Epistemic Probability[1]

In this paper I want to explore both the conceptual and epistemic place of epistemic probability within a foundationalist theory of justification. After arguing that one should take the relation of making probable holding between propositions to be one of the most important conceptual building blocks in epistemology, I examine different ways of trying to understand this key concept and explore the question of how our understanding of it will affect our prospects for finding a defensible version of inferential internalism.

Regress Arguments for Foundationalism:

As I indicated, I'm primarily interested in the place of probability within the framework of foundationalism. More specifically still, I'll focus on the way in which our understanding of probability is involved in the foundationalist's understanding of *inferential* justification. Before setting to that task, it might be helpful, however, to very briefly travel some familiar terrain and remind ourselves of the ways in which foundationalists typically argue for their view about the structure of justification. As we shall see later, the desire to avoid regress that so influences foundationalists might bear directly on the advantages that one particular theory of probability might have for certain versions of internalism.

[1] I read a version of this paper at an APA symposium and profited greatly from the comments of Jim Van Cleve and James Joyce.

The standard argument for foundationalism is the regress argument. But there are, however, two importantly different versions of that argument, an epistemic version, and a conceptual version.

The Epistemic Regress Argument for Foundationalism:

The more familiar of the two regress arguments for foundationalism is probably the epistemic version. Let us say that S's belief that p is inferentially justified if S's justification for believing P is *constituted* in part by the fact that S is justified in believing some proposition other than P, a proposition on which S at least partially bases his belief that P. The following principle seems patently obvious to both externalists and internalists alike. For S to be justified in believing P by basing that belief on some other proposition believed, say E1, S must be justified in believing E1. But if all justification were inferential, then the only way for S to be justified in believing E1 would be for S to be justified in believing yet another proposition E2, and so on *ad infinitum*. Finite minds cannot complete an infinitely long chain of reasoning, so if all justification were inferential we would be forced to the absurd (and epistemically self-defeating) conclusion that no-one is justified in believing anything at all to any extent whatsoever.[2]

The Conceptual Regress Argument:

[2] The conclusion of an argument can be said to be epistemically self-defeating if its truth entails that no-one could have justification for believing that it is true.

The epistemic regress argument rests on a number of controversial presuppositions. As the coherence theorist points out, the foundationalist assumes that justification is *linear* in character and as Peter Klein (1998) argues (in defending the view he calls infinitism), it also presupposes that having justification can only be constituted by a chain of reasoning that one *actually* completes, as opposed to one each link of which *could* be completed. If one allows that E's status as justifier requires only that one be capable of coming up with an appropriate argument for P having E as a premise, then even as finite epistemic agents we may have the potential to come up with an infinite number of different arguments, one for every link in an infinitely long chain of reasoning.

These sorts of responses to the regress argument presuppose that the problematic nature of the regress derives from the need to have an infinite number of different beliefs to serve as justifiers. As I have argued elsewhere (1995), there is another way of construing the problematic nature of the regress foundationalism seeks to avoid. Why *is* it a necessary truth that to be justified in believing P on the basis of E one must be justified in believing E? The most obvious answer, perhaps, is that one's inferential justification for believing P on the basis of E is partially *constituted* by one's justification for believing E. It is an *analytic* truth that one can be inferentially justified in believing P on the basis of E only if one is justified in believing E. But if this is correct, then the above account of *inferential* justification is viciously *circular* if it is intended to be an account of justification in general. Our understanding of inferential justification *presupposes* an understanding of the concept of justified belief (in just the way, for example, our understanding of instrumental goodness presupposes an understanding of intrinsic goodness). Put another way, reflection on the concept of inferential justification

suggests that any plausible definition of inferential justification will be *recursive*. Our understanding of the recursive definition of inferential justification will require an understanding of some base clause that invokes a concept of justification that is not inferential.[3] My desire to end a threatening conceptual regress is the basis of one of my fundamental concerns with Klein's infinitism.

To be sure, the argument is not decisive. Klein would argue that one might have some generic understanding of justification that we could employ in our attempt to understand inferential justification. Perhaps, for example, we could try some deontic conception of justified belief as belief that one ought to have given certain epistemic goals or ends.[4] Employing that generic concept one could still insist that all justification is inferential. But the virtue of a recursive analysis of justification is that one successfully completes two tasks at once. One shows how one ends vicious epistemic regress while analyzing the very concept of justification.

