Skeptical Doubt, Principal Bases, and Heuristics and Rational Norms

Jonathan Ichikawa

Penultimate draft: 27 March, 2025.
Forthcoming in *Asian Journal of Philosophy*

Abstract

A critical dicussion of Ángel Pinillos's (2023) *Why We Doubt*, focusing on questions about delineating the boundaries of skeptical doubt, the nature of evidence and its connections to sensitivity conditions, specifying principal bases for beliefs, and the connections Pinillos posits between skeptical intuitions and Bayesian norms.

Keywords: epistemology, doubt, skepticism, sensitivity, Bayesian epistemology, counterfactual conditionals

---

Ángel Pinillos's (2023) *Why We Doubt* explores a fascinating series of intersubdisciplinary philosophical questions about doubt. Its most central projects lie at the intersection between epistemology and cognitive science; due to the important role Pinillos attributes to reasoning about counterfactual conditionals, it also connects in interesting ways to metaphysics and the philosophy of language. It is a philosophically rich book, inspiring critical philosophical engagement. Following is some of mine.

1 Skeptical Doubt

Pinillos's central project is to explain human tendencies towards skeptical doubt. In particular, he seeks to provide a "unified cognitive explanation" for this tendency. (p. 9) Pinillos distinguishes this project from traditional first-order epistemic projects that focus on evaluating skeptical arguments; although he describes connections between the two projects, and takes a stand on some of the traditional project questions — signalling his commitment, for example, to the falsity of radical skepticism — Pinillos focuses primarily on cognitive explanations for skeptical doubt.

Pinillos seeks to explain skeptical doubt, but just what is skeptical doubt? The paradigms rise from familiar epistemological cases: tendencies to judge that one doesn't know that one isn't a brain in a vat, or that one's lottery ticket will lose, or that the company that runs most of the busses in town ran the bus that caused this accident. But since Pinillos's project is a *unified* cognitive explanation of skeptical doubt, evaluating its success requires a more precise specification of the explanans.

Some readers might be surprised, for example, that Pinillos stipulates that the skeptical doubt he is interested in is restricted to inclinations to deny *knowledge* in particular. Inclinations towards other skeptical judgments — like judging that one can't be *sure*, or that something isn't very *likely*, or that a belief wouldn't be *justified* — don't count as skeptical doubt, except insofar as they might incline one to deny knowledge. Cases of skeptical doubt are "inclinations to judge we don't know in a range of cases." (6) It is controversial whether questions about knowledge are particularly central in discussions about skepticism, and it is quite idiosyncratic to stipulate that skeptical doubt can *only* be about knowledge. I am unaware of any other epistemologist who uses the language in this way.

But setting aside the terminological question of whether denials of knowledge are the only instincts that count as skeptical doubt, I'd like to think a bit more explicitly about just *which* instincts to deny knowledge count as skeptical doubt. Pinillos also makes some idiosyncratic choices here, and they may implicate substantive matters.

First, Pinillos restricts his focus to denials of knowledge that stand in rational tension with other ordinary commitments, especially tendencies to *attribute* knowledge. (30) The tendency to judge that one can't know that one isn't a brain in a vat, for instance, is in tension with the tendency to judge that one can know the ordinary deliverances of perception. Pinillos doesn't say much about quotidian knowledge denials that are *not* in tension with ordinary knowledge ascriptions. For example, I am inclined to judge that I don't know exactly how many students will attend my class tomorrow morning, or whether this fair coin will land heads when I flip it. I also don't think Pinillos knows what I'm wearing today as I write this paragraph, for the humdrum reason that he's not here watching me and I haven't told him. I would be interested to hear more about why Pinillos's more restricted explanans is the methodologically appropriate choice. One wonders whether the broader category of *all* inclinations to deny knowledge might be a better candidate for having a unified cognitive explanation, compared to the potentially more gerrymandered category of inclinations to deny knowledge in those cases in which doing so creates a rational tension with other commitments.

