Confirmation Theory as a Branch of Inductive Logic: Historical and Philosophical Reflections

APA Eastern 2005 Inductive Logic Symposium Script
Branden Fitelson
Decermber 30, 2005

[The slides for this talk are available at `http://fitelson.org/apa2005_handout.pdf`.]

Slide #0 (Title). Before I get underway, I'd like to quickly thank a few people. First, Jonathan Vogel and John MacFarlane for working behind the scenes to make this thing happen. And, of course, David Christensen for chairing, and Patrick Maher and Jim Joyce for participating. I especially want to thank Patrick for his terrific feedback on my work this term, which has helped me to get much clearer on my project. Before we get started, does everyone have a handout? The handout contains all the slides I will be going through. That's almost everything I'm going to say. The script from which I am reading today will only occasionally embellish what's written on the slides (like right now, for instance). OK, now onto today's agenda. ▶

Slide #1 (Overview). I'll begin with some informal stage-setting, and then I'll slide into a brief historical survey of confirmation theory. We'll start with Hempel's theory of confirmation, which makes use of deductive relations, then we'll move on to Carnap who uses logical probabilities, and finally to Bayesians who work with subjective (or epistemic) probabilities. After this historical setup, I'll sketch my own current take on confirmation theory, which today I'm giving the catchy title "None of the Above" (I'll probably need something snappier for my book title). Anyhow, as I will explain, my current views have been significantly influenced by all of these historical accounts. OK, now for some informal stage-setting. ▶

Slide #2 (Stage Setting). Traditionally, confirmation theory has aimed to provide accounts (usually precise, formal accounts) of various sorts of relations of "inductive support" between propositions. There are three basic kinds of support relations ▶ [explain on screen]. ▶ My focus today will be on the quantitative and the ✍ #1
qualitative kinds. I'll say just a few brief things about comparative relations. [I actually think that comparative relations are the most interesting in the end, but it would take us too far afield to discuss them in any detail today.] ▶ We will see both logical and epistemic relations of "support". Indeed, the interplay between logic and epistemology will be my main focus. ▶ When I talk about inductive logic, I will have in mind a theory of "argument strength" or "argument goodness" which in some sense *generalizes* deductive logic. ▶ I would like to think of confirmation theory as a branch of inductive logic. ▶ But, it's controversial as to whether there is a logical notion of (non-deductive) argument strength or whether inductive support should be explicated as a logical concept. My trace of the historical development of confirmation theory aims to shed some light on *why* this is so controversial (and so tricky). ▶ There are two main interrelated themes in my talk today: the role and logical status of probability in inductive logic (confirmation theory), and the relationship between inductive logic (confirmation theory) and epistemology. The first part of the talk will be about inductive logic proper (what you might call theoretical inductive logic) and the second part will be about the *application* of inductive logic (i.e., the relationship between inductive logic and epistemology). OK, so much for the stage-setting. Now, we're ready for our historical survey, which begins with Hempel. ▶

Slide #3 (Hempel #1). Taking Nicod's inchoate remarks about instantial confirmation as his point of departure, ▶ Hempel formulated a rather sophisticated formal theory of confirmation as a logical relation between sentences in first-order languages. Hempel's qualitative account has the following form ▶ E confirms H if E deductively entails Z, where Z is a sentence constructed syntactically from E and H in a certain way. The details of this construction will not matter for our purposes. ▶ Hempel also gave quantitative and comparative accounts of confirmation, but I won't be discussing those today. ▶ Hempel's qualitative confirmation relation has various properties, including the following three, which will play a central role in my talk [explain on screen]. ▶ The entailment condition (EC) is rather uncontroversial. There are few ✍ #2
contemporary confirmation theorists who would deny it (except perhaps in a very small class of cases). Basically, (EC) says that entailment is a "best case" or a "limiting case" of confirmation, and this seems to be a fundamental desideratum for theories of confirmation (or inductive logic). The special consequence condition and monotonicity, on the other hand, are far more controversial. We will hear more about (SCC) later when we discuss Carnap. Monotonicity is perhaps the most important property in today's talk. ▶ Interestingly, Hempel's own intuitions about inductive support ran counter to monotonicity. In his discussion of the raven paradox, Hempel cautions us not to conflate the following two claims concerning inductive support. ▶

