FOUNDATIONS OF SKEPTICAL THEISM: CORNEA, CORE, AND CONDITIONAL PROBABILITIES

Stephen J. Wykstra and Timothy Perrine

Abstract: Some skeptical theists use Wykstra's CORNEA constraint to undercut Rowe-style inductive arguments from evil. Many critics of skeptical theism accept CORNEA, but argue that Rowe-style arguments meet its constraint. But Justin McBrayer argues that CORNEA is itself mistaken. It is, he claims, akin to "sensitivity" or "truth-tracking" constraints like those of Robert Nozick; but counterexamples show that inductive evidence is often insensitive. We here defend CORNEA against McBrayer's chief counterexample. We first clarify CORNEA, distinguishing it from a deeper underlying principle that we dub 'CORE.' We then give both principles a probabilistic construal, and show how on this construal, the counterexample fails.

The new "inductive atheism" argues that certain empirical features of evil are strong *inductive* (or "probabilistic") evidence against theism. A feature stressed by William Rowe, for example, is the "noseeum" character of much suffering. We can, try as we may, *see no* God-justifying[1] (good served by much suffering. And our *seeing no* God-justifying good served by an instance of suffering is, it is argued, strong evidence for there *being no* God-justifying good served by it—and hence also, by a further short step, for there being no God. Against this reasoning, so-called "skeptical theists" press this question:

> Granted, atheism makes the feature you cite—here, the noseem feature—entirely expectable. But isn't this feature also pretty expectable if it were the case that God exists? If God were to exist, shouldn't we expect—God being God and us being us—to often *not see* the goods He purposes for many evils? And if that's so, how can this feature be regarded as strong evidence that God *doesn't* exist?

The skeptical theist here employs a "neutralizing tactic"—a tactic for defusing alleged strong evidence—that we can find used in many contexts.[2] While this neutralizing strategy is intuitively appealing, it is not easy to adequately formulate the implicit principle on which it rests. One formulation has been Wykstra's CORNEA—the *Condition of ReasoNable Epistemic Access.*[3]

---

[1] By a "God-justifying good" served by for an evil, we mean a good that would suffice to justify an all-powerful, all-knowing, and entirely good Creator in allowing that evil. The question of whether we see any such good in a given case is independent of whether God exists and of whether that good is the actual reason justifying God in allowing the evil.

[2] The New Testament scholar F.F. Bruce thus opens *Jesus and Christian Origins Outside the New Testament* (Grand Rapids: Eerdmans, 1982) by citing a correspondent vexed by whether, in contemporary documents outside the New Testament, one finds any "collateral proof of the historical fact of the life of Jesus Christ." Bruce counters by asking (17): in "which contemporary writers during the first fifty years after the death of Christ, would you *expect* to find the collateral evidence you are looking for?" Bruce then argues that we should not, in the nature of the case, expect to find the sort of collateral evidence the correspondent finds lacking.

[3] For earlier discussions of CORNEA, see Stephen Wykstra, "The Humean Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance,'" *International Journal for the Philosophy of Religion*, 16 (1984), 73-93; "Rowe's Noseeum Arguments from Evil," in *The Evidential Argument from Evil,* ed. Daniel

Although CORNEA has morphed through many versions, all versions offer an acid test for whether one can reasonably regard some piece of evidence E as being strong[4] evidence for some hypothesis H. The test is this: ask whether, if H were *false*, E would likely be *different*. If the reasonable answer—what one is warranted in asserting—is "No," then one *isn't* entitled to regard E as strong evidence for H.[5]

Some critics, including Rowe himself, have tended to accept the CORNEA test, arguing only that the favored evidential feature passes the test. But others have called the test itself into question, including most recently Justin McBrayer.[6] McBrayer notes that CORNEA resembles "truth-tracking" or "sensitivity" theories of knowledge like that of Robert Nozick. Such theories place "sensitivity" requirements on knowledge—requirements usually stated using subjunctive conditionals. To say that my belief that the chimes on my patio are now ringing is "sensitive" means that if they *weren't* now ringing, I would not now be forming the belief that they are. Nozick's idea that a belief that *p* is knowledge only if it is sensitive to *p* is attractive, but it is now widely seen as falling to counterexamples—counterexamples arising, especially, for *inductive* knowledge.[7] And CORNEA, McBrayer thinks, falls to similar counterexamples. It is, he writes, "a *sensitivity* constraint on evidence, and inductive evidence is often *in*sensitive."[8]

McBrayer's counterexamples provide a challenging occasion to probe the foundations of skeptical theism. We here focus on two foundational aspects that, in light of McBrayer's treatment, need illumination. First, McBrayer—appropriating a recent suggestion of Wykstra—directs his critique against a reformulation of CORNEA that changes CORNEA into a criterion of when something "counts as [strong] evidence." To develop Wykstra's suggestion more faithfully, we distinguish two principles—the strongly internalist CORNEA, and a strongly externalist principle *behind* CORNEA that we call CORE:

> (CORE) In cognitive situation S giving new input E, E is levering evidence for hypothesis H only if it is true that if H were false, E would likely be different.

CORNEA expresses an internalist constraint on when *one is entitled to regard* E as strong evidence for H, whereas CORE expresses an externalist constraint on when E *actually is* strong evidence for H. Still, CORE and CORNEA share a common interest, for on each, the crux of the

---

Howard-Snyder, (Bloomington, IN: Indiana University press, 1996), 126-150; and "CORNEA, Carnap, and Current Closure Befuddlement," *Faith and Philosophy*, 24.1 (January 2007), 87-98.

4 Below, we follow Wykstra in defining "strong evidence" more precisely as "levering evidence."

5 CORNEA offers only a necessary condition for being entitled to regard E as strong evidence for H. If the reasonable answer is "Yes," then one may or may not be entitled to regard E as strong evidence for H.

6 Justin McBrayer, "CORNEA and Inductive Evidence," *Faith and Philosophy* 26:1 (January 2009), 77-83. All McBrayer citations refer to this paper.

7 The most well-known sensitivity theory of knowledge is that developed by Robert Nozick in *Philosophical Explanations* (Harvard: Belknap Press, 1981), 167ff.The kinship between CORNEA and Nozickean sensitivity is presupposed by Richard Swinburne's critique in "Does Theism Need a Theodicy?" *Canadian Journal of Philosophy*, 18.2 (June 1988), rebutted by Wykstra's 2004 APA paper "Skeptical Gambits and the Two Roads to Atheism: Swinburne vs. CORNEA." As Chad McIntosh has called to our attention, CORNEA also has close kin in a criterion (for "taking") given by Roderick Chisholm in his "'Appear,' 'Take,' and 'Evident'", reprinted in *Perceiving, Sensing, and Knowing*, ed. Robert J. Schwarz (U. Cal. Press, 1965), p. 484. For an account of why inductive knowledge poses problems for Nozick's account, see Jonathan Vogel's "Tracking, Closure, and Inductive Knowledge," in *The Possibility of Knowledge: Nozick and his Critics*, ed. Steven Luper-Foy (Totowa, New Jersey: Rowman & Littlefield, 1989), esp. pp. 204-13. For signs of an opposing groundswell, see Keith DeRose, "Insensitivity is Back, Baby!" *Philosophical Perspectives*, 24, Epistemology, 2010, 161-187.

8 McBrayer, 77. Italics ours.

matter, so to speak, concerns the truth or warranted assertibility (respectively) of a crucial conditional that we dub "*crux*":

(*crux*) if H were false, E would likely be different.

The second foundational aspect we aim to illuminate is how best to construe *crux*. Now *crux* is a conditional "were-would" sentence—what we call a *grammatical* subjunctive. Perhaps due to this grammatical feature, McBrayer sees it as a *logical* subjunctive or counterfactual, and so treats it using standard "closest possible worlds" semantics.[9] We propose taking it as a "non-counterfactual subjunctive"[10] best treated as a conditional probability, for which the tool of choice is the probability calculus.[11] This probabilistic treatment enables both CORE and CORNEA to handle McBrayer's chief counterexample; it also, we hope, improves the foundations skeptical theism, bringing it into relation with issues of broader epistemological interest.

1. The CORE Behind CORNEA

To develop the suggestion that behind the internalist CORNEA is a deeper externalist CORE, we do four things here. First, we locate two versions of CORNEA within the dialectic of Rowe's own development of his argument. Second, we explain the type of strong evidence—namely, "levering evidence"—that CORNEA is a constraint on. Third, we distinguish CORNEA from CORE, contrasting our version of CORE with a flawed version—which we dub "McCORNEA"—offered by McBrayer. And fourth, we introduce the probabilistic approach to *crux*, the crucial conditional in both CORE and CORNEA.

