**Notes on KAIL Chapter 10**
11/15/06 (B.F.)

This is a very long, technical, and difficult chapter. There is no hope of going through it all in one seminar. Instead, my comments (and Mike's comments) will focus on a few issues that are central to TW's discussion in the chapter. Let's begin by recalling TW's explication of "$E$ is evidence for $H$ for $S$", which we have already discussed in our last meeting:

(EV) $E$ is evidence for $H$ for $S$ iff $E$ is included in in $S$'s body of evidence and $\Pr(H \mid E) > \Pr(H)$.

Last week, we discussed the first component of (EV) — the "inclusion in $S$'s body of evidence" component. This week, we will focus on the second component — the "support as positive probabilistic dependence" component. There are many important questions that must be answered in order to get a handle on the intended sense of this second component. Here are five:

1. What *kind* of probability function is TW's "evidential" probability function Pr?
   - $\Pr(p)$ is clearly *not* $S$'s *actual* degree of belief (credence) in $p$ [$\Pr \neq C_a$]. For one thing, $S$'s *actual* degrees of belief $C_a$ can be *crazy* (even *incoherent*), and we want Pr to ground *non*-crazy claims about evidential support [although, it's not entirely clear *why* we want this from a traditional epistemological point of view, which cares about *full* belief and *knowledge* — see question (3), below, for further discussion of *that* issue]. Moreover, as we mentioned last week, if $S$'s credence in either $H$ or $E$ happens to be extreme (zero or 1), then $E$ and $H$ *cannot be* positively dependent on $C_a$. And, surely, $S$ *can actually have* extreme credences in either $E$ or $H$, for *any* $E/H$. As such, identifying Pr and $C_a$ would basically *trivialize* TW's "support" relation. Finally, TW wants his "support" relation to be *objective*, and identifying Pr and $C_a$ would make it patently *subjective* in a way TW doesn't want.
   - $\Pr(p)$ is not even to be identified with the *ideally rational* or *justified* degree of belief (for $S$) in $p$ [$\Pr \neq C_j$]. TW says that this "comes closer" to what he has in mind with his "evidential" probability. For one thing, at least $C_j$ is *objective* in ways that $C_a$ is not. For instance, even if $S$ *happens to* have extreme degree of belief in $E$ or $H$, it might be that this is *irrational* or *unjustified* (given everything $S$ knows). Moreover, $C_a$ might even be *incoherent*, whereas (presumably) $C_j$ cannot be. This gives $C_j$ a kind of objectivity (and also probative force) that seems to be required to undergird TW's objective "support" relation. However, TW wants to think of "support" as an objective relation *between E and H* [and *some* "background corpus" $K$ — see question (4) below]. But, if we identify Pr with $C_j$, then we can still get undesirable verdicts about "support", because an agent (albeit an ideally rational one) is "getting between" $E$ and $H$. For instance, there may be cases in which $C_j(E) = 1$, but we still want to say that $E$ supports (some) $H$ in TW's sense. [This is especially true in light of the fact that TW thinks it can be rational for $S$ to assign probability 1 to $E$ at $t$, and then later at some $t' > t$ assign probability less than 1 to $E$. Mike will talk about that set of issues.] So, not only is $C_j$ not to be *identified* with TW's "evidential" probability Pr, $C_j$ and Pr can *disagree on the values they assign* to some $p$'s. That is, TW seems committed to the view that $C_j(p) \neq \Pr(p)$ for some $p$'s. This is not unusual for "inductive" or "logical" probability in the tradition of Keynes, Carnap, and Maher (see Maher's recent paper on inductive probability for a very nice discussion of these and many other related issues). Carnap says $C_j$ and inductive or logical probability $\Pr_l$ are *conceptually distinct*, but intimately related. Specifically, Carnap takes $\Pr_l(H \mid E)$ to be a *logical* relation between $E$ and $H$, whereas $C_j(H \mid E)$ is not a logical relation, but an *epistemic* relation between $E$, $H$, *and* $S$. But, Carnap also endorsed a *bridge principle* that says ideally rational agents whose *total* evidence is $E$ [and who know the value of $\Pr_l(H \mid E)$!] should be such that $C_j(H) = \Pr_l(H \mid E)$. On this type of view (shared by Keynes and Maher), $\Pr_l$ and $C_j$ are conceptually distinct, but $\Pr_l$ *perfectly constrains* $C_j$ in the sense described by the bridge principle. Williamson seems to have a similar view, although he doesn't clearly endorse Carnap's bridge principle [which raises worries about the *relevance* of his Pr to broader epistemic questions about degree of belief, belief, and knowledge — see (2) and (3) below, and Mike's comments. It may *sound like* TW is endorsing a similar bridge principle in his ECOND, but ECOND is only constraining "evidential probabilities *for S*", which are *not* $S$'s *credences*. More on TW's mysterious ECOND later.]
   - The answer *seems* to be that $\Pr(p \mid q) = \Pr_l(p \mid q)$, where $\Pr_l$ is "logical" or "inductive" probability in the Keynes/Carnap/Maher tradition. And, as a result, Williamson's "support" notion is just Carnap's "initial confirmation" relation $\Pr_l(H \mid E) > \Pr_l(H \mid \top)$. There are many problems with the traditional notion of "inductive probability" that TW seems to be working with. For instance: (a) it is not at all clear that such things *exist*, and (b) even if they do exist, it's unclear how we can *know* enough about them. Maher defends the existence of "inductive" probabilities in his recent paper. But, he doesn't say enough about how we can know what their values are. In deductive logic, we postulate the existence of an entailment relation. But, we have a pretty clear sense that at least some instances of this relation actually obtain. For instance, we think that the proposition expressed by "the ball is white" entails the proposition expressed by "the ball is either white or black". Maher thinks that some instances of "inductive" probability relations are *equally* clear. He gives the following example as his main illustration:
     - The probability that a ball is white, given that it is either white or black, is 1/2.