Slide #4 (Hempel #2). (PC): that an object a is a non-black non-raven inductively supports that all ravens are black, relative to tautological corpus (or to put this in more epistemic terms, the observation that an object a is a non-black non-raven *evidentially* supports that all ravens are black, in contexts where nothing about the object a is antecedently known). ▶ (PC*) that an object a is a non-black non-raven inductively supports that all ravens are black, relative to a corpus which includes the information that a is a non-raven (or to put this in more epistemic terms, the observation that a is a non-black non-raven *evidentially* supports that all ravens are black, in contexts where it is antecedently known that a is a non-raven). Hempel's intuition was that (PC) is true, but (PC*) is false. [The idea here seems to be that observing *antecedently known* non-ravens doesn't tell you anything about the color of non-ravens. But, if you don't know anything about an object antecedently, and *then* you see that it is a non-black non-raven, then this observation rules that object out as a possible counterexample to the law, which makes the observation relevant to the truth of the law.] Hempel thought that this explains why people tend to think that (PC) is false — because they tend to conflate (PC) with the false claim (PC*). ▶ Unfortunately, Hempel's confirmation relation is monotonic, and there is no distinction in classical logic between "given that" and "and". That is, there is no distinction between the following two claims ▶ : $\boldsymbol{E \& K}$ entails Z, relative to (or, given) *tautological* corpus — *versus*, $\boldsymbol{E}$ entails Z, relative to (or, given) corpus $\boldsymbol{K}$. ▶ As a result, Hempel's theory makes it impossible for (PC) to be true while (PC*) is false — thus contradicting his own intuitions. ▶ I think Hempel's intuitions are plausible here. There is an intuitive difference between (PC) and (PC*) along the lines Hempel suggests. But, one needs a non-monotonic conception of confirmation to make sense of this. ▶ Moreover, Hempel's example suggests that the support relation may have more than two relata. There also seems to be something like a set of background conditions (or a background context) involved in relations of inductive support. I'll come back to that later. ▶ In any case, it seems that Hempel's account of confirmation is inadequate, even by Hempel's own intuitive lights. Continuing along our historical trajectory, ▶ Carnap (and others) used probabilities rather than entailments to provide accounts of confirmation. Because probabilities are non-monotonic (and sensitive to changes in background information), they are able to capture the sorts of features that arise in Hempel's discussion of the raven paradox. This brings us to Carnap's account of confirmation. ▶

Slide #5 (Carnap #1). In logical foundations of probability (LFP), Carnap aims to explicate various notions of inductive support. He clarifies his explicanda in various ways, including the following clarification of the qualitative concept of inductive support, which I am calling STAR. ▶ E gives some (positive) evidence for H. [Note Carnap's use of epistemic-sounding language here. Keep that in the back of your mind as we proceed.] ▶ In the first edition of LFP, Carnap explicates "the degree to which E confirms H" as "the conditional probability of H, given E". Given the framework Carnap sets-up at the beginning of the book, this leads naturally to the following three precise confirmation concepts ▶ [explain on screen]. ▶ On reflection, this qualitative ✍ #3 concept does not sound very much like STAR. The conditional probability of H given E can be very high even if E *lowers* the probability of H, i.e., even if E is *negatively relevant* to H (*e.g.*, diagnostic tests for rare diseases). Such cases don't sound like cases of "positive evidence" to me. I'll return to this dissonance later. ▶ Like Hempel, Carnap wanted an explication of confirmation as a logical relation between sentences in first-order languages. ▶ For Carnap, this meant that the probabilities themselves must be logical, i.e., that probability claims in inductive logic must be analytic and knowable a priori (and that they should generalize entailment). For Carnap, these probabilities were a kind of "partial entailment relation" between sentences. ▶ This leads naturally to the Carnapian program of trying to provide a logical explication of conditional probability itself. ▶ I am skeptical about that project (i.e., I am skeptical about the existence of logical probabilities). But, since that issue is orthogonal to my talk today, I will let Jim and Patrick talk about it. ▶

Slide #6 (Carnap #2). Instead, I would like to focus on two other questions that arise from (my reading of) Carnap's writings on confirmation. I'm calling the first question the "logicality of Pr" question: ▶

> **Logicality of Pr**. Must conditional probability *itself* be logical if (probabilistic) *confirmation* claims are to be logical? Or, to put this another way, must the confirmation theorist *qua inductive-logician* tell us *which probability assignment is salient* for the assessment of the strength of a given argument?

I'm calling the second question the "epistemology and inductive logic" question. ▶

> **Epistemology and Inductive Logic**. Carnap presumed (as do I) that there are *some* connections between deductive logic and epistemology. Are there analogous connections between confirmation theory (inductive logic) and epistemology?