The Principle of Inferential Justification:

S's having a justified belief in E might be a conceptually *necessary* condition for S's justifiably believing P on the basis of E, but it obviously isn't conceptually *sufficient*. At the very least S must *base* his belief that P on E. The analysis of the basing relation is problematic. Some insist that it is exhausted by causal connection between beliefs; others

[3] This is an argument Goldman made in "What is Justified Belief?" Both classic foundationalists and their newer externalist/reliabilist cousins hold very similar views about the need to embrace a foundationalist *structure* for justification.

[4] Foley (1987) pursues a variation on this strategy.

insist that it involves something else. In any event, it is again obvious that S's justification for believing P on the basis of E requires more than that S be justified in believing E and that S base his belief that P on E.[5] The most obvious candidates for the missing condition are the following:

1) There is a correct epistemic rule sanctioning the move from believing E to believing P

2) S is *aware of*, or has a *justified belief that* there is a correct epistemic rule sanctioning the move from believing E to believing P,

3) There is an appropriate logical or probabilistic connection between the propositions E and P, or

4) S is *aware of*, or has a *justified belief that* there is an appropriate logical or probabilistic connection between E and P.

Now on analysis there may be no difference between the contents of 1) and 3) (and therefore between 2) and 4)). One might argue that to claim that E makes probable P, for example, *in the sense relevant to epistemic justification*, is just a way of acknowledging that there is an epistemic rule licensing the move from believing E to believing P. Conversely, one might argue that all this talk about the correctness of epistemic rules is itself a convoluted way of talking about relationships between propositions. So, for example, there are rules of *deductive* logic that permit certain sorts

[5] As I'm using the locution, S can be justified in believing P on the basis of E even though S is not justified in believing P on the basis of his entire body of evidence.

of inferences, say modus ponens, but the most obvious answer to the question of what makes modus ponens a *correct* rule of inference is that there is a relation of entailment holding between the premises and conclusions of arguments having a certain form, where this entailment exists independently of any rule acknowledging it. Similarly, in the case of inductive reasoning, one might argue that to say that the premises of an inductive argument make probable its conclusion is just a way of pointing out that there is a correct rule licensing the inference to the inductive conclusion from the premises describing observed correlations. Conversely, one might argue that the rule itself merely acknowledges the relevant relation of making probable that holds between premises and conclusion of the argument whether or not there exists some rule that takes account of that relation.

While 1) and 3) may be alternative ways of making a common claim, it is important, I think, to decide the direction of the reduction. We need to figure out what the relevant conceptual building blocks are in trying to understand inferential justification. I'll return to this issue later. But, however, one decides the issue of conceptual priority, one still needs to decide whether it is the *existence* of the probability connection/correct rule or our *awareness* of the relevant connection/rule that is crucial to the possession of inferential justification. It is disagreement on this last point, I have argued elsewhere (1995), that gives us one (though only one) way of isolating a point on which self-proclaimed internalists and externalists often disagree. One might label the view that insists that for one to be inferentially justified in believing P on the basis of E one must be *justified in believing* that E makes probable P (where entailment can be

viewed as the upper limit of making probable) or, alternatively, that the inference from E to P is sanctioned by a correct epistemic rule, *inferential internalism*.

Leaving aside for a moment the correct analysis of probabilistic connection, what arguments can we offer for inferential internalism? It seems obvious that something like the view was simply taken for granted, explicitly or implicitly, by most of those in the history of philosophy who either argued for skepticism or took the problem of skepticism seriously. But given the enormous difficulties of meeting the skeptical challenge within the constraints of inferential internalism, why accept the view? We are surely rather liberal in our allowing that all sorts of people, including of course the philosophically unsophisticated (and the philosophically sophisticated with false views about justification) have all sorts of perfectly rational beliefs. Given that we seem not the least bit inclined to abandon the view that people have justified beliefs about the external world, the future, and the past, despite the difficulty they have coming up with any sort of reasons indicating the legitimacy of their inferences, should we not at least suspect that the ordinary understanding of justification requires nothing as strong as what is proposed by the inferential internalist?[6] Perhaps, but there also appears to be prima facie evidence to indicate widespread acknowledgement of the inferential internalist's requirements for inferential justification.