1.1. Rowe's Appears-Idiom and Two Versions of CORNEA

CORNEA proposes a necessary condition on the type of confirming evidence that Rowe-style arguments purport to offer. Rowe's earliest arguments couch this evidence using the appears-idiom.[12] A central strand of his argument can be represented as follows. Consider a case of a fawn horribly burned in a distant forest fire, lying in suffering for days before dying. Our cognitive situation is that:

(Rowe-1) We can, try as we may, *see no* God-justifying good served by the fawn's suffering.

This feature of our cognitive situation, Rowe thinks, entitles us to claim:

(Rowe-2) *It appears that* there *is no* such God-justifying good served by the fawn's suffering.

---

[9] On such a standard semantics, a broadly counterfactual conditional is true just in case the consequent is true in the closest antecedent world(s)—the closest world(s) where the antecedent is true. See Lewis, *Counterfactuals*, (Oxford: Blackwell, 1973), and Stalnaker, "A Theory of Conditionals," in *Conditionals*, eds. Frank Jackson, (Oxford: Oxford, 1991), 28-45. The differences between Lewis and Stalnaker will not affect our treatment here.

[10] Robert Stalnaker, "Indicative Conditionals," in Jackson, *Conditionals*, 136-54, especially 146. See also Alan Ross Anderson, "A Note on Subjunctive and Counterfactual Conditionals," *Analysis*, Vol. 12, No. 2 (Dec., 1951), 35-8.

[11] Cf. Wykstra's "CORNEA, Carnap, and Current Closure Befuddlement," 89, fn. 12: "…I originally meant the subjunctive formulation [in CORNEA] to express probabilistic implications of theories with Bayes' theorem in mind; the appropriateness of Lewis-Stalnaker semantics here needs scrutiny."

[12] William Rowe, "The Problem of Evil and Some Varieties of Atheism," *American Philosophical Quarterly*, 16 (1979), 335-41. Our representation here is meant to bring out the salient features relevant to a CORNEA defense, not to exegete the finer points of Rowe's argumentation.

And this appears-claim, he thinks, warrants the conclusion that:

(Rowe-3) *Probably,* there is no God-justifying good served by this suffering.[13]

We construe Rowe-2 as claiming that the noseeum feature of suffering is *prima facie* evidence against theism. And the step from Rowe-2 to Rowe-3 is best construed as drawing upon the norms governing how things *epistemically* "appear" to one.[14] It is licensed by an epistemic principle of justification that Swinburne calls the principle of credulity: when something epistemically *appears* to be a certain way, then it reasonable for one—barring defeaters—to believe that it *is* that way. Thus, if one lacks any defeaters, the noseeum feature of suffering can, by itself, be sufficient to reasonably warrant a shift from agnosticism to atheism.

The nerve of Wykstra's critique was to argue that Rowe's argument fails in its very first step—in the move from Rowe-1 to Rowe-2. While the move from "we see no such-and-such" to "it appears that there is no such-and-such" may seem innocuous, it is legitimate in some cases but illegitimate in others. For example, suppose your doctor drops a hypodermic needle on the floor, picks it up, looks at it carefully, and proceeds to try to use it on your arm. When you protest that it may be contaminated, he reasons as follows:

(Needle-1) We can, try as we may, see no viruses on the needle.

(Needle-2) Hence, it appears that there are no viruses on the needle.

(Needle-3) So probably (barring defeaters), there are no viruses on the needle.

You will certainly think your doctor has gone wrong in getting to Needle-3. But should you think that he has, in this situation, gone wrong in the very move from Needle-1 to Needle-2? Friends of CORNEA think you should.[15] For if there *were* viruses on the needle, then given the nature of viruses and of human vision, failing to see them is precisely and obviously what you (and the doctor) should expect. For this reason, not seeing such viruses in no way entitles the doctor to claim that there *appear* to be no viruses on the needle. CORNEA simply generalizes this intuitive constraint:

(CORNEA-1) On the basis of cognized situation s, human H is entitled to claim "it appears that p" only if it is reasonable for H to believe that, given her cognitive faculties and the use she has made of them, if p were not the case, s would likely be different than it is in some way discernible to her.[16]

To apply CORNEA to Rowe's argument, we must ask whether there are, on balance, good reasons to think that if God were to exist, then God-justifying goods connected to observed cases of suffering would, by virtue of God's nature and our human limitations, often not be evident to us. If there are such reasons, and Rowe is made cognizant of them, then it is not (or cannot

---

[13] In "Problem of Evil" Rowe, instead of using the term "probably," tends to put this as 'it appears reasonable to believe that..." As explained in the next section, we take Rowe-3 to be leaving open the possibility that other parts of one's total evidence might defeat the strong *prima facie* evidence provided by the facts of evil. Rowe-3 thus has an implicit "barring defeaters" clause making it still a claim about the *prima facie* or in-itself weight of the adduced data.

[14] Cf. Wykstra, "The Humean Obstacle," 85ff.

[15] McBrayer links his reformulation of CORNEA (discussed below) to the contention that one should meet evidential arguments in the appears-idiom by resisting the move from the second claim (e.g., Rowe-2) to the third claim (e.g., Rowe-3), while not resisting the move from first claim (e.g. Rowe-1) to the second (Rowe-2). But this implies that in the needle case, one will not resist the doctor's claim that there appear to be no viruses on the needle. We take this as a *reductio ad absurdum* of McBrayer's contention.

[16] Wykstra, "The Humean Obstacle," 85.

remain) reasonable for him to believe *crux.* And if this is so, then, by CORNEA, Rowe is not entitled to his claim that there appears to be no God-justifying good served by the observed cases of suffering. Such a claim is as ill-founded as our doctor's claim that there appear to be no viruses on the needle.

CORNEA was initially formulated using the "appears" idiom because this was prominent in Rowe's own formulation of his argument. But Rowe later reformulates his argument without using the appears-idiom. CORNEA is easily reformulated to follow suit, and thereby assumes a more general and powerful form.[17] It can thus be reformulated as:

> (CORNEA-2) For person P in a certain cognitive situation S, P is entitled to claim that new evidence E is levering evidence for H only if it is reasonable for P to believe that (*crux*) if H were false, E would, in the situation S, likely be different.

CORNEA-2 places the same condition on claims regarding levering evidence as CORNEA-1 places on epistemic appears-claims. While the two sorts of claims have identical import, it is preferable to work with this more general levering-evidence version of CORNEA. In what follows, the term CORNEA shall thus refer to CORNEA-2.

1.2. WHAT CORNEA CONSTRAINS: LEVERING EVIDENCE AND THE BY/ON DISTINCTION

CORNEA proposes a constraint on "Rowean" evidence—the type of confirming evidence that Rowe-style arguments purport to provide against theism. We call such evidence "levering evidence." As we use this term, evidence E qualifies as levering evidence for a hypothesis H only if E has the following three features.

First, to be levering evidence, E must consist of some alleged fact(s) or input such that acquiring or getting this input *properly changes*[18] the probability or credibility of hypothesis H from what it was prior to, or apart from, getting E.[19] Levering evidence is thus—to adopt the terminology of an earlier paper[20]—"dynamic." That is, it is *by* the addition of evidence, that the probability of a hypothesis is levered, or shifted, to a new value. Levering evidence, and dynamic evidence more generally, must be carefully distinguished from "static" evidence. E is static confirming evidence for H only if the probability of H, *on* E, is above some relevant threshold (e.g. .5). A hypothesis may be *made* more probable *by* some fact, even if the hypothesis *has* (statically) a low probability *on* that fact.[21]

---

[17] For a review of the dialectic, see Rowe "Friendly Atheism, Skeptical Theism, and the Problem of Evil," *International Journal for Philosophy of Religion*, (2006) 59: 79-92; Wykstra "Rowe's Noseeum Arguments from Evil." Wykstra discusses the special and general formulations in Wykstra "Suffering, Evidence, and Analogy: Noseeum Arguments Versus Skeptical Gambits," in *Philosophy through Science Fiction,* eds. Ryan Nicholas, et. al, (New York: Routledge, 2008).

[18] Here the normative "proper" carries the idea that although E may not actually change the person's belief regarding H (due to some cognitive defect, say), E is (epistemically) able to do so: it has what it takes to do so.

[19] Here the "prior to *or* apart from" allows that the "new" alleged facts may be either newly acquired facts or more familiar facts whose relation to the hypothesis is being newly considered or re-considered. (Here and elsewhere comments by Paul Draper were helpful.)

[20] "CORNEA, Carnap, and Current Closure Befuddlement," 91.

