

Article

Dretskean Sensitivity Wins the Lottery!

Kelly Becker 

Department of Philosophy, University of New Mexico, Albuquerque, NM 87131-000, USA; kbecker@unm.edu

Abstract

One major strand of post-Gettier epistemology concentrates on shoring up the connection between belief and truth by postulating a condition that includes a modal requirement. This approach abandons the traditional emphasis on reasons. A second major strand maintains the focus on reasons, aiming to distinguish those that are sufficient for knowledge from those that are not. The first strand has generated an endless parade of counterexamples. Some version of the second strand may be true but insufficiently informative. For example, perhaps knowledge really is undefeated justified true belief, but what is a genuine defeater, and what is justification? An early entrant in the post-Gettier competition, Fred Dretske, proposed a modal requirement that ties reasons, not belief, to the truth, thus incorporating key components of both strands. Drawing solely on resources in Dretske's presentation, I develop his notion of conclusive reasons in an account of knowledge that I call reasons-based Dretskean sensitivity (RDS), which resolves most, if not all, of the major problems afflicting its wayward cousin, Nozickian sensitivity, and indeed many other epistemological difficulties. This paper lays out the RDS resolution of Hawthorne's lottery puzzle.

Keywords: Dretske; Hawthorne; lottery; Nozick; reasons; sensitivity

A major strand of post-Gettier epistemology veers away from the traditional epistemological concentration on the quality of one's reasons. The motivation—reasonable, given the nature of Gettier cases [1], in which true belief with good reasons falls short of knowledge—is to eliminate epistemic luck (commonly indicted for sourcing the shortfall) by crafting conditions that somehow tether belief to truth. Prominent candidates for the requisite tethering mechanism include direct, causal relations [2], modal relations such as sensitivity [3,4] or safety [5–7], and explanatory relations [8–10]. Understandable though the motivation is, an implicit upshot of this paper is that the frantic quest for a tethering mechanism engineered to what turns out to be a fairly simple task, resulting in a fair number of Rube Goldberg contraptions,¹ stems from a misdiagnosis, namely, an ailing connection between *belief* and truth, which has led to decades of unchecked disease. A separate strand of thought retains the focus on reasons and tries to shore up the justification condition to achieve a correct analysis of knowledge, for example, by claiming either that justification may not include any false grounds [11] or that knowledge demands justification with no defeaters (Refs. [12,13], though there is much disagreement on what constitutes a genuine defeater). While the preceding only partially represents post-Gettier strategies, it suits my present purpose of revealing a gap in thinking that, set out in a certain light, was filled already in Dretske's rich and rewarding notion of conclusive reasons [14]. Dretske did not take the Gettier bait of forsaking reasons. Instead, he exploited modal notions to articulate the nature of (good) conclusive reasons—he tethered the tethering mechanism to reasons.²

In a series of papers [16–18], I develop Dretske's notion of conclusive reasons, interpreting them as reasons that are sensitive to, or track, the truth, to show how reasons-based Dretskean sensitivity (RDS) resolves the most difficult issues facing what has been



Academic Editor: Joseph Ulatowski

Received: 20 July 2026

Revised: 14 September 2026

Accepted: 15 September 2026

Published: 18 September 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

considered the *locus classicus* of sensitivity theory, namely Nozick's [4] *Philosophical Explanations*. Nozick's development of sensitivity set the modalized epistemology agenda for decades. His fatal misstep was to lose track of the importance of reasons. Critics point out that Nozickean belief-sensitivity is incompatible with knowing that one does not believe falsely [5,19]. RDS solves this problem at least as well as do the updated Nozickean responses based on revised Nozickean sensitivity (RNS; see for example [20])—the working standard of sensitivity that holds actual belief-forming methods constant in evaluating the sensitivity conditional [16]. Critics have shown that even RNS is either incompatible with [5,19,21] or makes an unprincipled hash of [22] inductive knowledge. RDS solves this problem, too [17]. RNS seems to have been demolished once and for all by Kripke's [23] version of fake barns, having already suffered a major blow early on by hologram (fake) vases, whose presence is caused by unseen genuine vases, a case raised and briefly but unsatisfactorily dismissed by Nozick (p. 190, [4]). Forbes (p. 45, [24]) recognized the latter case as deeply problematic. RDS resolves Kripke's fake barns—explaining why Henry knows *neither* that there is a red barn on the hill *nor* that there is a barn on the hill, compatibly with closure—and other case-based problems that are difficult for RNS [18].

All of that is well and good, but not especially relevant to the present topic, which focuses on the lottery and related problems. I mention the difficulties for RNS above only to give a sense of the power of RDS. Responding to those difficulties is mostly a matter of playing defense on behalf of sensitivity theory, whose specific troubles they are. But RDS is a promising approach to more general, more difficult epistemological puzzles. Here, I first set out RDS as a finished account (§I), without much background on how it was constructed because that story is told in the aforementioned papers. I will, however, briefly compare RDS to RNS. Second, I apply RDS to Hawthorne's [25] lottery problem (§II), weaving in important cases that are implicated to various degrees either by it or by related concerns of Hawthorne's. I conclude with general remarks about RDS in comparison to competing theoretical frameworks.

1. Reasons-Based Dretskean Sensitivity (RDS)

Rather than offering RDS in a groundwork-first, account-last format, appropriate to a first construction and presentation, I state the account first, accompanied by a few comments.