Consider the palm reader who predicts that I will have a long life based on the belief that I have the proverbial long "life line" on the palm of my hand. It seems

[6] Of course, one might argue that their difficulties reflect an implicit awareness that talk of inference in this context is highly misleading. One might also argue that one can "see" a probabilistic connection without being able to articulate it.

obvious that a *sufficient* condition for rejecting the palm reader's inference as rational is that the palm reader has no reason to believe that the length of a line on one's palm has anything to do with the probability of one's living to a ripe old age. If my high priest predicts that the war I am planning will go badly based on the observation that the entrails of a recently dissected bird are bloody, an epistemically rational person will surely demand evidence for supposing that features of entrails are correlated with success in battle kind before conceding the rationality of the priest's prediction. These commonplace examples and indefinitely many others like them surely indicate that we do in fact embrace the inferential internalist's account of what is necessary for inferential justification. We may pick and choose when we decide to make an issue of someone's lacking reason to believe in a legitimate evidential connection. We do sometimes take for granted the justification of certain beliefs and the legitimacy of certain inferences when we are primarily concerned with the justification of other beliefs and inferential connections. But once we take seriously a skeptical challenge to a commonplace inference, we must apply the same standards that we insist on applying to astrologers and fortune tellers.

Mike Huemer (2002) objects to the above argument for inferential internalism. He argues that the examples I use to make initially attractive the principle are misleading in that they inappropriately characterize the evidence from which one infers the relevant conclusion. Even palm readers don't think that they can legitimately infer their predictions from propositions describing the character of a person's palm *and from that information alone*. I suspect that the priests at Delphi didn't think of themselves as inferring truths about battles from the appearance of entrails *and from that alone*. It

should be a truism that much of the argument we actually give outside of a philosophical context is highly compressed, highly enthymematic. As we ordinarily use the term ”evidence”, we certainly do characterize litmus paper’s turning red in a solution as evidence that the solution is acidic. The approach of very dark clouds is evidence that there will soon be a storm. Footprints on a beach is evidence that someone walked on the beach recently. But it is surely obvious upon reflection that one’s evidence for believing that the solution is acidic, for example, is not the color of the litmus paper *by itself*. To legitimately draw the conclusion one would need an additional *premise*, most likely a premise describing a correlation between the color of litmus paper in a solution and the character of that solution.

Once one realizes that the reasoning in the examples I discussed above is enthymematic, one is positioned to respond to that appearance of an argument for inferential internalism. For the reasoning described above to be legitimate, it *is* indeed necessary to have some justification for believing that there is a connection palm lines and life expectancy, the approach of dark clouds and storms, footprints on a beach and the recent presence of people, but *only* because propositions describing connections or correlations of the relevant sort are implicitly recognized as critical *premises* from which the relevant conclusions are drawn. As we saw earlier, internalists *and* externalists alike typically share the foundationalist’s insight that inferential justification is parasitic upon the justification we possess for believing the relevant premises of our arguments. If the palm reader is relying on an unstated, but critical premise describing correlations between palm lines and length of life in reaching her conclusion, she will, of course, need justification for believing that *premise* (in exactly the sane unproblematic sense that she

will need justification for believing the premise describing the length of the palm line itself). But that in no way suggests that when we have fully described all of the relevant premises from which a conclusion is drawn, we should require that the person who draws that conclusion needs additional evidence for believing that the premises make probable the conclusion. The existence of the relevant connection between premises and conclusion is enough.

While Huemer is right, I think, in arguing that the examples discussed above do not support inferential internalism, it is not clear to me that one can't make just as strong a case for inferential internalism focusing on non-enthymematic reasoning. Consider the case of someone who infers P from E where E logically *entails* P. Is the inferential internalist right in maintaining that in order for S to believe justifiably P on the basis of E, S must be aware of the fact that (or at least have a justified belief that) E entails P (or alternatively, that the inference in question is legitimate)? The answer still seems to me obviously "Yes." We can easily imagine someone who is *caused* to believe P as a result of believing E where E does in fact entail P, but where the entailment is far too complicated for S to understand. Unless S "sees" that P follows from E, would we really allow that inference in question generates a justified belief? Or to make my case a bit stronger, would we allow that the person who reaches the conclusion has philosophically relevant justification or ideal justification—the kind of justification one seeks when one searches for philosophical assurance.[7]