Two points about this. First, it's not equally clear *to me* that this is true (*e.g.*, think about what happens if we partition the colors differently, *etc.*). But, putting this to one side, even if we grant that *some* inductive probability statements of the form $\Pr(p \mid q) = r$ are as clear as any statement asserting an entailment, this hardly suffices to get Maher and TW off the hook. Here's an easier example: $\Pr(p \mid p) = 1$, for any contingent $p$. *These* seem as clear to me as any entailment claim. But, we need a lot more than these toy/trivial examples to do the work TW needs to do (and the work Carnap and Maher want to do as well). The real problem comes in when we try to give equally clear examples involving statements like $\Pr(p \mid \top) = \Pr(p) = r$, where $p$ is a contingent claim. That is, it's the "a priori"/"prior"/"relative to tautological background" inductive probabilities that are the tough ones to motivate. And, these are needed in order to undergird TW's account of "support" (or Carnap's notion of "initial confirmation", relative to the "logical ether"). Presumably, $\Pr(H)$ is supposed to be something like "the probability of $H$ *conditional on no (or tautological) background corpus.*" It is *this* concept that is a total mystery to me, especially if $H$ is something like "I am in the good case". Unfortunately, these are exactly the sort of $H$'s that *matter* here. Williamson says nothing to ground his "prior evidential distribution" $\Pr(\cdot)$, and Maher's examples don't establish clear cases of it. Carnap spent many years trying to *logically construct* $\Pr(\cdot \mid \top)$, without success.

2. What is the connection (if any) between Pr and $S$'s *degrees of belief*? This is related to question (1). We've partially answered this already. As I noted above, TW doesn't say anything as specific as what Carnap says about this [although, see question (3) for something *close*]. But, TW must assume that the agent's credence function (even if it's a perfectly rational one) is sometimes *not* to be used as "evidential". Last week, I argued that since TW accepts $E \subseteq K$, he can *never* conditionalize the evidential probability function on $S$'s total evidence $K$, without undermining $E$'s ability to support anything (for $S$). This still stands. Nothing TW says here responds to that point. It may *appear* that his ECOND is relevant, but it isn't. Bayesians already don't think that all evidence must have prior probability 1. So, the fact that TW allows evidence to have non-extreme prior probability is not a new idea, and it doesn't fix the problem here anyhow. *If* the evidential probability function Pr is *conditionalized* on (anything that entails) $E$, *then* $\Pr(E) = 1$. This is just a fact about *probability* — what agents learn or forget is beside the point. TW's story about "updating", which allows "$S$'s evidential probability" to "evolve" from assigning probability 1 to $E$ to assigning probability less than 1 to $E$ is a *non-sequitur*. On this way of formulating things (which just adds confusion — see *fn.* 1), the problem kicks-in "at the time" when the "agent's evidential probability function" $\Pr_{\alpha}$ *does* assign probability 1 to $E$. At *that* time, $E$ *can't* support $H$ for $S$. *That*'s the problem. That this may not plague a *different* $\Pr_{\alpha'}$ which $S$ may "have later" is *irrelevant*.