▶ Carnap thought the answer to the first question was “YES”. I disagree with this. Not only because I am skeptical about the *existence* of logical probabilities, but also for independent reasons having to do with their epistemological cash-value, which I will explain below. In the end, I will sketch an alternative approach that doesn’t *require* the probabilities used in confirmation theory to be logical. ▶ Carnap also thought the answer to the second question was “YES”. With this, I wholeheartedly agree. I think it is crucial that inductive logic be applicable to actual knowledge situations — that it not be completely detached from epistemology. I will try to gesture toward something positive here, but there will be some serious challenges along the way. Before we get to that, we need to return to the dissonance we saw in Carnap’s initial qualitative conception of confirmation. ▶ In chapter 6 of LFP, Carnap gives a counterexample to Hempel’s special consequence condition, which presupposes a different qualitative conception of confirmation — which is more ($\star$)-like. ▶

Slide #7 (Carnap #3). This other approach to qualitative confirmation says that E confirms H iff E and H are positively probabilistically correlated (under a suitable logical probability function, of course). ▶ This probabilistic relevance concept is one which *violates* the special consequence condition. On the other hand, Carnap’s previous probability threshold conception *entails* the special consequence condition. Thus, these concepts really are drastically different, and in ways that have historically important consequences for confirmation theory. ▶ In the 2nd edition of LFP (largely in response to Popper who brought some of these things to his attention) Carnap includes a preface which acknowledges an “ambiguity” in the discussion of the qualitative conception of confirmation that appears in the first edition. In that preface, Carnap also says that the qualitative *relevance* concept is the more interesting of the two. I agree (as do most contemporary probabilistic confirmation theorists). Here’s what Carnap says about the two concepts ▶ [explain on screen]. ▶ **U**nfortunately, the first edition is almost entirely about firmness, and ▶ the second edition does not rigorously ✍ #4 develop or defend a precise theory of confirmation as increase in firmness. ▶ My project today can be seen as an attempt to extend some of Carnap’s ideas about firmness to increase in firmness. As it turns out, this is a non-trivial task, especially when it comes to Carnap’s ideas about the *application* of confirmation theory to knowledge situations (i.e., to relating confirmation as increase in firmness and epistemology). ▶

Slide #8 (Carnap #4). Before moving on to epistemological issues, let’s think a bit more about the *theory* of confirmation as increase in firmness. There are *many* functions $\mathfrak{f}$ that satisfy the basic sensitivity to relevance constraint, which I am expressing here as the desideratum $\mathcal{R}$ ▶ [explain on screen]. **I**ndeed, ✍ #5 the many functions that satisfy $\mathcal{R}$ lead to drastically (and surprisingly) different comparative theories of confirmation as increase in firmness. For this reason, we will need to narrow down the field of candidate functions with further inductive-logical desiderata. ▶ Another inductive-logical desideratum that immediately comes to mind is the requirement that confirmation measures should generalize entailment (they should be measures of partial entailment, in some sense). We can state this desideratum, which I will call $\mathcal{D}$, as follows ▶ [explain on screen]. ▶ **K**emeny and Oppenheim (who were writing before the 2nd edition of ✍ #6 LFP was published) used this inductive-logical desideratum (and others) to select the following explicatum for (what Carnap later called) degree of confirmation as increase in firmness. ▶ I have written F as a function of the *likelihoods* [$\Pr(E \mid H)$ and $\Pr(E \mid {\sim}H)$], rather than the conditional and unconditional probabilities of H, because this makes F easier to understand. ▶ But, F *can* be so expressed. Moreover, F satisfies $\mathcal{R}$, $\mathcal{D}$, and various other *comparative* inductive-logical desiderata, which I have discussed elsewhere [explain on screen]. **I** think F (or something very close to it) is the best explicatum for degree of confirmation as increase ✍ #7 in firmness (although, my desideratum/explicatum argument in favor of F would be rather different than Kemeny and Oppenheim’s). ▶ Anyhow, once we have F in hand, we can then use F to define comparative and qualitative notions of confirmation as increase in firmness in the standard way, which is what I will do. That (for the most part) completes the theoretical inductive logic part of this talk. Now, we’re going to move on to the *applications* of inductive logic, that is, to the epistemological part of the talk. ▶