[7] I argue in "Achieving Epistemic Ascent" that one might accommodate at least some externalist intuitions by allowing derivative concepts of inferential justification. Perhaps in certain contexts we will concede the justificatory status of a belief if it was inferred from premises that do make probable the conclusion, particularly if we think that the existence of the probability connection is causally connected to the person's

A great deal hinges on how we understand the critical concept of someone's believing one proposition on the basis of another. There is some plausibility to the claim that genuine inference involves more than mere causal connection. But the source of this intuition will, I think, provide little solace for the inferential externalist. In fact, I suspect that we may not concede that there has been a genuine inference unless there has been a veridical or nonveridical "perception" of a connection between that from which P is inferred and P. But, of course, this "perception" will be just what the inferential internalist claims is the awareness of connection that is necessary for inferentially justified belief.

There is at least some concern that inferential internalism will lead to regress. The view does at least remind one of Carroll's (1895) famous dialogue between the Tortoise and Achilles.[8] Paraphrasing liberally, the Tortoise admits at one point that P is true and also that if P is true then Q is true but doesn't see why that's a reason to believe Q. Obligingly, Achilles plays the game and adds an additional premise: If P is true and if it is true that if P is true then Q is true then Q is true. Even so, wonders the Tortoise, why does that premise coupled with P and (if P then Q) give one reason to believe Q. Isn't the inferential internalist in the position of the Tortoise who keeps insisting that even when one has in one's possession evidence that entails one's conclusion, epistemically rational belief in the conclusion requires yet additional reason for supposing that the entailment holds? Like the Tortoise the inferential internalist *does*

---

willingness to form belief in the conclusion. Relaxing our standards still further, I suggest, we may even allow that a belief is justified if it is caused by a *fact* that is the truth-maker for a proposition that makes probable the proposition believed.

[8] Jim Van Cleve suggested to me that inferential internalism might give rise to the regress.

require for justified belief in a conclusion something other than mere justified belief in premises which do, in fact, entail the conclusion. The inferential internalist does insist that the person possessing the justification be aware of the entailment. But that requirement for inferential *justification* does not, obviously, suggest that the *argument* needs additional premises in order to be good argument. But even if we make that distinction, we should recognize that to avoid vicious regress the inferential internalist may need to ensure that the relevant awareness of connection between premises and conclusion does not itself require *inference* from still additional premises. But we'll have more to say about that shortly.

I have dwelt on Huemer's argument against inferential internalism not just because I'm interested in the question of whether or not the argument succeeds. It seems to me that Huemer's insightful argument reminds us of certain features of our talk of evidence that will be important to keep in mind when we evaluate the plausibility of certain views about the nature of probability.

The Analysis of Epistemic Probability:

Whether or not we adopt inferential internalism, we need an analysis of the probability connection that by itself or as the object of awareness is partially constitutive of inferential justification. Earlier we suggested two ways in which one might think of the epistemic claim that one proposition or conjunction of propositions makes probable another. We might think that the truth of the probability claim derives from the existence of a correct *rule* sanctioning the inference in question. Or we might think that our

understanding of the probability connection is prior to and legitimizes our putting forth the rule that sanctions the relevant inference.

The rule-oriented approach itself can be thought of in quite different ways. On the most radical view, the rules of inference might be thought of as analogous to the rules of a game, rules that are themselves neither true nor false (though the claim that they are the rules of the game is either true or false). Rules thought of this way are more like imperatives that prescribe, permit, and prohibit certain inferential "moves." What we obviously need from the proponent of such a view, though, is some answer to the question of what makes a given rule "correct", or the right rule to employ.[9] There are many possible responses to this question. On one extreme one might advance a thoroughgoing subjectivist/relativist position, according to which there is no possible non-question begging answer to the question of how to choose rules, since any answer would itself presuppose the legitimacy of certain rules. Such a view seems to lead inevitably to a kind of philosophical anarchy.

Chisholm and his followers are inclined to take as basic certain *epistemic* concepts that apply to beliefs and would employ these concepts in turn in defining the correctness of epistemic rules, and more generally the concept of epistemic probability.[10] So on Chisholm's view there are facts about what it is reasonable to believe that

---

[9] There is an exactly analogous question that arises for rule utilitarians. One can perhaps define morally justified and morally unjustified action by reference to rules that require, permit, and prohibit certain sorts of actions, but we are in desperate need of a criterion for choosing between alternative rules.

