

Slide 17:

Question: How is our inductive logic *relevant* or *useful* from an epistemological point of view? How does our logical concept of *confirmation* relate to salient epistemic concepts like evidential support or justified degree of belief?

Analogous question: How does *deductive entailment* relate to salient *epistemic* concepts like conclusive evidence or justified belief or knowledge?

Answer: there are *some* connections between deductive entailment and conclusive evidence and/or justified belief; but, they are pretty weak. For instance,

(i) If *S* knows that *P*, and *S* knows that *P* entails *C*, then *S* is (*ceteris paribus*) justified in believing that *P* constitutes *conclusive evidence* that C.

And, *perhaps*:

(*ii*) If *S* knows that *P*, and *S* knows *P* entails *C*, and *S* believes *C on the basis of P*, then *S* is (*ceteris paribus*) justified in believing *C*.

The same *kind* of thing can be said about confirmation and various salient *quantitative* doxastic states. But, the connections here are even more tenuous…

At this point – instead of going through **Slide 18 –** go to the *diagrams* instead, and use *them* to explain the deductive *vs* inductive cases:

Slides 19 & 20 – The Diagrams
…show overview slide again… Click "-models" on the bottom of the present slide…

Conclusion: To sum up: I began by describing a serious problem in the foundations of Bayesian confirmation theory --- the problem of measure sensitivity. Then, I traced out some historical and conceptual developments in inductive logic. This lead to a new, "Neo-Carnapian" (and *relational*) account of inductive logic, which furnished an elegant solution to the Bayesian measure-sensitivity problem.

Our account diverged from Carnap's in two important ways: First, we rejected "The Received View" that the confirmation relation is to be identified with the conditional probability relation – a relation which is insensitive to the *probabilistic relevance* of the premises to the conclusion. Instead, we took confirmation to provide a logical foundation for *evidential support* (*not* conditional probability), and we discovered that likelihood-ratios are uniquely suited, among all historical proposals, to provide this logical foundation. Second, we rejected Carnap's assumption that if *confirmation* is to be a logical relation, then it must be defined in terms of a "logical" probability function. This assumption, together with Carnap's formalistic views on *analyticity and logic* prevented him from "tolerating" the use of *arbitrary* probability models in his inductive logic, and forced him to use only those probability models that can be constructed out of the syntax of certain formal languages in certain ways. Nonetheless, in the *spirit* of Carnap's account, we arrived at an objective and non-contingent (albeit, *relational*) inductive logic, which is flexible enough to undergird the application of probability models to the interpretation of evidence in a wide variety of both Bayesian and non-Bayesian settings. THE END