[21] On the By/On distinction see Wykstra "CORNEA, Carnap," 91. There is a further distinction to be made here between the "isolated" static probability of H on E alone, and its "holistic" static probability on E plus one's other background information. We here have in mind the holistic sense.

For example—to adapt an example from Elliott Sober[22]—suppose that I am playing poker and want to know if the card I am about to be dealt will be the Jack of Hearts. The dealer is careless, and I see that the next card is red. The fact that the card is red "confirms" in a dynamic sense that the card is the Jack of Hearts, potentially doubling its probability—from 1 in 52 (about .02) to 1 in 26 (about .04). Nevertheless, the hypothesis that the next card is the Jack of Hearts is still very improbable *on* the fact that the next card is red. So in a static sense, the fact that the card is red does not "confirm" the hypothesis that the next card is the Jack of Hearts, since the probability is below the relevant threshold of (say) .5.

The second feature of levering evidence E is that E must properly change the probability or credibility of the hypothesis *sharply*. To give this content, we follow Wykstra's earlier article[23] in distinguishing between three "square" doxastic states—namely, square belief, square *dis*belief, and square *non*-belief (that is, *squarely suspended* or *withheld* belief). We also adopt his simplifying assumption[24] that "square belief" can be associated with giving a proposition a probability (or degree of credibility) fairly close to 1 (say, .99 or higher), "square disbelief" with giving the proposition a probability close to 0 (say, .01 or below), and "square non- belief" with giving the proposition a probability of around .5.[25] We then call a doxastic change a "sharp" change if it is from one *square* doxastic state to another—in particular, from square *non*-belief to square belief.[26] By "levering evidence" we thus mean evidence for or against a hypothesis that is of sufficient strength to shift the rational credibility of a hypothesis from one square state to another. We formulate *crux* as a condition on levering evidence *for* a hypothesis. It is also a condition on levering evidence *against* a hypothesis, since to claim E is levering evidence against a hypothesis is to claim that E is levering evidence *for* the denial of that hypothesis.[27]

Third, we intend this notion of levering evidence—in keeping with Rowe-style arguments—to capture what might be called the *prima facie* (rather than *'ultima facie'*) bearing of E on H.[28]

---

[22] Sober gives a similar example in *Did Darwin Write the Origin Backwards: Philosophical Essays on Darwin's Theory* (New York: Prometheus Books, 2011), 145-6. Sober uses his example to target "the special consequence condition of confirmation" as formulated by Hempel in his classic 1945 paper "Studies in the Logic of Confirmation" *Mind*, 54. 1-26 and 97-121. It should be noted, however, that under Carnap's tutoring, Hempel revised his early ambiguous formulation, and that the improved principle—now restricted to static confirmation—is correct. See Hempel's *Aspects of Scientific Explanation* (New York: MacMillan, 1965), especially "Studies in the Logic of Confirmation," fn. 40 and "Postscript (1964) On Confirmation," 49-50.

[23] Wykstra, "Rowe's Noseeum Arguments from Evil," 130-2.

[24] This is, we stress, a simplifying assumption in which the vague term "associate" is used advisedly. Given the vagaries of the concept of belief, it may be preferable to speak of squarely "accepting" a proposition, for *identifying* belief with being above any probabilistic threshold leads quickly into familiar lottery paradoxes. We think our results will be sustainable however these tricky issues are eventually negotiated. (Comments by Glenn Ross were helpful to us here.)

[25] We emphasize here that the values we are assigning are elucidatory; that square beliefs, along with square non-belief, represent more of a range of values.

[26] Levering in this way is by no means an unrealistically high bar for evidence to meet. The simplest testimonial evidence is, for example, able to exceed this bar without difficulty. A colleague informing me that he has three children properly levers me to believing that he has three children, even though I was in a state of square agnosticism prior to that.

[27] We thank an anonymous referee for urging more clarity on this point.

[28] Thus Wykstra, "The Humean Obstacle,": "[Rowe's] overall claim… is that the evidence of suffering supports atheism in what we might call a "qualifiedly strong" sense, viz., it strongly supports atheism *provided* that independent assessments of theistic arguments show these to be as weak as Rowe, based on his independent study of them, gives us his word they are" (79).

Here we may think of ourselves as like Blind Lady Justice. *Eventually,* she wants to be in a position to weigh all the evidence for and against a particular hypothesis in the two pans of her scales. But *initially,* she wants to determine how weighty some individual pieces of evidence are. She thus starts from some squarely neutral or "even" position, and asks how much some item of interest can tilt her balance from that even position. An answer to this question—the one, we think, Rowe means to be addressing—leaves open the further question of how the scales will tilt when any opposing evidence is placed on the *other* pan.

1.3. THE *REAL* CORE OF CORNEA: CORE VS. MCCORNEA

CORNEA posits a constraint on when someone is *entitled to claim,* in some cognized situation, that certain data or input is levering evidence for some hypothesis.[29] This constraint is 'internalist,' for it requires that it be, for this person, *reasonable to believe* that the following test condition is satisfied: if H were false, then E would likely be different. CORNEA's constraint is thus doubly internalist: it is a constraint on when a person is *entitled to* make a levering-evidence claim, and the constraint itself is that it be *reasonable for that person to believe*[30] a particular conditional claim (namely, *crux*).

As McBrayer notes, Wykstra has recently suggested that "behind CORNEA" is a deeper idea:

> the key idea behind CORNEA is a proposed test for whether some alleged evidence seriously 'supports' —in a sense to be clarified presently— some hypothesis H. [The test is this: ask whether, if H were false, E would likely be different. If the answer is "No," then E can't seriously support H.][31]

McBrayer, picking up on this suggestion, urges that CORNEA should itself be seen as "a restriction on when any 'cognized situation' counts as evidence" for a hypothesis. He offers his own rendition of CORNEA, which we will call McCORNEA:

> McCORNEA: A subject S's cognitive situation C is evidence for P *only if* it is reasonable for S to believe that were P false then C would [likely] be discernibly different.[32]

And he supports this rendering of CORNEA by selectively quoting Wykstra: "Wykstra concedes as much: "the key idea behind CORNEA is a proposed test for whether some alleged evidence seriously 'supports' . . . some hypothesis H."[33]

But McCORNEA has three problems. First, whereas Wykstra's "key idea" refers to a constraint on evidence that *seriously supports* a hypothesis, McCORNEA makes this a constraint on evidence *simpliciter*.[34] Second, whereas McBrayer intends to be offering a formulation of CORNEA, Wykstra is not. He is attempting to articulate the "key idea *behind* CORNEA." And

[29] Or "entitled to regard" it as levering evidence.

[30] In the sense used here, it can be reasonable for a person to believe a proposition even if he or she does not believe this in any conscious way; see Wykstra's dispositional explication in "Humean Obstacle," 87.

[31] Wykstra, "CORNEA, Carnap," p. 88. McBrayer's quotation omits the parts we have here put in brackets, which will become important below.

[32] McBrayer, 81 (italics ours). McBrayer renders what follows "…reasonable to believe" as a counterfactual in David Lewis' notation. We have restored it to a grammatical subjunctive. McBrayer also omits the crucial "likely" that we have here put in brackets.

[33] McBrayer, 81, citing Wykstra's "CORNEA, Carnap," 88.

[34] Here by his ellipses, McBrayer is omitting Wykstra's "—in a sense to be clarified presently—". By this phrase Wykstra is alerting the reader that "seriously supports" will be explicated *via* his forthcoming definition of *levering* evidence.

this key idea concerns a constraint on when some datum *is* levering evidence—to be distinguished from CORNEA's constraint on when the datum can be *reasonably regarded* as levering evidence. Third, whereas on McCORNEA, the constraint itself is internalist (that it be *reasonable for a person to believe* the *crux* subjunctive), on Wykstra's suggestion, the constraint is an externalist one (that *crux* be *true).*

The key idea that McBrayer imperfectly appropriates, then, is this: that "behind" the internalist CORNEA is a deeper externalist principle, one that lends plausibility and support to CORNEA. And just as CORNEA is doubly internalist, so this deeper principle is doubly externalist. It is externalist, first, in *what* it is a constraint *on*—namely, on when something *is* evidence of the "seriously supporting" (i.e., levering) type. It is externalist, second, in what the constraint *is*—namely, that it *be true that*, if H were false, then E would likely be different.[35] The deeper CORE principle, put side by side with CORNEA, is then as follows:

CORNEA	CORE
For person P in cognitive situation S, P is *entitled to claim* that new evidence E is levering evidence for H only if it is *reasonable for P to believe* that (*crux*) if H were false, E would, in the situation, likely be different.	In cognitive situation S giving person P new input E, E *is* levering evidence for hypothesis H only if it is *true* that *(crux)* if H were false, E would likely be different.[36]

1.4. UNDERSTANDING CORNEA'S CONDITIONAL

We've seen that both CORE and CORNEA have, at their heart, the conditional proposition we've dubbed *crux*:

> *crux*: if H were false, E would likely be different.