[10] So, for example, in Chisholm (1977), the primitive "more reasonable to believe than" is used to define "beyond reasonable doubt," "has some presumption in its favor," "is acceptable," "is certain," and "is evident."

supervene on certain properties of believers. The noninferential justification of certain beliefs supervenes on non-epistemic features of various conscious states (such as what I seem to remember or seem to perceive). But once we get justified beliefs, there are additional synthetic necessary truths that tell us which other beliefs are justified provided that we have certain justified beliefs. We can, if we like, summarize these truths with rules that permit, require us, or prohibit us from forming certain beliefs when we have others.

The pragmatist might suggest that the legitimacy of a rule is a function of how well we get on employing the rule. But depending on how one understands getting along well, that view might collapse into a view I'll talk about shortly, a view that understands the critical concept of probability in terms of *frequency*. If, for example, a rule of inference is legitimate only in so far as it generates more true beliefs than false beliefs (when the input beliefs are true), then the theory that appears to take rules as its conceptual cornerstone is better understood as a version of the reliabilism that has become so popular primarily through the writings of Alvin Goldman.

Finally, one might take the legitimacy of a rule to be a function of internal relations of making probable (more about this shortly) that hold between propositions. But if one adopts this approach one will again be in a position to eliminate reference to rules in the analysis of inferential justification. It will be our understanding of the relations that hold between propositions that is key to our understanding inferential justification.

In what follows I want to focus on the last two approaches to understanding epistemic probability. I'm not sure many are willing to embrace the philosophical

anarchy that accompanies a view in which arbitrary epistemic rules occupy center stage. And while there may be synthetic necessary truths about what one is justified in believing when one is justified in believing certain other propositions, it is hard for me to believe that the justificatory status of inferentially justified beliefs is not fundamentally derived from relationships between that which is believed. Put another way, it is surely a feature of the *arguments* whose premises and conclusions are believed that is key to understanding the justificatory status of the beliefs formed in the conclusions as a result of justified belief in the premises.

Suppose we agree that a key to understanding inferential justification is an understanding of the relation of making probable that holds between the premises and conclusions of arguments. What's the best way of understanding that relation? That debate has a long history, one that predates, but in many ways foreshadows, the now more familiar internalist/externalist controversies in epistemology.[11] Painting with a very broad stroke, one can attempt to analyze probability claims in epistemology on the well-known model of relative frequency that is offered as a way of interpreting claims about the probability of an individual or event having a certain characteristic. On a very crude interpretation of the frequency theory, to say of something that it is probably G is always elliptical for a more complex relativized claim of probability. One must refer the individual about which the probability claim is made to some reference class, say F. The more perspicuous statement of the probability claim is one about the probability of a's being G relative to its belonging to the class F. On the crudest and least plausible version

[11] One of the most interesting debates that has clear implications for the internalism/externalism controversy can be found in Keynes (1921) and Russell (1948) Part V.

of the view, the truth conditions for the claim of relative probability are determined by the percentage of F's that are G. The higher the percentage of F's that are G the more likely it is that something is G relative to its being F. Of course, we very often don't *explicitly* supply the relevant reference class for a probability claim. A great deal of time and energy has been spent by philosophers trying to figure out what reference class is the appropriate *default* for ordinary probability claims. Is it a class that is ontologically homogeneous, or one that is epistemically homogeneous?[12] If epistemically homogeneous, with respect to *whose* knowledge is the homogeneity defined? I'm not sure that there are unambiguous answers to these questions, and in any event I'm not concerned with these questions now. It is also fairly obvious that the relative frequencies that constitute the truth conditions for probability claims are not actual frequencies—one will inevitably need to turn to counterfactuals with all of the problems that that move meets. My main concern, here, however, is the extent to which one can incorporate the alleged insight of a relative frequency theory of probability into an analysis of the *epistemic* probability that we are assuming holds between propositions.