   [At this point, we should also ask what (if anything) $\Pr(H \mid E) > \Pr(H)$ has to do with what happens (or should happen) to $S$'s degrees of belief when $S$ *learns* (exactly) $E$. After all, much of the chapter *seems* to be about *learning* (or *updating some* kind of "personal probability"). Mike will address this question in his remarks. Here, I highly recommend Mike's dissertation research, which gives nice formal models (much better models than Williamson's!) of forgetting and other cases that are problematic for traditional Bayesian approaches to learning. Anyhow, Mike will tell us about this stuff.[1]]

3. What is the connection (if any) between $\Pr(H \mid E) > \Pr(H)$ and the epistemic status of $S$'s *beliefs* regarding $H$ and $E$ (and/or their "objective probabilities"). This is related to questions (1) and (2). I discussed this a bit in my handout last week. TW suggests we should "proportion our belief in a proposition to its probability on our evidence". What does that mean? Is this an endorsement of Carnap's bridge principle? It sounds like it *may* be. But, why isn't TW saying something different here? Specifically, why isn't he saying "proportion our *beliefs about probabilities* to our *knowledge of* $\Pr/\Pr_l$"? Isn't this more coherent with the rest of the book? NOTE: these are *radically* different "epistemic rules"!

4. Which propositions are to be treated as "background certainties" by Pr? Which propositions $K$ are such that $\Pr(K) = 1$? This is a crucial question. We know that $\Pr(E)$ and $\Pr(H)$ *cannot* be extreme, *if* $E$ supports $H$. But, it would be nice to have some *independent* grasp on the stuff that *is* included in the set of "background certainties" $K$. If we already need to know independently whether $E$ supports $H$ in order to determine $K$ (hence Pr), then this indicates that "support" may be *philosophically prior* to "evidential probability", which makes TW's probabilistic explication of "support" unhelpful.

5. How much *awareness* or *access* to TW's "support" relation must an agent have in order for $E$ to count as evidence for $H$ *for S*? Let's say "evidential" probabilities *do* exist, and that there is even some story about them that is as compelling as our story about entailments. Is it sufficient that $\Pr(H \mid E) > \Pr(H)$ in order for $E$ to support $H$ ***for S***? Think about the extreme case in which $E$ *entails* $H$. Is *even this* sufficient for $E$ to count as evidence in favor of $H$ ***for S***? Doesn't $S$ need to have *some awareness of* the fact that $E$ entails $H$? What if it is *impossible* for $S$ to know that $E$ supports $H$? Or, what if $S$ believes (falsely, but reasonably) that $\Pr(H \mid E) < \Pr(H)$? Some internalists (Fumerton) want more *cash value*.

[1] There is a ton of confusion in the literature about the relationship between *atemporal* objective probabilities $\Pr(H \mid E)$ [or confirmation relations $\mathfrak{c}(H, E)$], *synchronic* degrees of belief (or credences) $C$ of an agent $S$, and *diachronic* learning relations involving the agent $S$ (which have to do with $S$ evolving from one credence function $C$ to another $C'$ as a result of learning in between). TW has only added to this confusion by inventing yet another type of "personal probability": "$S$'s evidential probabilities in $\alpha$", which are distinct from $S$'s *credences* in $\alpha$ and also from the "prior evidential" Pr that *defines* "support" for TW. This does nothing to aid our understanding of TW's account of "support", nor does it help us understand *learning/updating*.

## Some Comments on *KAIL* Chapter 10

Mike Titelbaum, 11/15/06

Some of the comments below refer to ideas developed in my paper "Unlearning What You Have Learned."

### Section 10.1

- On p. 210, it does not follow from "The probability of a hypothesis on our evidence does not always coincide with the credence which a perfectly rational being with our evidence would have in it" that "We therefore cannot use decision theory as a guide to evidential probability." There is a false assumption here that the only way to do decision theory is to work with the concept of a perfectly rational being.
- I don't understand the argument in the first full paragraph on p. 211. Among other things, a thoughtful Bayesian would not try to develop a theory of probabilities on evidence by analyzing rational beliefs about probabilities on evidence. Facts about evidential probabilities would develop from facts about rational degrees of belief. This will come up again with rule-following in Section 10.3.
- Besides the fact that I'm dubious that one could actually develop an intuitive objective theory of "the intrinsic plausibility of hypotheses prior to investigation," Williamson sets that theory up in a way that makes the task especially difficult. He seems to want to bar us from asking what this theory is a theory *about*. (It's certainly not about any notion we use in ordinary conversation!) This makes it hard to understand why certain moves within the theory are allowed, such as the "familiar bargain" he strikes on p. 212. Moreover, Williamson cuts off connections between intrinsic plausbility and intuitive notions that might help us figure things out about intrinsic plausibilities, such as how it would be rational to respond to learning various pieces of evidence or the degree to which one proposition confirms another. The problem isn't that intrinsic plausibility is "vague," it's that he's rendered it a bizarrely free-floating notion. (This makes his two analogies on p. 211 especially inapt.)
- Notice that Williamson's $P$ mandates a credence for any proposition on any logically consistent evidence set. That is, for any logically possible set of propositions $e$ and any proposition $h$, there exists a unique $P(h \mid e)$. (The "mandating" terminology comes from my "Unlearning.")