Slide #9 (Carnap #5). I agree with K&O that F has the proper *form* for a confirmation as increase in firmness measure (that is, F gives us the right kind of story about *how confirmation relations depend on probability assignments*), but since K&O follow Carnap in requiring their probabilities to be logical, their approach is too Carnapian for me. As we’ll see momentarily, my reasons for moving away from the Carnapian program are not merely because of my skepticism about the existence of logical probabilities, but also because of independent problems involving the application of a Carnapian (logical) conception of confirmation as increase in firmness to actual knowledge situations. ▶ This brings us to Carnap’s views on the relationship between logic and epistemology. ▶ In the case of deductive logic, Carnap endorsed the following epistemic bridge principle connecting entailment and knowledge (this is just a naive closure principle, and I'm not endorsing it, I'm just reporting it) ▶ :

> If E is known by person X at time t, then H is likewise known by X at t [provided that $E \vdash H$ is also known by X at t].

▶ In the case of inductive logic, Carnap endorsed the following (analogous) epistemic bridge principle connecting confirmation as firmness and credence (again, not endorsing, just reporting). ▶

> If E *and nothing else* is known by person X at time t, then the degree of belief justified by the knowledge of X at t is r [provided that $\mathfrak{c}_f(H, E) = r$ is also known by X at t].

▶ The telling "and nothing else" clause is required *because* confirmation as firmness is non-monotonic. Generally speaking, if we strengthen the evidence, we can *undermine* probabilistic relations of confirmation. Thus, whenever we determine the degree of confirmation in an actual knowledge situation (i.e., a context), we must be sure to take into account our total evidence (in that context). This "requirement of total evidence" will come back to haunt us when we try to formulate an analogous bridge principle for confirmation as increase in firmness. ▶ Before we do that, let's take a quick glance at Bayesian confirmation theory. ▶

Slide #10 (Bayesianism). Most contemporary Bayesians are skeptical about the existence of logical probabilities. I'm inclined to agree, but I'll let Jim and Patrick discuss that, since I have other Carnapian fish to fry. ▶ As a result, most modern Bayesians have simply given up on the traditional project of confirmation as a branch of inductive logic. ▶ Instead, they take aim at an explicitly epistemic notion of inductive support. Their qualitative confirmation relation is defined as follows: ▶

> E confirms H *for agent X at time t* iff E and H are positively correlated under X's credence function at t.

▶ This is *formally* quite similar to Carnap's increase in firmness concept (because it is couched in terms of correlation), but it is subjective and epistemic rather than logical and objective. ▶ One nice thing about this approach is that it forges a very tight connection between confirmation theory and epistemology. One might even call this taking "epistemology on the cheap". [you might see me as taking "logic on the cheap" later.] ▶ One crucial presupposition that the Bayesians share with Carnap is that all inductive support relations supervene on one kind of probability. They just disagree about which kind of probability that is (Carnap thinks it's logical probability, and Bayesians think it's epistemic probability). I will reject this supervenience assumption *altogether*, since I think inductive logic should be neutral on the kind of probability that is appropriate in any given application of inductive logic. ▶ [comp/quant Bayesian issue not on the agenda].

RE-CAP. So far, we have seen three theories of confirmation. Hempel's theory was based on deductive entailment, which caused his confirmation relation to be monotonic, and that seemed inadequate (even by his own lights). Carnap's theory was probabilistic, which overcame this and other shortcomings of Hempel's approach, but which also added a couple of new complications of its own. The first complication is that Carnap used *logical* probabilities, and it is unclear whether such things exist. The second complication is that because Carnap's relations are non-monotonic, he needed a new principle to govern applications of confirmation theory — the requirement of total evidence, which will play a starring role in our sequel. Finally, the Bayesian theory gives up on a logical account, in favor of an explicitly epistemic account, based on credential correlation. This historical survey sets the stage for my proposal. [water, maintain speed] ▶

Slide #11 (Me #1). Before sketching my account, I will describe the central problem that has motivated my move away from Carnapian and Bayesian approaches to confirmation theory. This has to do with the application of inductive logic to actual knowledge situations. In deductive logic, we make a distinction between validity and soundness, and a related distinction between the theory of deductive logic and its application to making inferences in knowledge situations (or contexts). I will frame my discussion in terms of "all things considered" assessments of the "goodness" of arguments in contexts (or knowledge situations). Such assessments will generally have both logical and non-logical components. ▶ In the deductive case, the ▶ logical component is the determination of whether the argument is valid, and ▶ the non-logical component is the determination of whether the premise of the argument E is true (I say "generally" non-logical here, because in some contexts E may *itself* be a logical claim — I'll bracket such cases). ▶ How do we extend this to include the inductive case? ▶ Presumably, we still need to determine whether E is true in the context. ▶ What *else* do we need to do? I guess Carnap would have said that the other thing we need to do is to determine the degree of confirmation that E provides for H. Traditionally, this would have been seen as determining whether a certain kind of partial entailment relation holds between E and H. As always, this must be done in accordance with the requirement of total evidence, which says (quote from Carnap): ▶

> **The Requirement of Total Evidence**. In the application of inductive logic to a given knowledge situation, the total evidence available must be taken as a basis for determining the degree of confirmation.