One could borrow at least the spirit of the relative frequency interpretation of probability and apply it to relations between propositions in the following way. We could suggest that in claiming that P is probable relative to E we are simply asserting that E and P constitute a pair of propositions, which pair is a member of a certain class of

[12] Let see that a reference class F is ontologically homogeneous with respect to some characteristic G when there is no way of dividing the class further such that the frequency with which things are G relative to membership in the subclass is different from the frequency with which things are G relative to F. A reference class F is epistemically homogeneous with respect to G when as far as we know the class is ontologically homogeneous.

proposition pairs such that, when the first member of the pair is true, usually the second is. Thus in saying that a's being G is probable relative to its being F and most observed F's being G, I could be construed as claiming that this pair of propositions is of the sort: Most observed X's are Y and this is X/This is Y, and most often it is the case that when the first member of such a pair is true, the second is. Similarly, if I claim that my seeming to remember eating this morning (E) makes it likely that I did eat this morning (P), I could be construed as asserting that the pair of propositions E/P is of the form S seems to remember X/X, such that most often when the first member of the pair is true, the second is.

The above view obviously resembles, at least superficially, the reliabilist's attempt to understand justified belief in terms of reliably produced beliefs. And it encounters many of the same difficulties. Just as the relative frequency theory of probability must inevitably move beyond actual frequencies in defining probability, so both the above account of epistemic probability, and the reliabilist will inevitably be forced to move beyond actual frequencies in order to define the relevant epistemic probability/reliability. Just as reliabilism must deal with the generality problem, so the above approach to understanding epistemic probability as a relation between propositions must deal with the problem of how to choose from among alternative ways of characterizing the class of propositions pairs to which a given pair belongs. In evaluating the reliability of beliefs produced by memory, for example, the reliabilist must decide whether or not to lump together faint and vivid apparent memories, apparent memories of events in the distant past and events in the recent past, apparent memories of emotions and apparent memories of memories, vivid memories that occur in young people and

vivid memories that occur in old people. A frequency approach to understanding epistemic probability can make the same sorts of distinctions between pairs of propositions, and consequently has the same sorts of decisions to make. Just as many reliabilists are troubled by the implications of their view for what to say about worlds in which demons consistently deceive epistemically "faultless" believers, so a frequency theory of epistemic probability must deal with similar alleged counterintuitive consequences about what is evidence for what in demon worlds. Lastly, both reliabilism and the frequency theory of epistemic probability will be anathema to the inferential internalist who is convinced that one needs *access* to probability connections in order to gain philosophically satisfying inferential justification. The inferential internalist who is a foundationalist will need to end a potential regress when it comes to gaining access to probabilistic connections. If one's model for foundational knowledge is something like knowledge of truths made true by facts with which one is directly presented, there seems no hope that one will get that kind of access to either the reliability of a belief-forming process or a probability relation (understood in terms of frequency) holding between propositions.

One of the historically most interesting alternatives to the frequency interpretation of epistemic probability is a view developed some eighty years ago by Keynes (1921). Keynes wanted to model epistemic probability on entailment. He held that just as one can be directly aware of entailment holding between two propositions, so one can also be directly aware of a relation of making probable holding between two propositions. There are, of course, obvious differences between entailment and making probable. From the fact that P entails Q it follows that the conjunction of P with any other proposition entails

Q. From the fact that P makes probable Q, it doesn't follow that P together with anything else makes probable Q. But for all that, we could still take making probable to be an a priori internal relation holding between propositions (where an internal relation is one that necessarily holds given the existence and nonrelational character of its relata). P and Q being what they are it cannot fail to be the case that P makes probable Q. (It might also be true that P, R, and Q being what they are it cannot fail to be the case that (P and R) makes probable not-Q).

Which view of probability is correct? One might approach an answer to this question by looking at the most uncontroversial upper limit of making probable—entailment. But it quickly becomes apparent that entailment is a double-edged sword when it comes to serving as a paradigm for understanding probability. The Keynesean will, of course, be right to stress the fact that entailment is an internal relation knowable a priori. But the frequency theorist (or the reliabilist) can equally stress that valid deduction is a paradigm of a conditionally reliable belief-producing process (a paradigm of pairs of proposition types such that when the first member of the pair is true, the second is as well).