There is also a question to ask, in settling the subject matter, whether different motivations for denying knowledge all count as "skeptical doubt". Here, Pinillos is explicit, but as in the case of the knowledge restriction, his stipulation is a bit idiosyncratic. Pinillos stipulates (9–10) that only *rational* and *discerning* tendencies to deny knowledge count as skeptical doubt. If one tends to deny knowledge because they think it makes them seem cool, this isn't a rational tendency. (Or at least it's not the kind of *epistemically* rational motivation at issue here — if one is right that being a skeptic would make one look cool, one might be prudentially rational in doing so.) And if one tends to deny knowledge in a given instance because one tends to deny knowledge in *every* instance, this isn't a discerning tendency, so that also doesn't count as skeptical doubt. Interestingly enough, I think it's a consequence of these restrictions that by Pinillos's lights, Pyrrhonian skeptics don't experience skeptical doubt — both because their skepticism is universal, and because it is motivated by the pragmatic goal of ataraxia.

I do not raise this as an objection to Pinillos; I do see the empirical motivation for restrictions along these latter lines. But I do think it would be easy for a reader to misunderstand the scope of Pinillos's subject matter.

2 The Skeptical Algorithm, Evidence, and Sensitivity

The centrepiece of Pinillos's cognitive explanation for skeptical doubt is a principle he calls "Principal Base Sensitivity," or *PBS*. He characterizes it thus:

PBS If S knows p and E is the principal base for S's belief that p, then E is sensitive to p. (65)

Pinillos supplements his statement of *PBS* with these two definitions:

Sensitive Base E is sensitive to $p =_{df}$ if p were false, then E would have a different truth value.

Principal Base E is the principal base for S's belief $=_{df}$ E is a base for the belief and it is the principal cause for the belief among all the bases. (65–66)

Pinillos's invocation of sensitivity is inspired directly by the rich epistemological tradition of invoking sensitivity as a necessary condition for knowledge. But Pinillos is careful to explain that he does not posit PBS as a *true description of epistemic normativity*; rather, it's a psychological posit about the way humans tend to reason. Most of the book is an application of this idea.

Given the definition of a base being sensitive, it is clear that Pinillos is assuming that principal bases must have truth values. His discussion fits most naturally with the idea that evidence is propositional, and that a given proposition counts as evidence (and so is eligible to be a principal base) only if it is known.[1] Pinillos intends to be less committal on this front, but I am not sure he can afford to be. He writes on p. 70:

> I assume that bases have truth-values. This means they are intimately connected to propositions. Facts and beliefs have this property. Arguably, mental states like perceptual experiences do so as well. For example, I may form a belief that the pencil in front of me is yellow. This belief may be based on a perceptual experience or a seeming with a certain content: that the pencil in front of me is yellow.

I have no quarrel with the idea that perception can have propositional contents, and therefore that there is a sense in which perceptual experience can have a truth value. But I don't think such states can possibly do the work that Pinillos wants principal bases to do in *PBS*. The idea of *PBS* is supposed to be that when one knows p, one has a principal base that is sensitive to whether p, and that when one doesn't, this will tend to generate skeptical doubt. But if one identifies the truth-value of one's perceptual experience with the truth-value of its content, then all veridical perceptual experiences are trivially sensitive to their contents in Pinillos's stated sense: they would have a different truth value, were their contents false. This means that perceptual experiences could *never* be insensitive principal bases, contradicting much of what Pinillos says about a variety of cases in the book.

On p. 155, for example, Pinillos gives:

> Consider a situation where you enter a restaurant and see a red table in front of you. You then form the belief that it is not the case that the table is white with tricky red lighting shining on it. Let's call this belief content, H. Suppose this belief is based on your perceptual experience of the table, which we call E. Since you never looked up to check for tricky lighting, many of us judge that you don't know H. PBS predicts failure to know.

[1] See e.g. Williamson (2000, ch. 9).

But if we understand the truth conditions of E along the lines given on p. 70, *PBS* does not predict that there will be an inclination to deny knowledge. When one looked at the red table, one had a perceptual experience or seeming with the content, *the table in front of one was red*. And that content would have been false, if it had been a white table with red lights. This way of thinking about sensitivity wrongly implies that perceptual experience *is* a sensitive base in this instance.

I think the solution to this problem would have been to follow e.g. Lewis (1996) in insisting that the evidence in these cases is the fact that one *has* experiences with a particular content, rather than with the content itself. This fact, of course, would have remained true in the counterfactual scenario described, and so would have been better able to play Pinillos's intended role. As far as I can tell, this change would have cohered reasonably well with the rest of Pinillos's book.