This immediately raises a problem for the application of confirmation as increase in firmness. If we know that the premise of the argument is true, then this undermines our ability to determine that E confirms H in the increase in firmness sense. Here's why. ▶ Let K_C be the knowledge we have in a given context C in which we are assessing an argument from E to H. And, let's assume that K_C already contains the information that E, because we already know in the context that the premise E is true. In such a case, conditionalizing *any* probability function (logical or otherwise) on our total evidence K_C (as the requirement of total evidence requires) has the effect of *canceling* (or screening-off) any correlation between E and H. ▶

Slide #12 (Me #2). Therefore, if we know in a context that E is true, we *cannot* — on a Carnapian approach — determine *in that very context* that E increases the firmness of H. ▶ This is problematic. The non-logical component of our all things considered assessment of the argument's "goodness" has *interfered* with (or undermined) its (Carnapian) logical component! ▶ Interestingly, this undermining effect does not plague the application of *firmness* confirmation, since it doesn't prevent the logical probability $\Pr(H \mid E\&K_C) = \Pr(H \mid K_C)$ from being greater than *some fixed threshold value* (say $\frac{1}{2}$). This is probably why Carnap does not discusses this problem in LFP. You see, Carnap never re-wrote any of the main text of LFP in light of the above "ambiguity" between firmness and increase in firmness. Instead, he just added the brief preface, which does not trace out all of the consequences of thinking about confirmation in the increase in firmness sense. In particular, he did not discuss in detail how the requirement of total evidence is to be applied when we are working with confirmation as increase in firmness. If Carnap *had* discussed this issue in LFP, I suspect the history of confirmation theory would have been significantly different. More than 15 years after the second edition of LFP was published, Clark Glymour discovered a ▶ similar problem which infects the Bayesian approach to confirmation. Glymour's problem is now infamously known as "the problem of old evidence." [Bayesians have tried to cure their strain of this disease by using various kinds of philosophical and mathematical tricks. But, I am now convinced that none of these tricks can solve the underlying epistemological problem. I'd be happy to explain (if time permits later in the session) why I think none of the existing Bayesian responses can do the trick.] The Carnapian and Bayesian problems share a formal component: that *no probability-assignment which assigns probability one to E can reflect a correlation between E and H*. Thus, any correlation-based account of confirmation that respects the requirement of total evidence by *conditionalizing* the salient probability function on E in contexts where E is known will face this problem. ▶ So, how in the world can anyone who knows in a context that E is true determine in that very context that an inductive argument from E to H is "all things considered good" in *some* "increase in firmness sense"? ▶ *That* is one of the central challenges motivating my project. I'll look at a concrete example of this phenomenon below. ▶ But, first, I will outline the basic ideas behind my approach. ▶

Slide #13 (Me #3). What we need is a way to be able to know (in a context C) that E is true *and at the same time* determine (in C) that — in *some* sense that is *epistemically salient* — E makes a positive difference to the probability of H. Two steps. ▶ **Step 1**. Confirmation as increase in firmness is a *three*-place relation [$F_{\Pr}(H,E)$] between E, H, *and a probability assignment* Pr, where Pr is a *non-logical parameter*, which gets fixed *contextually* in epistemic *applications* of confirmation as increase in firmness. ▶ Note: on the present account, *confirmation* claims $F_{\Pr}(H,E) = r$ are *logical* and *objective*. That is, they are *analytic*, knowable *a priori*, and they *quantitatively generalize entailment and refutation* (and this is all Carnap meant by "logical" in LFP). Moreover, our confirmation claims have all of these properties without requiring the probability assignment itself to be "logical". This is because, on the present account, confirmation is always *relative to an assignment of probabilities*. I concede that (in some sense) this is getting "logic on the cheap." But, this does not trivialize the task of the confirmation theorist *qua inductive logician*, since the problem of determining the proper form of the confirmation function remains (and it is non-trivial and important). I'll return to that issue below. ▶ **Step 2**. When assessing in a context the "all things considered goodness" of an argument from E to H (*i.e.*, when assessing whether *E evidentially supports H* in a context), one must now determine *three* things in the context in question. And, I submit that two of these determinations will (in general) be *non*-logical in nature ▶ [explain on screen]. ▶ Now, let's look at a concrete "problematic context." ▶ ✍ #8