Against the Keynesean, one might argue that it is patently absurd to suppose that making probable is an internal relation holding between propositions. Such a view yields the absurd consequence that claims about evidential connections are necessary truths knowable a priori. If *anything* is obvious it is that the discovery of evidential connections is a matter for empirical research. But while the objection might *seem* initially forceful, *one must remember the point we conceded in considering Huemer's objections to inferential internalism*. There is certainly no necessary connection between litmus

paper's turning red in a solution and the solution's being acidic, between dark clouds and storms, between footprints on a beach and the prior presence of people. But then, on reflection, we decided that it is misleading to characterize the litmus paper, dark clouds and footprints as *the* evidence from which we infer the respective conclusions. What we call evidence in ordinary parlance is just a *piece* of the very elaborate fabric of background information against which we draw our conclusions. So we shouldn't *expect* to find Keynesean probabilistic connections holding between, for example, the proposition that the litmus paper turned red and the proposition that the solution is acidic. Where *should* we look for a plausible example of Keynes's relation of making probable?

The obvious, though perhaps not all that helpful, answer is that we should look for it wherever we have what we take to be legitimate, *non-enthymematic* and non-deductive reasoning. One needn't, and probably shouldn't, insist that even if probability connections between propositions are know*able* a priori that they are *easy* to know a priori. On some views, all mathematical truths are knowable a priori but as we painfully learned in math classes, their a priori character doesn't necessarily make the final for the course easy. Keyneseans have been given considerable grief for the fact that they may have come up with bad examples of alleged necessary truths about probability.[13] Various formulations of the principle of indifference, for example, are notoriously seductive but also notoriously problematic. The difficulty has always been to find the "right" way to characterize the continuum of alternative hypotheses whose probability can then be "divided" equally. If I know that something is either red or not-red, but don't have any

[13] Gillies (2000) seems to rely on counterexamples to putative probability connections to attack the Keynesean view.

evidence that bears on the thing's specific color one might suppose that it is just as likely relative to that evidence that it is red as that it is not-red. But a bit of reflection tells us that it is unreasonable to treat being red and being not-red the same way. There are many more ways of being not-red than there are ways of being red. The examples that give rise to paradox suggest, however, that there may be no unproblematic way of finding the appropriate way to divide up the properties along a continuum.[14] If it should turn out that there is no useful principle of indifference available to the epistemologist, it doesn't follow, of course, that a Keynesean conception of epistemic probability is doomed. The Keynesean should simply look elsewhere for plausible examples of propositions standing in the relation.

The trouble, of course, is that philosophers don't agree with each other about what constitutes legitimate but deductively invalid reasoning. Notwithstanding difficulties posed by Goodman's new (now not so new) riddle of induction, one might look at the relationship between the premises and conclusion of an enumerative inductive argument. Less plausibly, perhaps, one might think about the connection between the proposition that I seem to remember having an experience and the proposition that I had the

[14] Consider a well-known example. You know that I drove the mile between point A and point B traveling somewhere between 30 mph and 60 mph and thus taking somewhere between 1 minute and 2 minutes to make the trip. So what's the probability that I was going between 30 – 45 mph? It's surely just as likely as that I was going between 45 – 60 mph—the probability must be ..5. And what's the probability that it took me somewhere between a minute and a minute and a half. Also the same as that I took between a minute and a half and two minutes--.5. But it turns out that you can't assign .5 probability to both the hypothesis that I was traveling between 30 and 45 mph and the hypothesis that it took me between a minute and a minute and a half.

experience. Still more problematically, we might suggest that there it is some sort of synthetic necessary truth that when I seem to see some physical object that is red and round that makes likely that there is some object that is red and round. How plausible is it to suppose that there are necessary truths asserting that our putative evidence in the above examples makes at least prima facie problem the conclusions?

Well how do we generally assess the plausibility of the claim that a certain proposition is necessarily true? We often start by asking ourselves whether we can conceive of a situation in which the proposition in question is false. And here, it seems we are in a position no more, but no less plausible than that critic of externalist analyses of justification who invokes demon world scenarios in order to cast doubt on the externalist's view. The critic of reliabilism, for example, asks you to consider a possible world in which our sensory experiences have been produced by a demon bent on inducing in us a massively mistaken system of beliefs. The victim of the demon, the argument goes, surely has as much reason to believe propositions about the external world as do we with our phenomenologically indistinguishable sense experience. Because the reliabilist seems committed to the view that the unreliable belief-producing process in the demon world yields unjustified belief, while the reliable belief-producing process in the world as we take it to be yields justified beliefs, the reliabilist has an implausible view.