The *crux* proposition is a grammatical subjunctive, a conditional in the subjunctive mood. Due to this grammatical feature, McBrayer takes it as expressing a logical subjunctive, or "counterfactual conditional." Consequently he evaluates it using the possible-world semantics pioneered by Robert Stalnaker and David Lewis. But here caution is in order. As Stalnaker himself notes, grammatical subjunctives sometimes express, not logical subjunctives, but instead "non-counterfactual subjunctives." His example of a non-counterfactual subjunctive is:

> If the butler had done it, we would have found just [i.e. exactly] the evidence we did find.[37]

Stalnaker envisions this sentence as uttered in a context where the speaker—perhaps a detective—is arguing that the evidence confirms that the butler *did* do it. In calling this subjunctive "non-counterfactual," Stalnaker does not simply mean that its antecedent is not

---

[35] In addition to the well-known constraint on *knowledge* found in Nozick's *Philosophical Explanations* (172ff.), there is also a little-discussed section (248ff.) in which Nozick also defends a sensitivity constraint on *evidence.* He proposes that in order for E to be evidence for a hypothesis H, it must be true that "if H weren't true, then E wouldn't hold." Nozick's constraint is distinct from CORE in at least the following way: it has a wider domain—over any evidence—whereas CORE is a constraint only on levering evidence.

[36] A slightly different—but not equivalent—formulation is "In cognitive situation S giving new input E, E is levering evidence against hypothesis H only if, were H true, E would not likely be the same."

[37] Stalnaker, "Indicative Conditionals," 146. The striking affinity of CORNEA with Stalnaker's example needs no comment. The same affinity is found with Alan Ross Anderson's examples such as "If Jones had taken arsenic, he would have shown just exactly those symptoms which he does in fact show." See Anderson's "A Note on Subjunctive and Counterfactual Conditionals," 37.

contrary-to-fact. He means more deeply that such grammatical subjunctives should not be evaluated using a Lewis-Stalnaker closest-possible-worlds semantic.[38]

We agree with Stalnaker that some grammatical subjunctives are not best understood as logical subjunctives. Further, his example strikingly illustrates the natural use of grammatical subjunctives to express the evidential implications so vital to hypothesis testing. It is natural to say, for example: "If Huygen's vortex theory of gravity were true, it is entirely expectable that the planets would revolve around the sun in the same direction; whereas if Newton's theory were true, this phenomenon would be rather improbable."

As we see it, the most promising way to think about these evidential implications is as conditional probabilities. The examples are then saying that the probability of finding just the evidence that was found, conditional on the hypothesis that the butler did do it (and our other background information), is high, and that the conditional probability of the planets revolving around the sun in the same direction is high on Huygens's theory, but not on Newton's theory. We thus suggest that, in the hypothesis-testing context, grammatical subjunctives are often best understood as expressing evidential implications of hypotheses, and that such implications are best understood as conditional probabilities.

To be sure, understanding grammatical subjunctives in this way also raises some vexing issues about whether the formalism of the probability calculus—and the formal notion of conditional probability that is defined within it—has any meaningful *general* application to the arena of hypothesis testing. Some thinkers—Richard Von Mises, for example—argue that it is wrong-headed to think that the probability calculus has any relevance to evaluating historical evidence for hypotheses about *unique* events or processes (say, the hypothesis that Kennedy was killed by Lee Harvey Oswald, or that Jesus of Nazareth rose from the dead). The probability calculus and formal notion of conditional probability, as von Mises sees it, has relatively narrow application—it applies only, he thinks, to "mass phenomena" involving events of an indefinitely repeatable type, allowing us to assign long-term ratios of favorable outcomes within an indefinitely large class.

Others, like Rudolf Carnap, think that the probability calculus and notion of conditional probability has broader application. On Carnap's view, for example, the probability relation can be modeled by logical relations of "partial entailment" between propositions, even where these involve no mass phenomena or discrete repeatable events to which frequency-ratios can be assigned. Such partial-entailment relations, Carnap thinks, bear intimately on the rational credibility of hypotheses. Working within this Carnapian tradition, Richard Swinburne thinks that even if there is only one unique universe, we can meaningful speak of the conditional probability of (say) the phenomenon of beauty on the hypothesis that God created that universe, and compare this with conditional probability of beauty on a naturalistic hypothesis. And even if we can only attach comparative estimates to these conditional probabilities, Swinburne thinks we can properly and usefully employ the probability calculus to help us discern, in a comparative

[38] For Stalnaker's argument, see "Indicative Conditionals." A full discussion of Stalnaker's treatment of conditionals is beyond the scope of this paper. The interested reader should consult his "A Theory of Conditionals" and "Indicative Conditionals," both in Jackson, *Conditionals*.

though admittedly non-quantitative way, how the probability—interpreted as "rational credibility"—of each hypothesis changes as a widening range of evidence is taken into account.[39]

While our sympathies lie with these broader applications, we concede to von Mises that in these broader applications, it is often not easy to say what relationship between propositions is being mapped onto the conditional probability relation. Happily, however, McBrayer's chief counterexample to CORNEA involves a lottery scenario—just the sort of case involving those limit frequencies that, as von Mises helps us see, affords the best traction for the conditional probability relation within the probability calculus.

Taking *crux* in this way, we'll understand "If H were false, E would likely be different" as saying that the conditional probability of E, given the falsity of H, is low—at least below .5. Put in standard notation, this says that P(E|~H) < .5. So instead of approaching *crux* by using a possible world semantics, we use the probability calculus as our tool of choice, making use of Bayes' theorem and weighted averages. Taken simply as a theorem in the probability calculus, Bayes' theorem is usually put as:

$$P(\mathrm{A} \mid \mathrm{B}) = P(A) * \left[\frac{P(B \mid \mathrm{A})}{P(B)}\right]$$

But as Swinburne and Wesley Salmon have taught us,[40] this formula can be applied to the relationship between a hypothesis H and alleged confirming evidence E. On this application, we take the "absolute" probabilities P(A) and P(B) as P(H|k) and P(E|k)—that is, as the background probabilities of H and E on some relevant background information k. Bayes' theorem then tells us that P(H | E & k)—the new probability of H when E is added to k—is as follows:

$$P(H \mid E \,\&\, k) = P(H \mid k) * \left[\frac{P(E \mid H \,\&\, k)}{P(\mathrm{E} \mid k)}\right]$$

As we'll see, Bayes' theorem will be a powerful tool for applying CORNEA and CORE to McBrayer's chief counterexample scenario.

2. MCBRAYER'S LOTTO ARGUMENT

Having clarified CORNEA and CORE, we now consider how each principle fares against McBrayer's counterexample. His paper puts the most weight on his first "lottery" counterexample, so we shall do so as well. We shall, following McBrayer's lead, finesse his initial scenario so as to make his counterexample applicable to the more mature and sober versions of CORNEA (and CORE).

McBrayer initially formulates his lottery argument as follows:

> Though I hold a ticket, I believe that I will lose the lottery. I have inductive evidence for this claim. I know that the odds of winning are one in a million. Is my evidence sensitive to the fact that I will lose the lottery? Go to the closest world in which I win. I just get lucky and pull the right ticket. Is it reasonable to

---

[39] See, for instance, Swinburne, The Existence of God, (Oxford: Oxford University Press, 2nd. ed. 2004), 14-9; and his "Introduction" to Bayes's Theorem, ed. Richard Swinburne, Proceedings of the British Academy, 113, (Oxford: Oxford University Press, 2002), 1-20.

[40] For Swinburne, see the aforementioned works. For Salmon see, *inter alia*, his "Bayes' Theorem and the History of Science," in *Minnesota Studies in the Philosophy of Science*, vol. 5 *Historical and Philosophical Perspectives of Science*, ed. Roger H. Stuewer (Minneapolis: University of Minnesota Press, 1970), 68-86, and "Rationality and Objectivity in Science, or Tom Kuhn Meets Tom Bayes," in *Minnesota Studies in the Philosophy of Science*, vol. 14, ed. C. Wade Savage (Minneapolis: University of Minnesota Press, 1990), 175-204.