3 Principal Bases

Consistent with his cognitive explanatory ambitions, Pinillos describes the human tendency to reason according to *PBS* as an *algorithm*: "a rule where we specify the inputs and outputs. In our case, the inputs involve the details of an agent's epistemic situation with regard to one of their potential beliefs that p, and the output produces an inclination to judge the agent fails to know p. (15)" This kind of mechanistic conception — which Pinillos later describes as occurring within a "black box" (111), is part of the reason why Pinillos locates his discussion within cognitive science.

From time to time I had occasion to wonder, however, whether Pinillos's suggested input–output model might have glossed over some of the important questions it was designed to answer. In particular, I was not always satisfied with the book's discussion of principal bases. Sometimes I suspected that the choice about identifying the principal base was where most of the skeptical action might end up being — in which case, it is that prior question, not the application of *PBS*, that becomes key for explaining doubt.

Pinillos argues, plausibly enough, that the idea of *bases* is a relatively intuitive one. (66, 69) I'm comfortable discussing bases, even without a clear statement of their necessary and sufficent conditions. But I do not think that Pinillos's key notion of a *principal* base is similarly grounded in common sense, despite his (70–1) suggestion that an intuitive grasp suffices. He writes on p. 71:

> We easily recognize that the main or principal cause of the billiard ball entering the pocket is it beign struck by the cue ball (as opposed to my decision to play pool tonight). We easily recognize that the main or principal cause of the student passing the exam is the studying and not her sitting on a blue chair to take an exam. Of course, there will always be hard cases, but their existence should not detract too much from the felt plausibility of the idea.

But I worry that Pinillos has overestimated the specificity of ordinary intuition in these matters, and underestimated the prevalence and importance of the difficult cases. It's one thing to say a student passed the exam because she studied, rather than because she sat on a blue chair; but if you ask whether the principal cause was her *studying* or her *knowing the right answers*, or her *maintaining focus during the exam*, I think most people can do little but shrug their intuitive shoulders. When it came to principal bases for beliefs in certain scenarios, I often didn't know how to tell what the principal base was, and so didn't know what predictions *PBS* made.

Here is an example. On p. 238, Pinillos applies *PBS* to this case:

> Imagine a judge who decides a prisoner is guilty of rioting based on specific evidence (witness report) plus (insensitive) statistical evidence (95 percent of prisoners in his ward rioted). Neither evidence is sufficient for her to make up her mind. But together, the evidence is compelling. ... [T]he statistical evidence adds something to the case of guilt, so it should be possible that the statistical evidence adds enough to the body of evidence to secure knowledge. If this line of reasoning is right, then this may also be a genuine counter-example to PBS (*assuming the statistical evidence is the principal base here*). For those who still feel some lingering reluctance to say that the judge can know in this case, I submit that this inclination may be produced by the PBS heuristic. (emphasis added)

It is far from intuitively obvious to me that it makes sense to assume, as Pinillos does, that one of the two reasons the judge jointly considers to be probative is "the" principal base. But it is necessary to do so, for *PBS* to do the work Pinillos intends for it. If both considerations counted as part of the principal base, then the base would be sensitive.

Later in the same chapter, Pinillos suggests some even more surprising principal bases. He imagines leaving an ice cube outside in the hot sun, then considering an hour later whether it has melted. Vogel (1987) uses this case as a counterexample to a sensitivity condition on knowledge, since it is intuitively plausible that in this case, one knows that the cube has melted, but that if it hadn't melted, one would still believe it had (since one hasn't checked on it in the past hour). Pinillos argues that *PBS* does not make the incorrect prediction about this case; the key move is in specifying the principal base. He says (244) that the principal base is that the ice cube was left in the heat.

The truth conditions for "the ice cube was left in the heat" are deceptively subtle, and plausibly context-sensitive. There is one reading of something being "left" somewhere that requires only that someone left it there: one put it there and walked away. But there is another that requires, in addition, that it wasn't subsequently recovered. (If I set a book down on the table and walk away, then a few minutes later you come collect the book, was the book left on the table?) This distinction ends up playing a critical role in Pinillos's discussion of the ice cube case. He argues that the principal base — that the ice cube was left in the heat — after all *is* sensitive to the proposition that the ice cube has melted, because, as he puts it, "if the ice cube hadn't melted, it would have (had) to have been removed from the heat." (244) In other words, Pinillos does not treat the principal base in this case as the directly perceived historical fact that the ice cube was left in the heat in the sense that someone put it there and walked away — he thinks it's the temporally stronger fact that it was left in the heat in the sense that nothing subsequently removed it. This gives Pinillos's discussion of this case a surprisingly externalist, neo-Moorean character. The principal base is a temporally emergent fact, with substantive contingent entailments about the future, relative to the time at which the principal base was perceived.