This objection to reliabilism is actually neutral with respect to the epistemic status of beliefs about the external world based on perception. It states only that *whatever* epistemic status such beliefs have in the world of veridical perception, they surely have that same status in the demon world. But if one adds the premise that both our beliefs

and the beliefs of the demon's victims are *justified*, then one seems to be very close to endorsing the view that one can't *conceive of* having the kind of perceptual evidence we have without that evidence making probable the truth of what we believe, at least in the sense of "making probable" relevant to the possession of epistemic justification. Since many internalists (and even some externalists) have felt the force of the demon world objection to reliabilism, the view that making probable is an internal relation between propositions should be at least initially attractive to many epistemologists.[15]

At present, I am arguing only for the very modest conclusion that the Keynesean approach to understanding epistemic probability is a view worth considering seriously. When one distinguishes partial "evidence" from genuine evidence (the body of propositions in its *entirety* from which we infer conclusions), and when one keeps firmly in mind the ways in which making probable would differ from entailment (even if it is an internal relation between propositions), it is not that hard to take seriously the idea that one couldn't seem to remember having done X, for example, without that rendering probable having done X.

The view that there is an internal relation of making probable that holds between propositions is also just what the inferential internalist desperately needs in order to avoid vicious regress. While the classic foundationalist recognizes the need to cauterize the chain of reasoning that threatens to extend infinitely into the past, the inferential internalist needs to fear not one but an infinite number of infinite regresses. Just as one's

[15] Goldman himself was clearly troubled by the problem in his early paper on reliabilism (1979). Concern with the general problem led him first to "normal worlds" reliabilism (1986), and ultimately to bifurcate the concept of justification (1988).

inferential justification for believing P must be traced ultimately to something one is noninferentially justified in believing, so one must find evidential connections one can justifiably accept without inference. If one infers P from E one must not only be justified in believing E, but one must be justified in believing that E makes probable P. One might be able to infer that E makes probable P from some other proposition F, but then one must not only be justified in believing F, one must be justified in believing that F makes probable that E makes probable P. If inferential internalism and foundationalism are true, then unless we are to embrace a fairly radical skepticism, we must find *some* proposition of the form E makes probable P that we can justifiably believe without inference. Since most foundationalists will concede that there are at least some propositions of the form E entails P that one can know without inference, the closer we can make our analysis of making probable resemble our analysis of entailment, the more plausible will be the claim that we can know without inference propositions of the form E makes probable P.[16]

Richard Fumerton

University of Iowa

[16] Notice that the Keynesean approach to understanding probability does not require that for one to be inferentially justified in believing P on the basis of E one must know or be able to formulate *general* principles of probability. One might be able to see the connection between particular propositions without seeing how to generalize. An analogous point holds of entailment. One can see that P entails Q without being able to see that entailment as an instance of modus ponens, modus tollens or any other general *kind* of entailment.

References

Carroll, Lewis (1895). "What the Tortoise said to Achilles." *Mind*, 4, No. 14, 278-80.

Chisholm, Roderick (1977). *Theory of Knowledge*, 2nd Edition. Prentice-Hall.

Foley, Richard (1985). *The Theory of Epistemic Rationality*. Harvard University Press.

Fumerton, Richard (1995). *Metaepistemology and Skepticism*. Rowman and Littlefield.

------. Forthcoming. "Achieving Epistemic Ascent" in *Ernest Sosa and His Critics.* Blackwell.

Gillies, Donald. (2000). *Philosophical Theories of Probability*. Routledge.

Goldman, Alvin (1979). "What is Justified Belief?" in *Justification and Knowledge*, ed. George Pappas. Dordrecht.

------. (1986). *Epistemology and Cognition*. Harvard University Press.

------. (1988). "Strong and Weak Justification." Pp. 51-69 in *Philosophical Perspectives.* Ed. James Toberlin. Ridgeview Publishing.

Klein, Peter (1998). "Foundationalism and the Infinite Regress of Reasons," *Philosophy and Phenomenological Research*, LVIII, 919-26.

Huemer, Mike. 2002. "Fumerton's Principle of Inferential Justification." *Journal of Philosophical Research* 27, 329-40.

Keynes, John. 1921. *A Treatise on Probability*. Macmillan.

Russell, Bertrand. 1948. *Human Knowledge: Its Scope and Limits*. Simon and Schuster.