> believe that my cognitive situation in the actual world would be discernibly different from my cognitive situation in the possible world in which I win? No—things would look just the same to me. So, my cognized situation in this case is not evidence for the claim that I will lose the lottery.

Let's here refer to the ticket-holder as "Holt," so as to keep his beliefs distinct from McBrayer's claims about him. McBrayer portrays Holt as both knowing that

(1) The odds of (Holt) winning are one in a million.

and as taking (1) to be "inductive evidence for"

(2) Holt will lose the lottery.

Our intuition, McBrayer thinks, is that in this scenario, (1) is inductive evidence for (2). But is CORNEA consistent with this intuition? It is consistent, on McBrayer's line of thought,[41] only if it reasonable for Holt to believe the crucial conditional:

crux: "If Holt were to win the lottery, then Holt's evidence—(1)—would [likely] be discernibly different."[42]

McBrayer thinks that *crux* is neither true nor reasonable for him [or Holt] to believe. For in the closest world where Holt wins the lottery—the world where, by luck, he has simply drawn the winning ticket—(1) isn't any different. Things "would look just the same." And that things would look the same is, McBrayer says, something that "I [and Holt] know full well." Thus, it is not reasonable for him [or Holt] to believe *crux*. In this way, neither CORNEA nor CORE seems consistent with our best epistemic intuitions.

2.1. First Improvements

McBrayer, alerted by an anonymous referee, goes on to improve this scenario in several important ways. He notes—in accord with our own exegesis above—that in its later and more mature formulations, CORNEA is formulated as a constraint not on "evidence *simpliciter*" but on a specific sort of evidence, namely, "dynamic" evidence of "levering" strength.[43] But in his scenario as currently specified, proposition (1) does not clearly function as levering evidence. For the scenario does not describe Holt as being in some initial doxastic state, and then shifting to some new doxastic state; nor does it describe the shift so as to make clear that it is a proper shift from square non-belief to square belief; nor does it specify some specific cognitive input that effects this shift. The counterexample thus, he notes, might be dismissed as attacking only a straw man.

But McBrayer argues that his scenario is easily improved so as to make Holt's evidence a clear case of levering evidence—while still clearly failing CORNEA's test. Here is his finesse of the scenario to that end:

---

[41] McBrayer counterexample, while directed against the misbegotten McCORNEA, can without loss (or gain) be redirected against CORE, and in expressing McBrayer's "line of thought," we do just this.

[42] McBrayer, 83.

[43] McBrayer suggests that that CORNEA was originally meant to be a constraint on "evidence *simpliciter*," and that only later was its restricted to dynamic levering evidence. But the distinction between static and dynamic evidence, and between weak and strong dynamic evidence, was central to the original exposition of Rowe's case. And while Wykstra claims in "The Humean Obstacle" that Rowe's evidence isn't "even weak evidence" for atheism, he later qualifies this as over-reaching, and it is not clear that the over-reaching version is built into his earliest formulation of CORNEA.

> [Holt] is given a lottery ticket in ignorance of how many tickets are sold. Perhaps Holt has the only ticket, or perhaps there are a million tickets. Being rational, Holt withholds belief concerning the proposition that he will win the lottery. Later Holt learns that the odds of winning are one in a million. Based on this new information, Holt disbelieves that he will win the lottery. [Holt's] cognitive situation in this case warrants a revision from non-belief to disbelief. The evidence is therefore levering evidence.[44]

2.2. SECOND IMPROVEMENTS

The added details in McBrayer's finesse go *some* distance toward giving a situation in which there is a shift from square non-belief to square belief (or square disbelief),[45] and thus a case of levering evidence. But they do not, we think, give enough detail to specify a *specific cognitive input* that is affecting this shift.[46] As stated, the scenario still specifies the basis of the shift as Holt's "knowing that"

(1) The odds of (Holt) winning are one in a million.

But this is entirely too vague. It gives no concrete depiction of what Holt's specific new "cognitive input situation" is supposed to be.[47] The belief (or "knowledge") that (1), after all, normally rests on a very complex bramble of considerations put in place over a long period of time. One of these is the information Holt has as to the size of the lottery; but in grounding (1), this works in tandem with many other considerations that are normally part of one's background beliefs—e.g., that the lottery tickets are distributed by a humanly fair non-rigged process, that no angels or other supernatural beings are giving one special dibs on winning lotteries of this sort, and so forth.

We will improve the scenario by specifying that one of these considerations is a new cognitive input situation, so that it is intuitively clear that this input will, when added to other normal background considerations already in place, significantly raise the probability or degree of confidence of (1). The simplest way to do this is to suppose that what Holt receives, as new input, is reliable information about the *size* of the lottery. Our improvements thus extend McBrayer's finesse by adding to the improved scenario that Holt gets specific new input from a reliable source that (1') the lottery is a million ticket lottery, and that prior to learning this, Holt has in place other normal background beliefs—of the sort just mentioned—so that this new input does greatly increase his degree of confidence in (2).

2.3. THIRD IMPROVEMENTS

Despite the important finesses made so far, the improved LOTTO scenario is still flawed. The scenario is meant to be one in which Holt's cognitive situation clearly effects and warrants a shift from square non-belief to square belief. But does such a shift occur in the scenario as described so far? Notice, in particular, Holt's initial doxastic state:

---

[44] McBrayer, 85.

[45] Square disbelief that he has a winning ticket is of course the same as square belief that he has a losing ticket.

[46] CORNEA was from the start expressly formulated as a constraint on how belief should be dynamically altered by the addition of *specific new cognitive* input. Here see Wykstra's use of the example of Tom and his worm-sandwich in "The Humean Obstacle," 80.

[47] And this leaves us equally at sea about what it would take, concretely, for Holt's situation to be "discernibly different."

> [Holt] is given a lottery ticket in ignorance of how many tickets are sold. Perhaps Holt has the only ticket, or perhaps there are a million tickets. Being rational, Holt withholds [i.e., suspends] belief concerning the proposition that he will win the lottery.

Here we can read the word "perhaps" in two very different ways. On the most natural reading, the two possibilities mentioned are *illustrative*. Holt knows he holds a lottery ticket, but he has no idea how large the lottery is: perhaps it is from one-ticket lottery, or perhaps from a two ticket lottery, or perhaps a three-ticket one… up to (let's say) a one-million ticket lottery.[48] But read in this way, the story does not give a scenario in which Holt is in an initial state of square non-belief, that is, of squarely suspended belief. Such a state is, we saw, a confidence level associated with a probability of around .5. But on the present reading, Holt's situation is tantamount to one in which it is equally probable that his ticket is from a one-ticket lottery, or that it is from a two-ticket lottery, ...up to a million ticket lottery. A little calculation shows that in this situation, the probability of holding a winning ticket is nowhere near .5: instead it is something like 1 in 70,000, making the probability of holding losing ticket something like 69,999 in 70,000, or .999 986.[49] Holt's initial situation thus does not warrant a state of square non-belief at all; what is initially warranted is instead a confidence level associated with a .999986 probability. And his new evidence, instead of levering him from a doxastic state of around .5 to something well over .99, nudges him from a .999986 doxastic "square belief" state to a .999999 doxastic "square belief" state.

To remedy this problem, we shall re-structure the scenario so that Holt's initial situation is one in which the two possibilities mentioned by McBrayer (i.e., "perhaps I have the only ticket, or perhaps there are a million others") are not illustrative, but instead exhaustive. Holt, let us suppose, knows that his ticket is from one of two lotteries—either from a million-ticket lottery, or from a single ticket lottery. The scenario does now put Holt in an initial state of squarely

---

48 On this reading, the scenario puts Holt in a situation of *modest* ignorance: while ignorant about the *actual* size of the lottery, he has some idea of the distribution of probabilities within the range of possible sizes. This is tantamount to Holt's having a ticket that he knows has been randomly chosen from a "distribution" of a million distinct lotteries, one of which is a one-ticket lottery, the next a two-ticket lottery, etc., up to a million-ticket lottery. But if Holt's situation is one of *radical* ignorance, where he has no idea of the distribution in a possible range, it might plausibly be claimed (as Harry Plantinga and Glen Ross have noted) that Holt cannot meaningfully associate any probability—neither high, middle, or low—to having in hand a winning ticket: he must see it as being indeterminately probable, or "aprobable." In this event the probability calculus seems to no longer apply, although we think it plausible some analogue of the CORNEA constraint should still apply to evidence that levers him out of that "aprobable" state into one of square belief or disbelief. We here confine ourselves to the more tractable "modest ignorance" construal of the scenario, hoping to return to this more difficult case at a later date. We thank Mike Bergmann for pressing us on this point.