I do not object to the epistemic plausibility of such externalist approaches to knowledge; I'm actually quite sympathetic to them. But Pinillos's treatment of principal bases here was potentially concerning to me, in two ways. First, it was difficult to square it with his earlier (85–8) discussions of neo-Moorean responses to skepticism, like using one's perceptual knowledge that one has hands as a basis for concluding that one isn't a brain in a vat. There, Pinillos explained the intuitive resistance to such responses by denying the tendency to attribute such a strong principal base. But doing so is structurally just like Pinillos's response to the ice cube.

Second, relatedly, it is looking more and more like much of the substantive epistemological action depends on the specification of the principal base. This risks undermining Pinillos's distinctively *explanatory* ambitions for *PBS*. Whether *PBS* makes predictions of doubt depends on one's assessment of the principal base. But the elusiveness of the specification of principal bases raises an alternate possibility: perhaps people tend to identify weaker principal bases when they are inclined to doubt. If so, *PBS* isn't explaining the inclination to doubt; it's just responding to it.

4 Function, Belief, and Credence

One of the most creative and surprising projects in the book is Pinillos's attempt to connect the heuristic skeptical rule *PBS* to a genuine Bayesian norm of rationality. This primarily happenss in Chapter 5. Pinillos thinks that, as in the case of other psychological heuristics, we tend to use *PBS* because it is part of a generally effective strategy towards helping us comply with a more complicated, and genuinely binding, norm. This works because, Pinillos says, *PBS* "approximates" a true principle of rationality. (128)

The genuine norm Pinillos has in mind is this one:

Norm For all subjects S, evidence E, propositions p, and times i and f, with i before f, where n is a degree of rational confidence below which knowledge is impossible, and C is a rational credence function, if

(a) $C_i(E \mid \neg p) = 1$,

(b) $C_i(p) < n$, and

(c) E is the strongest proposition learned after i (up until and including time f),

then S does not know p at f.[2]

Pinillos describes *Norm* as a "Bayesian" principle. Since its consequent is a claim about knowledge, it is really a kind of bridging principle that links rational credence to knowledge. Perhaps Pinillos had in mind the genuinely Bayesian principle that would result from replacing the denial of knowledge in the consequent with the conclusion that $C_f(p) < n$. (The implication from this conclusion to the denial of knowledge is trivial, given the stipulation about n.)

It is prima facie surprising to suggest that the central role for *PBS*— a heuristic about knowledge, which has its home in the binary belief-centred approach to epistemology — is to approximate a Bayesian principal about rational credence. The connections Pinillos draws are substantive and interesting, but I didn't leave his discussion wholly convinced. I'll start my critique with a couple of smaller points. I'm not sure these get at the heart of the problem, but I do think they're important for thinking clearly about any possible connections here.

First, I don't actually think *Norm* is true in generality. Its intuitive plausibility depends in part on the implicit assumption that learning new evidence and conditionalizing on it is the *only* permissible way to change one's credence functions over time. Suppose S's total body of evidence changes in some other way between i and f — suppose, for instance, that S *forgets* a certain body of evidence B during that interval, and so no longer conditionalizes on B at f.[3] If B misleadingly

[2]This is my own formulation of Pinillos's principle *Norm*, stated a bit more precisely than his p. 138 statement.

[3]See Talbott (1991) or Williamson (2000, 219) on forgetting and the limits of conditionalization to model rational belief change.

contraindicated p, then between i and f, S's total evidential state might change in a way that brings p above n, even if E is the strongest proposition learned, consistent with both $C_i(E \mid \neg p) = 1$, and $C_i(p) < n$.

I think this by itself is a relatively minor problem for Pinillos; it's easy to imagine him responding by adding a further condition to the antecedent of *Norm*, specifying that no other change to the total body of evidence has occurred. That would solve my first complaint, but it would further underscore my second one, which is that the further one goes in stipulating idealized conditions to define the scope of *Norm*, the less clear it will be that the heuristic *PBS* will be a good approximation of it. I raise some challenges to this idea in the next two sections.