### Section 10.2

- If you stop to think about it, there's no reason Williamson should particularly care about a *diachronic* principle like Conditionalization. He's interested in the probability of a proposition on a particular body of evidence. Given his theory, that's $P(h \mid e)$. ECOND doesn't add anything to that: it just says that the evidential probability of $p$ <u>for you</u> is $P(p \mid e_\alpha)$ where $e_\alpha$ is your evidence. (On p. 220 he says that $e_\alpha$ is "the conjunction of all old and new evidence" for you. Since he's allowing the possibility of forgetting, I don't understand what "old evidence" for you means unless it just means evidence that is among the evidence you currently have but isn't among the evidence you most recently learned.)

  Now we might inquire about the relation between the evidential probability of $p$ for you and the evidential probability of $p$ for you five minutes ago. But we might also inquie about the evidential probability of $p$ for you and the evidential probability of $p$ for you five minutes from now, or the evidential probability of $p$ for you and the evidential probability of $p$ for someone else. What we're really asking about is a comparison between the evidential probability of $p$ on one evidence set (call it $e_1$) and the evidential probability of $p$ on another evidence set (call it $e_2$).

  Here we can employ a simple theorem of the probability calculus that if $e_2$ is a logically consistent set of propositions, $P_1(h) = P(h \mid e_1)$, $P_2(h) = P(h \mid e_2)$, and $e_1 \subseteq e_2$, then $P_2(p) = P_1(p \mid e_2 - e_1)$ for any proposition $p$. (For more on this point and a bit of clarification about the notation, see

"Unlearning.") So in the situation where $e_2$ is your current evidence, $e_1$ is your evidence from some time ago, and $e_1 \subseteq e_2$, BCOND follows (as Williamson notes on p. 220). But there's really nothing (or at least nothing of interest) going on here that's particularly *diachronic*. What's doing the work is $P$'s universal mandating of credences: the fact that there's one universal probability function that gives everyone their evidential probabilities, and the fact that this probability function yields a precise probability for any proposition conditional on any consistent sets of propositions.

- His protestations against operationalizing epistemology notwithstanding, Williamson complains about Jeffrey Conditionalization that it fails to explain how to obtain $P_{new}$. But ECOND combined with E=K has similar problems when applied to the black-and-red balls example. Williamson rejects as "Cartesian" (a four-letter word for him) a Bayesian modeling approach that takes as certain not "A black ball went into the bag" but "I seemed to see a black ball going into the bag" at $t_{old}$. But on his approach, we first have to somehow determine *extrasystematically* that "A black ball went into the bag" is no longer evidence at $t_{old}$ before we can apply ECOND. And even then, to figure out what $P_{new}$ should look like we need something in $e_{new}$ very close to "I seemed to see a black ball going into the bag," which presumably was also in $P_{old}$. So why not take advantage of that evidence all along?

- On the propositionality of evidence: I may be wrong on this point, but Williamson seems to want in general to understand evidence that is represented to the agent as a sentence involving an indexical as being the proposition expressed by the sentence in context (see p. 218 on this point). But suppose I represent my evidence to myself right now as "It is sunny today" in a context where my evidence narrows down the possibilities for "today" to Saturday or Sunday and makes Saturday twice probable as Sunday. What proposition represents my evidence in that context? "It is sunny on Saturday or Sunday" will not do the job when plugged into $P$. Williamson might respond to this example by introducing centered propositions, but we might be able to create a similar example involving a non-self-locating indexical that couldn't be handled with this move.

- What's the point of Williamson's discussion on p. 221 of relativizing to background evidence $f$?

**Section 10.3**

In his treatment of rule-following on p. 223, Williamson discusses the problem that when trying to follow a rule you don't always know whether you know that the condition for the rule obtains. He suggests that you follow the rule and rely on the fact that *most* of the time you know whether the condition obtains. Williamson's response maintains the conception of rule-following as a second-order behavior; it understands rule following as something you consciously and explicitly do with the rule in mind. A better response might be that the self-relfective activity of determining whether one is following a rule often has little to do with the process of following the rule itself.

**Section 10.6**

Quick thought on Williamson's puzzling examples, especially the inscribed rock: One can prove the Reflection Principle from a universally-mandating probability function and various conditions (again, see "Unlearning"). The condition that fails in the rock example is the condition that the set of possible packages of evidence you might receive in the future forms a partition. But that fails only because Williamson is representing the evidence in a particular way, namely $e_2$ & $e_3$ & $e_4$ & $e_5$ or some such. If you represented the evidence in a way that explicitly included not only these propositions but also the process by which you learned the propositions, you could recover the "partitionality" of the possible future evidence packages. Then Williamson's puzzling result for the inscribed rock would disappear and Reflection would hold once more. I suspect there's a way to do a similar thing for his puzzling accessibility example.