49 It's easiest to see this if we imagine Holt knowing there are one million possible lotteries, ranging incrementally from a one-ticket lottery to a one-million (or two-million, etc.) ticket lottery. The probability that the ticket he holds is a winning ticket would simply be a summation, for each of these lotteries, of the probability that a ticket is a winning ticket given that it is from that lottery, multiplied by the probability that it is from that lottery. Thus, for the first 1-ticket lottery, we have a probability of 1 (there is a probability of 1 that he has the winning ticket given that it is from that lottery) times one-millionth or $10^{-}6$ (the probability that it is from that lottery). The weighted summation is $(1/1 * 10^{-6}) + (1/2 * 10^{-6}) + (1/3 * 10^{-6})$ ...+ $(1/10^{-6}) * 10^{-6}$). This is equal to $[1/1 + 1/2 + 1/3 + \ldots 1/10^{-6}] * 10^{-6}$. The series in brackets, $[1/1 + 1/2 + 1/3 + 1/10^{-6}]$, is equal to about 14.39. The weighted summation is thus $14.39 / 10^{6}$, or roughly 14 in a million, which is roughly 1 in 70,000—pretty long odds. We thank Harry Plantinga in Calvin's Computer Science Department for grinding out the 14.39 using a quickly written and executed Pascal algorithm.

suspended belief about whether he has a losing or winning ticket. And consequently, his new evidence—that his ticket is from a million ticket lottery—does now dynamically shift Holt from square non-belief to square belief that the ticket he holds is a losing ticket.[50] Here, then, is our triply-improved version of McBrayer's LOTTO scenario, with some details added for ease of reference:

> Holt has purchased two tickets, one for a one-ticket raffle at a Dutch church picnic (with a used Psalter hymnal prize), and the other from a million-ticket Catholic raffle. He stores them in his room. He then finds one ticket missing—his thieving brother Klep slipped into the room and stole one of them. Holt knows this much, but he doesn't know which ticket Klep stole. In this initial evidential situation, the odds that he holds the Dutch ticket are thus 50/50. Holt's initial state is one of squarely suspended-belief about whether he holds the losing (or winning) ticket.
>
> But Holt now gets new input: Klep, in his usual compulsively honest way, breaks down and confesses that the ticket he stole is the Dutch ticket. Getting this new information properly levers Holt from square non-belief into square belief that the ticket he holds is a losing (because Catholic) ticket. Klep's testimony thus qualifies as levering evidence.

We think this triply-improved scenario puts McBrayer's scenario in its most formidable form. Now Holt's new cognitive input does seem to function as levering evidence: it shifts Holt—properly, it seems—from squarely suspended belief to square belief that he holds a losing ticket. At the same time, this evidence does not seem to satisfy the constraints that CORNEA and CORE—using *crux*—put on levering evidence. And it does not seem to for the very reasons that McBrayer—adopting a counterfactual reading of *crux*—gives: in the closest possible world where Holt—by luck—has a winning ticket, his new inductive evidence (Klep's testimony to stealing the Dutch ticket) is exactly the same. Hence, if Holt were to have the winning ticket, his evidence would *not* be different—instead, it would be the same. So it seems that Holt's evidence *fails* both the CORNEA and CORE tests.

With the triply-improved scenario in mind, we thus put McBrayer's argument as follows:

(Lotto-1) If CORE [CORNEA] is right, then in the improved lottery scenario, Holt does *not* have [and is *not* entitled to claim that he has] levering evidence that his ticket is a losing ticket.[51]

(Lotto-2) But in the improved lottery scenario, Holt *has* [and *is* entitled to claim he has] levering evidence that his ticket is a losing ticking.

(Lotto-3) So CORE [CORNEA] is false.

3. CORE AS CONDITIONAL PROBABILITY: WHAT BAYES' THEOREM REQUIRES

According to CORE, E is levering evidence for H only if *crux*—if H were false, then E would likely be different—is true. In his LOTTO counterexample, McBrayer interprets *crux* as a logical

---

[50] It's worth noting that if effects this shift by a two-staged process: first it shifts his belief about the odds of any given ticket-holder's winning the lottery; this in turn shifts him to a new degree of belief that his own ticket is a loser.

[51] It is of course possible that Holt's evidence fails to be levering evidence because it fails some *other* necessary constraint on levering evidence. So we here mean Lotto-1 to be elliptical for something like 'if CORE and CORNEA are true, then Holt's evidence isn't levering evidence *because it fails their tests*.'

subjunctive, using Lewis-Stalnaker semantics to evaluate it. But as noted above, *crux* can also be interpreted as a conditional probability. Putting these two interpretations[52] of CORE side by side:

<table>
<tr><td colspan="2">(CORE) In cognitive situation S giving new input E, E is levering evidence for H only if it is the case that: (crux) If H were false, then E would likely be different.</td></tr>
<tr><td>(C-CORE) In cognitive situation S giving new input E, E is levering evidence for H only if it is the case that: (c-crux) in the closest possible world(s) in which not-H is true, E is [likely] not true.</td><td>(P-CORE) In cognitive situation S giving new input E, E is levering evidence for H only if it is the case that: (p-crux) the conditional probability of E on not-H— viz, P (E | not-H & k)—is below .5.</td></tr>
</table>

Now McBrayer cases threaten C-CORE (and CORNEA) because in them, E seems to be levering evidence even though the *c-crux* requirement evidently fails to be satisfied. We will argue here that interpreting CORE as P-CORE deflects this threat, and puts us in a position to see why, in cases like these, counterfactual conditionals behave differently than conditional probabilities. In Section 3, we show that Bayes' theorem itself entails—in perfect accord with P-CORE—that *p-crux* must be satisfied if E is to be levering evidence for H. In Section 4, we then use an expanded form of Bayes' theorem to show that—and, more importantly, *why*—*p-crux* is indeed *true* in McBrayer's LOTTO scenario, so that P-CORE is satisfied in this scenario. (We also show that on this treatment, *p-crux* is *reasonable to believe* in this scenario, so that the probabilistic version of CORNEA—P-CORNEA, as it were—is also satisfied.)

Since our focus in Sections 3 and 4 will be entirely on the probabilistic versions of CORE, and *crux*, these sections will use the terms "CORNEA" and CORE" and "crux" to refer to their probabilistic P-versions.

3.1. DICTIONARY OF ABBREVIATIONS

When we take *crux* as a conditional probability, CORE says that E is levering evidence for H only if, in the situation, we wouldn't likely get evidence E on not-H—that is, only if it the conditional probability of E on not-H is below .5. Here, then, is the probabilistic rendering of CORE:

> (P-CORE) In cognitive situation S giving new input E, E is levering evidence for H only if it is the case that: the conditional probability of E on not-H— *viz*, P (E | not-H & k)—is below .5.

To see how P-CORE handles McBrayer cases—and begin to see why it works differently from C-CORE—we view the question through the lens of a standard application of Bayes' theorem:

$$P\,(H \mid E \,\&\, k) \;=\; P\,(H \mid k) \quad * \quad \left[\frac{P\,(E \mid H \,\&\, k)}{P\,(\mathrm{E} \mid k)}\right]$$

NEW PROB of H = OLD PROB OF H * "THE *QUOTIENT"*

[52] It's a bit imprecise to call P-CORE an "interpretation" of CORE. C-CORE does use standard semantics so as to "interpret" CORE, but P-CORE merely invokes the notion of conditional probability, without offering any interpretation (logical, frequentist, propensity, etc) of the term "probable." We pursue this more in Wykstra and Perrine "Foundations of Skeptical Theism: On the (Un)importance of Being Sensitive" in a forthcoming volume *Skeptical Theism: New Essays* (Oxford: Oxford University Press), edited by Trent Dougherty and Justin McBrayer.

On this application, Bayes' theorem tells us that the "new probability" of a hypothesis—namely, P(H| E & k)—is equal to the "old probability" of H on k alone—namely, P (H | k)—times a special fraction that John Maynard Keynes called "the relevance quotient," and that we shall just call "The Quotient." In it, the numerator P(E | H & k) represents how probable, or "expectable", the evidence E is, on the assumption that hypothesis H (together with our background information k) is true. The denominator P(E | k) represents how likely the evidence E is, merely on the assumption of our background knowledge k by itself. We shall refer to these as, respectively, the "hypothetical expectability" of the new evidence E (its likelihood assuming the hypothesis is true) and the "background expectability" of E (its likelihood on the mere assumption of the background information alone).