5 Strongest Evidence vs Principal Base

I suggested in section 2 that Pinillos's notion of a principal base isn't always as clear or precise as it might appear. We form beliefs in response to large and complex bodies of evidence, and the project of specifying a particular E to identify as "the" principal base will inevitably be a bit arbitrary. It also points to a major important difference between *PBS* and *Norm*. The latter is given in terms of *the strongest proposition learned between two times*; it encodes literally everything the subject has come to know. *PBS* can only serve as an approximation of *Norm* to the degree to which the principal base is an approximation of that body of knowledge. And I didn't see much reason to think there would often be even a close approximation here.

Pinillos acknowledges this distinction, but says less about it than I would have expected. (He has a bit more to say about a different difference than the one I emphasize here, namely that *Norm* but not *PBS* depends in important ways on one's prior probabilities.) He says on 140 that one of the reasons it would be cognitively costly to apply *Norm* directly is that "it would requires [sic] us to determine the strongest proposition learned during a time interval," and that since *PBS* doesn't do that, it "is an excellent candidate for being a heuristic short cut for *Norm*." But absent a richer discussion about how principal bases are specified, I don't see any reason to be confident that the principal base will even plausibly approximate the strongest proposition learned. And indeed, if it will, then once again, much of the interesting heuristic action may be in the project of identifying the principal base, rather than in applying *PBS*.

6 *PBS*, *Norm*, and *Norm-Approximate*

Another major difference between *Norm* and *PBS* is that where *PBS* invokes a sensitivity condition, *Norm* makes use of a probabilistic one. Pinillos takes these to be quite closely related. When he introduces the principle *Norm*, he does so alongside what he calls an "approximate" version of it; this helps see why he thinks there's a connection to *PBS*.

Here are both principles, paraphrased in my own words to better bring out the structural parallels. His own formulations are on p. 138.

Norm-Approximate For all subjects S, evidence E, and propositions p, if

(a) S knows for sure that E is insensitive to p, and

(b) S didn't know p before learning E,

then S does not know p after learning E.

Norm For all subjects S, evidence E, propositions p, and times i and f, with i before f, where n is a degree of rational confidence below which knowledge is impossible and C is a rational credence function, if

(a) $C_i(E \mid \neg p) = 1$,

(b) $C_i(p) < n$, and

(c) E is the strongest proposition learned after i (up until and including time f),

then S does not know p at f.

There are three differences between *Norm* and *Norm-Approximate*, corresponding to the three conditions in *Norm*. In condition (a), *Norm* invokes insensitivity instead of probability; in (b), *Norm* cites a low prior rational credence instead of a prior lack of knowledge; and condition (c) was absent from the approximated version. My focus in this section is on the relationship between sensitivity and probability, and the corresponding distinct versions of condition (a).

Pinillos draws a connection between these two versions of condition (a) on his p. 137, where he says that the *Norm* version is entailed by the approximated version: "Let us assume that S *knows for sure* (confidence at 1) that E is insensitive to p at some initial time i. It follows that $C_i(E \mid \neg p) = 1$."

Prima facie, this is quite a surprising claim. There are very ordinary apparent counterexamples. I have been working on this symposium contribution for the past couple of hours, and have not checked the news, so I don't know whether new tariffs have been announced in the last hour. Let p be the proposition that new tariffs have been announced. I am about to flip a coin; so I know that either the coin will land heads (H) or it will land tails (T). And I am certain that the result of the coin flip is insensitive to whether tariffs have been announced: both H and T are insensitive to p. But both $C(H \mid \neg p) = 1$ and $C(T \mid \neg p) = 1$ are obviously false for me. If one of them were true, then if I checked the news and learned that there are no new tariffs, that would make me rationally certain of the result of the coin flip.

So why does Pinillos think there is a stronger connection here? I am not entirely sure. One might wonder whether he intended the entailment to be restricted to the case where E is already known or observed. But I don't see that this could be Pinillos's intent, because in his more careful statement of *Norm* he explicitly stipulates that E is learned *after* i, the time at which the conditional probability is said to be 1.