Slide #14 (Me #4). When we are assessing an invalid argument in a context, there may be some *prima facie* salient probability assignments (in the context) under which E and H are correlated, and yet others under which they are not correlated. This kind of non-unanimity cannot happen with valid arguments, since all probability assignments will agree on whether E and H are correlated in cases where E entails H. ▶ Here's a concrete example, which also illustrates the "old evidence" phenomenon I described above. Keep in mind that this is just an example. It's not meant to be canonical. It's just meant to be a context that is problematic for the Carnapian and Bayesian approaches, and that we (at least I) have strong intuitions about. ▶ [explain on screen] ▶ It seems clear to me that John's assessment in this context should be that E provides some ✍ #9 positive evidence for H. This suggests that Pr (and not Pr′) is the salient assignment in this context. ▶

Slide #15 (Me #5). This also suggests (to me) the following contextual bridge principle. Before I state the principle, I want to emphasize that I am just stating the Carnap-style rendition of it. There are bound to be various things that would need to be modified to make this principle more plausible (as was also the case with Carnap's other bridge principles). Anyhow, here's what the naïve, Carnap-style principle looks like: ▶

> If X knows that E is true in C and that Pr is the (most?) evidentially salient Pr-assignment wrt E and H in C, then X knows that E provides some (positive) evidence for H in C [provided that X also knows $F_{\text{Pr}}(H, E) > 0$].

▶ In this **qualitative** case, *any relevance measure* will do in place of F, since they all agree on judgments of positive and negative relevance. ▶ However, if we extend this to **comparative** relations such as "E favors H_1 over H_2 (in context C)", things get much more interesting from an inductive-logical point of view, since those assessments will (in general) *depend on* the *form* of our chosen confirmation measure F. As I have argued elsewhere recently, some important debates in the philosophy of statistics concerning the proper interpretation of statistical evidence can be seen (and clarified) as debates about the proper form of the confirmation as increase in firmness function (here, I have in mind debates among statisticians concerning the so-called "law of likelihood"). That topic would take us too far afield today. But, I mention it to remind you that the determination of the proper form of the confirmation function remains a non-trivial project for confirmation theorists *qua inductive logicians* to complete. ▶ Now, for two important remarks about my bridge principle. First, I'm *not* claiming that *all* (rational) evidential judgments (in *all* contexts) can be explicated or grounded in this way. It may be that some are (inherently) probabilistically inexplicable (either because there is no salient probability assignment available in the context, or for some deeper reason having to do with the nature of evidential relations). ▶ Second, in some contexts, the most evidentially salient probability assignment *does* assign probability one to E. For instance, if the pregnancy test is *known* to *always* give positive results in the context, then the degree of support *should* be zero, because (intuitively) there is no evidential support in such contexts! ▶ The $64 question is: What *determines evidential salience*? Of course, I have no *theory* of this (as that would require taking a stand on many deep epistemological issues — and this is just a session on inductive logic after all!). But, I will say a few things. ▶

Slide #16 (Me #6). The first thing I want to reiterate is that just because a probability assignment is evidentially salient in a context, this does not imply that it would be rational to adopt it as a credence function in that context. ▶ This is especially true in cases of "old evidence" where doing so can be a very bad idea. ▶ Second Thing. Accuracy is something that *can* contribute to evidential salience. If we've decided that it's the reliability of the test apparatus that matters (in some context), then we'll want to avoid models of its error characteristics that are known (in the context in question) to be wildly inaccurate. ▶ Third, which assignments one takes to be salient in which contexts will also depend on one's epistemological commitments concerning the nature of evidential relations. ▶ For instance, externalists may tend to see objective probability assignments as more evidentially salient than subjective ones, whereas internalists may want to require that the agent have a certain kind of special (or privileged) access to the probability assignment in question in the context in question [for example, I think some recent debates about testimonial evidence can be reconstructed as debates about which probability assignments are evidentially salient in various kinds of testimonial contexts]. ▶ Finally, the inductive logician *qua logician* needn't resolve such issues. Their job is to determine the proper form of the confirmation function, not to provide an (epistemic) theory of the evidential salience of probability assignments. OK, that's the end of my story today. Thank you for listening.