In applying Bayes' theorem to the LOTTO scenario, we will use the following abbreviations:

k = Holt's background information. This includes (but is not exhausted by) the following salient points: that he bought two tickets for lotteries or raffles as described above, that his brother Klep stole one of these tickets, leaving the other in his possession; that Klep, while a kleptomaniac, is compulsively honest in truthfully confessing his misdeeds, and so on.

E = Holt's new evidence (Klep's confession that the ticket he stole is the Dutch ticket).

H_L = the hypothesis that the ticket Holt holds is a losing ticket.

H_W = the hypothesis that the ticket Holt holds is a winning ticket.

In the next section, we shall also have occasion to use two further abbreviations:

H_{Wd} = the hypothesis that the ticket Holt holds is a winning ticket from the Dutch raffle.

H_{Wc} = the hypothesis that the ticket Holt holds is a winning ticket from the Catholic lottery.

CORE says that Holt's new evidence E counts as levering evidence for H_L only if: were not-H_L true, then E would likely be different. This constraint, interpreted as a conditional probability, requires that P (E| ~H_L & k) be low. Now ~H_L is, in the scenario, the same as H_W. So the constraint just requires that P (E| H_W & k) be low.

By working backwards from the fact that E *is* levering evidence for H_L in the LOTTO scenario (a point on which we and McBrayer agree), we now aim to see what requirement *Bayes' theorem* puts on this levering evidence, so as to compare it with the CORE requirement.

3.2. THE BAYESIAN REQUIREMENT: STEP ONE

We begin, then, by agreeing with McBrayer that in Holt's situation, E most definitely *is* levering evidence for $H_{L.}$ In the improved scenario, Holt is reasonably shifted by his new evidence from square agnosticism to square belief. Holt begins in a state of squarely suspended belief, assigning to P(H | k) a probability of .5; his new evidence—being levering evidence—then properly boosts this to a probability of .99 or better.[53]

Plugging these into Bayes' theorem, we get

$$.99 = .5 * \left[\frac{P(\mathrm{E}|\mathrm{H_L} \ \& \ k)}{P\ (\mathrm{E}|k)}\right]$$

Working backwards, it is evident by inspection that in this scenario, E's being levering evidence for H requires The Quotient itself to be nearly 2.

[53] Again, we remind the reader that the values we assign here are not exact, nor need they be for present purposes.

3.3. THE BAYESIAN REQUIREMENT: STEP TWO

But what, in turn, does *this* require?

We can make one more step by noting that in the Lotto scenario, the numerator of The Quotient—namely, P (E| H_L & k)—is close to 1. H_L is, after all, the hypothesis that Holt holds a *losing* ticket, and our background information k includes the claim that the Dutch ticket is a sure *winner*. It follows that if Holt has a *losing* ticket, he must have a Catholic ticket. And Holt's having a Catholic ticket entails that Klep must have stolen the *Dutch* ticket. But given (from our background information) Klep's compulsive honesty, Klep stealing the Dutch ticket implies that his confession will be to stealing the Dutch ticket—which is just E. So, P (E| H_L & k) is 1.[54] But this means the Bayesian Formula now becomes

$$.99 \; = \; .5 * \left[\frac{1}{P\ (\mathrm{E}|k)}\right]$$

Working backwards one more step, it is evident by inspection that in this scenario, E's being levering evidence for H requires that The Quotient's *denominator*—namely, P (E|k)—must be about .5.

3.4. THE BAYESIAN REQUIREMENT: VOILA!

But what, in turn, does *this* require?

To answer this, we must look closely at what the denominator P(E|k) signifies. Earlier we referred to the denominator as the "background expectability" of E on the assumption of our background information—k—alone. We contrast the denominator with the numerator, which is the "hypothetical expectability" of E on the assumption that the hypothesis of interest—here H_L—is true. However, this contrast must not mislead us into thinking that P(E|k) is the probability of E on the hypothesis that H_L is *false*. Rather, the value of P(E|k) is an average of E's hypothetical expectability on *both* H_L and H_W, where those values are each corrected by a "weighting factor" of how likely, on their own, H_L and H_W are. That is, the value of P(E|k) in an expanded form is:

$$P(E|k) = \; P(E\ |\mathrm{H_L}\&\mathrm{k}) * P(\mathrm{H_L}|k) + P(\mathrm{E}|\mathrm{H_W}\,\&\mathrm{k}) * P(\mathrm{H_W}\ |k)$$

In this expanded form, there are five terms, and four of them have values determined by the scenario. We saw in our last step that in order for E to be levering evidence, P (E |k) must be about .5. We also know from that step, that $P(E|\mathrm{H_L}\&\ \mathrm{k})$ —the hypothetical expectability of E on H_L—is 1. We know, from our first step, that $P(\mathrm{H_L}|\mathrm{k})$ is .5. And we know that P(H_W | k) is also .5, since H_W is the denial of H_L. Plugging these four values in gives us:

$$.5 \qquad \cong \qquad [1\ *\ .5]\ +\ [\,\boldsymbol{P}(\mathbf{E}\,|\,\mathbf{Hw}\ \&\ \mathbf{k}))\quad *\ .5\,]$$

Working backwards one more step, it is evident by visual inspection and a little arithmetic that for E to be levering evidence, the hypothetical expectability of E on H_W (here put in boldface) must be very low—close to zero, comparatively speaking. That is:

$$[\,P(\mathrm{E}|\mathrm{Hw}\ \&\ \mathrm{k})\ \cong\ 0\ \,]$$

In other words, an application of Bayes' theorem shows that E's being levering evidence for H_L requires that the conditional probability of E on H_W (along with our background knowledge)

[54] Note that it is not a requirement of positive levering evidence (shifting one from square agnosticism to square belief) that the numerator of the Quotient have a value close to 1. What is required is the value *of the Quotient* be around 2; in some cases this may be because the numerator has a value of, say, .04, and the denominator a value of .02.

be very low—well below .5. But this more than satisfies what the CORE constraint requires! For the CORE constraint says that E is levering evidence for H *only if*, were H false, it is *unlikely* that one would get input E—in other words, that the conditional probability of E on not-H is below .5.

This result should renew our confidence that CORE is fundamentally right-headed. For using only Bayes' theorem, with no appeal at all to CORE, we have shown that in this McBrayer scenario, Bayes' theorem *entails* the constraint on levering evidence that is imposed by CORE. Thus a Bayesian approach also dictates that unless this CORE constraint is satisfied, E cannot be levering evidence for H_L.

4. BUT *IS* IT SATISFIED? BRIGHT LIGHT FROM AN OBSCURE CORNER

But *is* the above constraint—that the conditional probability of E on H_W be below .5—in point of fact *satisfied* in the scenario? So far we have shown only that *if* E is to be levering evidence in the scenario, *then* this constraint must be satisfied. But *is* it? Can we show this simply by focusing on the details of the scenario? We are now in a position to address this crucial question.

4.1. IS CORE SATISFIED?

To evaluate whether the conditional probability $P(E|\ H_w \ \&\ k)$ is indeed low, it is crucial to notice that there are two ways in which Holt can possess the winning ticket—namely, by possessing a winning Catholic ticket, *or* by possessing the winning Dutch ticket. For this reason, determining $P(E|\ H_w \ \&\ k)$, the expectability of getting Klep's testimony conditional on Holt's holding a winning ticket, requires that we consider *both* the Catholic and Dutch ways of Holt's holding a winning ticket. The expectability of getting E (Klep's testimony) given H_w (that Holt holds a winning ticket) will be a weighted average. It will be weighted sum of E's hypothetical expectabilities on *each* of the ways of Holt's holding a winning ticket, with each *multiplied by* a corresponding "weighting factor" of how likely that way of holding a winning ticket is. That is:

$$P\ (E|Hw\ \&\ k) = [P\ (E|\ Hw_d \ \&\ k) * P\ (Hw_d\ |k)] \quad + \quad [P\ (E|\ Hw_c \ \&\ k)\ * P\ (Hw_c\ |k)]$$
$$= [\quad \text{A d d e n d \# 1} \quad] \quad + \quad [\quad \text{A d d e n d \# 2} \quad]$$

Clearly, for $P(E|H_w \ \&\ k)$ to be very low, each addend must be very low. And in the actual scenario, as we shall now see, both addends are in fact very low—though for very different reasons.

In the second addend, the first factor is $P(E|\ Hw_c \ \&\ k)$. This is the conditional probability of evidence E—Klep's testimony to having stolen the Dutch ticket—on the hypothesis of Holt's having a winning Catholic ticket. This value is very high: 1, or nearly 1. But it is *weighted* by the second factor, $P(Hw_c\ |k)$. This is the *background* probability (on k alone) of Holt's having a winning Catholic ticket. And this value is very low. The background probability that Holt has the winning Catholic ticket is, after all, equal to the probability on k of having a Catholic ticket (namely, .5) *multiplied by* the probability on k of that ticket winning (namely, .000001). This value is .0000005—that is, one in two million. And since the first factor of Addend #2 is equal to 1, the value of Addend #2 is itself one over two million—which is very low indeed.