Moreover, we can generalize my case above to provide counterexamples even against this restricted entailment claim. Suppose as before that I haven't turned on the news, don't know about recent tariff announcements, and flip a coin. This time I observe that H is true — the coin lands heads. And for specificity, let's stipulate that there have indeed been new tariff announcements in the past hour, so p is also true, albeit unknown to me. As before, I know perfectly well that coin flip results do not co-vary with tariff announcements. I know H, but I also know that H is totally insensitive to whether p. Should I be fully confident that *if there hadn't been new tariff announcements, the coin would have landed heads*? Not obviously. On one family of approaches to subjunctive conditionals, the way to evaluate this question is to imagine rolling back history, changing the tariff announcement decision, and letting the world otherwise run its course according to the laws of nature, including perhaps the probabilistic and indeterministic processes that govern coin flips. On such approaches, I should consider it highly uncertain whether the coin would have landed heads, had there been no tariff announcements, even though I remain certain that H is not sensitive to p.[4] In his p. 137, fn. 2, Pinillos rules out these approaches to subjunctive conditionals, committing to idea that there is always a unique closest possible world in which a given proposition is true, and that a counterfactual is true if and only if its consequent is true in the unique nearest world where its antecedent is true. This is a commitment to the principal of so-called "conditional excluded middle," and it is a substantive and controversial commitment.

But there are cases where even conditional excluded middle won't help. These are cases where one can be arbitrarily confident that the evidence is not sensitive to a hypothesis, but where the conditional certainty Pinillos describes is implausible, even if conditional excluded middle is true.

Consider a gamble that plays out in two stages. First, a fair coin is flipped. The result of that flip determines which of two fair six-sided dice is rolled. If the coin lands heads, the red die is rolled; if tails, the blue die is rolled. The result of the game is the result of whichever die is rolled. This is a pretty uninteresting game for gambling purposes — the coin flip does nothing interesting to the probabilities within the game. But it is interesting for evaluating the connections between counterfactuals and probability.

Suppose the game is played and the coin lands heads (C_H), the red die is accordingly rolled, and it lands on 6. So the result of the game is 6 (G_6). Suppose one observes the dice roll but not the coin flip, and so knows that G_6, and wonders whether C_H. G_6, one knows perfectly well, would be a terrible basis for accepting C_H; the result of the game is insensitive to the result of the idle coin flip with which the game begins. One can and should be sure that G_6 is insensitive to C_H. But — and here is the problem for Pinillos — one should not be at all confident that if C_H had been false, G_6 would have been true. Not even if we accept conditional excluded middle and the idea that there is always a unique nearest possible world in which a given subjunctive condition is true. In the case contemplated, one's credence in C_H should be 0.5. If we assume, with the tradition Pinillos commits to, that the conjunction of $\neg C_H$ and G_6 suffices for the truth of the subjunctive conditional that $\neg C_H \mathrel{\Box\!\!\rightarrow} G_6$, one can be certain that this conditional is true *if* C_H *is false*.

But in this case, one should have a 0.5 credence that C_H is true, and therefore that the red die was actually rolled. And in this scenario, then the counterfactual possibility one considers when supposing $\neg C_H$ is one in which the blue die was rolled instead of the red one. And even if we assume that there is a unique closest world where that's what happened, there is no reason to suppose it's one where the blue die came up 6, and so there's no reason to suppose that G_6 would have been true. So this is a case where I know that G_6 is insensitive to C_H, but where I should think that the probability of G_6 given C_H, is less than 1. (In fact, we can be precise: it should be $(1 \times 0.5) + (\frac{1}{6} \times 0.5) = 0.583$, if we assume conditional excluded middle. If we don't, it should just be $\frac{1}{6}$.)

These are some of the reasons I am skeptical that *Norm-Approximate* approximates *Norm*, which in turn provides further pressure against Pinillos's idea that *PBS* approximates *Norm*.

Bibliography

Edgington, Dorothy. 1995. "On Conditionals." *Mind* 104:235–329. ISSN 0026-4423. Publisher: [Oxford University Press, Mind Association].

Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74:549–567. doi: 10.1080/00048409612347521.

[4] See Edgington (1995).

Pinillos, N. Ángel. 2023. *Why We Doubt: A Cognitive Account of Our Skeptical Inclinations*. Oxford, New York: Oxford University Press. ISBN 978-0-19-887196-5.

Talbott, William. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies* 62:135–150. doi:10.1007/bf00419049.

Vogel, Jonathan. 1987. "Tracking, Closure, and Inductive Knowledge." In Luper-Foy Steven (ed.), *The Possibility of Knowledge: Nozick and His Critics*, 197–215. Rowman & Littlefield.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press. ISBN 0198250436 (alk. paper).