Addend #1 is also very low, though for opposite reasons. Its weighting factor is $P(Hw_d\ |k)$. This is the probability of Holt's having a winning Dutch ticket on k alone. The probability of having a Dutch ticket on k is .5, and since there is only one Dutch ticket, the probability of having a winning Dutch ticket on k is also .5. So the weighting factor here is .5. But what that factor weights is $P(E|\ Hw_d \ \&\ k)$— the conditional probability of E, on the hypothesis that Holt holds a winning Dutch ticket. And this probability is extremely low. For the hypothesis that Holt holds the winning Dutch ticket entails that Klep holds a Catholic ticket. And so on this hypothesis—given our background knowledge of Klep's compulsive honesty—the conditional probability of Klep's confessing to holding the *Dutch* ticket is extremely low. Addend #1, which multiplies this by the weighting factor of .5, is thus also extremely low.

The sum of Addend 1 and Addend 2 is thus very low. And since that sum is equal to $P(E| H_w \& k)$, this means that the conditional probability of Klep's testimony, on the hypothesis that Holt holds a winning ticket, is very low. And this more than satisfies what CORE requires. So the data of Klep's testimony *passes* CORE's acid test for levering evidence. This means that McBrayer's LOTTO argument fails at its first premise:[55]

> (Lotto-1) If CORE is right, then in the improved Lotto Story, Holt does not have levering evidence that his ticket is a losing ticket.

Here Lotto-1 (as applied to CORE) rests on the claim that the data of Klep's testimony does not satisfy the CORE constraint. On the conditional probability interpretation, this is false, and McBrayer's Lotto Argument is unsuccessful.

4.2. AND IS CORNEA SATISFIED TOO?

We've argued that on the conditional probability construal, *crux* is true: the data of Klep's testimony thus meets the CORE constraint, and CORE does not fall to the improved LOTTO counterexample. But what about CORNEA? For the internalist CORNEA to stand, what matters is not whether *crux* is true, but whether *crux* is something *reasonable for Holt to believe*.

Let's imagine Holt anxiously wondering whether he should regard Klep's new testimony as levering evidence, and thus be levered into regretfully but squarely believing that his ticket is a loser. Suppose Holt engages in the following soliloquy:

> Hmm. Klep confessed to stealing the Dutch ticket? Well, what should I expect, regarding Klep's testimony, if I were to have a *winning* ticket? Should I in that event expect Klep's testimony to be different than it is—i.e., should expect that he would have confessed to stealing the *Catholic* ticket? Or should I, in that event, expect Klep's testimony to be the same—to confess that the ticket he stole is the Dutch ticket?
>
> This is a bit tricky to figure out retrospectively, because as it happens, I already know what my brother *did* testify. So let's imagine I *don't* know this—that his confession is in a sealed envelope, that I rightly expect it to be as compulsively honest as usual, but that I haven't opened the envelope yet. And now suppose that my wife has just called me and told me that I've won one of the lotteries, but the reception is bad and I didn't hear which one. So having gotten the information that I have a winning ticket, how expectable is it, given that information, that Klep's letter in the sealed envelope says that the ticket he stole is the *Dutch* ticket?
>
> Intuitively, it is compelling clear that it isn't likely *at all*: on the supposition that I've got a winning ticket, the odds are vastly in favor of his having stolen—and so of confessing in his letter to have stolen—the *Catholic* ticket (leaving me with the Dutch ticket).
>
> But why is this? Well, there's no doubt some complex way to calculate this using probability theory, using the total probability theorem, weighted averages, and the

[55] As earlier, we take the truth of Lotto-1 to depend on the claim that Holt's evidence fails the CORE test.

> like. But I will leave that to others; I haven't studied that stuff since college. I'd just put it this way. Prior to opening Klep's letter, on my background information, it is 50/50 whether I hold the Dutch ticket or the Catholic ticket. But on new information that I have a winning ticket, this changes: my holding the Dutch ticket becomes vastly more probable than my holding the Catholic ticket; and—consequently—it becomes vastly more probable that the ticket stolen by Klep, and confessed to in his letter, is the *Catholic* ticket, not the Dutch ticket. As things actually stand, of course, he has confessed to having stolen the Dutch ticket; but if I were to have a winning ticket, it would be extremely likely that things would stand differently—that his confession would be to having stolen the Catholic ticket.

This line of thinking is reasonable, and it is one that Holt—a competent rational adult—could utilize. It thus makes it reasonable for Holt to believe that if he had a winning ticket, his evidence would likely be different. The CORNEA constraint is thus met. And Lotto-1 (as applied to CORNEA) is false:

> (Lotto-1) If CORNEA is right, then in the improved Lotto Story, Holt is not entitled to claim that his ticket is a losing ticket on the basis of Klep's testimony.

So, McBrayer's LOTTO scenario fails as a counterexample to CORNEA.

5. Two Evaluational Contexts

We've argued that McBrayer's counterexample does not tell against the probabilistic renderings of CORNEA and CORE. Why then does it seem so plausible against the counterfactual renderings of CORNEA and CORE?

We think that the counterfactual treatment creates a strong pressure to conflate two distinct contexts, which we'll call the pre- and post-evaluational contexts. Put simply, the pre-evaluational context is where data is yet to be evaluated; the post-evaluational context is where it has been evaluated. In the LOTTO scenario, Holt begins in a pre-evaluational context, where it is equally probable that he holds the Dutch or Catholic ticket. He then receives, as new data, the alleged evidence of Klep's testimony. Because Holt immediately discerns its evidential bearing, he immediately moves to the post-evaluational context, in which he has been (properly) levered to square belief that he holds a Catholic (and hence almost certainly losing) ticket. Now CORE is a norm addressing the question of whether some data alleged to be levering evidence actually is. In applying CORE (and CORNEA), it is thus crucial to answer the test question regarding the truth or reasonableness of *crux* within the *pre*-evaluational context: to answer it in the *post*-evaluational context would be to beg the very question at issue.

McBrayer's counterfactual construal makes this question-begging mistake a very natural one. In the LOTTO scenario, the test question takes the form "If it were false that Holt has a losing ticket (i.e., were he to have a winning ticket), would his new data likely be different?" To evaluate this as a counterfactual, one must determine the closest antecedent world(s)—i.e., the world(s) in which Holt has a winning ticket that are *most similar* (on a relevant similarity-ordering) to the real word. McBrayer takes the closest antecedent world to be one where Holt happens to have a winning Catholic ticket. He does so, presumably, because it is what Klep's testimony so strongly and obviously indicates. This shift to the post-evaluational context, while begging the question, is abetted by the standard Lewis-Stalnaker semantics. On that semantics, which antecedent world is "closest" is determined by how the real world actually is. When we are unsure or undecided about some relevant feature of the real world (like whether Holt has the Dutch or the Catholic ticket), using this semantics can thus create a strong pressure to pin down what the real world is like by making use of the new alleged evidence. The probabilistic treatment creates no such pressure. For when we ask the crucial question when the antecedent is epistemically "forked" (i.e., when we are undecided as between the Dutch way or the Catholic way of his holding a winning ticket), this treatment allows us to use weighted averages to take into account *both* prongs of the fork.

Where does this leave us on whether inductive evidence must be, in a Nozickean counterfactual sense, "sensitive" to the way the world is? Our results leave this as a vexing question. To the extent that Nozick's sensitivity requirement is tied to logical subjunctives, the answer will depend on whether the semantics for such subjunctives can be modified so as to incorporate the key strength of the conditional probability interpretation—namely, its use of Bayesian analysis and weighted averages—so as to handle "epistemic forks" between divergent ways in which the antecedent of a subjunctive conditional can be true or false. Our results here open the possibility that sensitivity accounts could be rejuvenated by the injection of this Bayesian strength into the possible worlds account of subjunctive conditionals.[56]

Calvin College and *Indiana University, Bloomington*

[56] This paper has benefitted from interactions with students in Wykstra's modal logic interim classes, participants in Jon Kvanvig's annual research conferences at Baylor, colleagues in the Tuesday Philosophy Colloquium at Calvin College, graduate students and faculty at Purdue University, and two commentators—Glenn Ross and Paul Draper—at a 2010 Pacific Division APA Symposium. We especially thank two anonymous *Faith and Philosophy* referees for their extensive comments.