1.1. RDS

RDS:

- (1) *S* knows that *p* only if *S* has a conclusive reason *R* for believing that *p*.
 - (2) *S* has a conclusive reason *R* for believing that *p* iff:
 - a. *R* is a conclusive reason for believing that *p*.
 - b. *S*'s belief that *p* is based either on a relevant mental state ψ whose content is *R* or on an *R* experience.
 - c. (i) *S* knows that *R* is the case*, or
(ii) *R* is an experiential state of *S*'s.
 - (3) *R* is a conclusive reason for believing that *p* iff, were *p* false, *R* would not be the case.
- [*To flesh out the recursive nature of RDS: on the basis of some mental state ψ_i , whose content is a conclusive reason *R_i*, that is, satisfies clauses (2) and (3), but such that *R_i* need not itself be a conclusive reason for *p*.]

1.2. Explanatory Comments

For (me and) Dretske, a reason need not be cognitively accessible. A reason, *R*, when it is the content of a relevant mental state, is the basis of *S*'s belief that *p* or the basis of another reason in a chain of reasons leading to *p*. *R* is externalistically construed,

characterizable by some proposition, for example, that such-and-so is the case, or that S has experience(s) of a specific kind (whether or not S can adequately describe or even remember the experience(s)).

To be conclusive, where R is not an experiential reason (R_e), R must be known (c.i). To be known, R must have its own conclusive reason. The regress stops (for empirical knowledge) at an experiential state—a reason based in experience— R_e (c.ii).

Nozickean sensitivity (if p were false, S would not believe that p) and revised Nozickean sensitivity (RNS: if p were false, S would not believe that p, using S's actual world belief-forming method, M) both characterize sensitivity in terms of belief—the first version bluntly, the second version in a more nuanced, or “method-relative” way. RDS characterizes sensitivity in terms of reasons. In addition, RDS does not say that if p were false, S would not *have* their reason R. It requires only that if p were false, R would not be the case, and this can be so regardless of whether S would believe or otherwise have R if p were false (and regardless of whether S would believe that p). This is a huge competitive advantage for RDS, as we shall see. Importantly, for a foundational, experience-based reason, R_e , there is no possibility that S has it while it is not the case. Such a reason is characterizable generally and roughly as *S's experience of such-and-so type*. To say that R_e would not be the case, then, is to say that S would not have that experience.

Very importantly, counterfactual conditionals are not transitive. It can be true that if p were false, q would be false, and if q were false, r would be false, but false that if p were false, r would be false. (If I wore red suspenders, my outfit would look silly. True. If I were a fireman, I would wear red suspenders. True. If I were a fireman, my outfit would look silly. False. I'd be wearing firefighting gear, with firefighter-red suspenders to hold up the heavy flame-retardant pants.) So, S can know p on the basis of reason R, if R is sensitive to, or tracks, p. S can know R on the basis of R_e (an experience-based reason) if R_e tracks R. And this can all be the case even if R_e does not track p, that is, even if S would have the same experience if p were false.

2. RDS, the Lottery, and Other Salient Cases

RDS is comparatively lean. It is also agile in handling epistemological problems, which together with its leanness implies elegance. It does not generate the handwringing about how to construe methods that has vexed RNS. (RDS leaves room for questions about how to understand the epistemic basing relation—among others, about required causal relations—but this is an issue for everyone in epistemology.) It does not raise thorny questions about the space of possible worlds relevant to its satisfaction, as does safety (or what needs to be held constant in applying the safety conditional,³ or which sets of propositions are relevant to the assessment of the safety of a particular belief (see [27] on global safety)). It does not require further considerations about practical stakes or context-dependent truth conditions or other bells and whistles whose main purpose is either to get desirable results regarding epistemological puzzles or (inclusive) to capture ordinary usage of “knows.” Let's see.

2.1. Basic Lottery

Why does S not know that their lottery ticket has lost, in advance of the announcement of the winner, despite the extremely long odds of winning? p: My ticket is a loser. R: The odds that it wins are miserable. If p were false, R would still be the case. R is not a conclusive reason, per RDS. (Nozickean) RNS also handles this case easily. I rehearse the example because it serves as a warm-up to the main program and because Dretske's own treatment of the lottery ([14], p. 4) reveals that he is thinking of conclusive reasons as ones that are sensitive rather than safe. (I mention this because I have been told in conversation

that Dretske should be read as endorsing safety.) Now to the main program: Hawthorne's lottery puzzle.⁴

2.2. Hawthorne's [25] Lottery Puzzle

1. The Puzzle

Safari. S knows that p (an ordinary proposition): I [S] will not be able to afford to go on safari this year. S does not know that q (a lottery proposition): My lottery ticket is a loser. However, p entails q, S knows this, and (let us assume that) S has no trouble deducing the obvious conclusion. So how can S know that p but not q?

One fairly obvious suggestion is that S knows neither p nor q, citing the facts that knowledge is closed under known entailment and that S obviously does not know that q. (Also, many prominent theories imply that S does not know that p: The belief is neither (a) safe nor (b) sensitive, (c) there is a defeater for S's justification, and (d) p's being true does not *explain* why S believes that p. Could they *all* be wrong?) One of Hawthorne's concerns is that, if we accept that S does not know that q, and we accept that knowledge is closed under known entailment, hence that S also does not know that p, then we need an explanation for why explicit lottery propositions are distinct from propositions that can be treated in a lottery-like way, such as *N.N. will not have a heart attack in the coming days* (p. 3, [25]), *my car has not recently been stolen from Lot F, where I parked it* (p. 4, [25]), and *N.N. has not died in the last few minutes* [25]. If we lack knowledge of all such lottery-like propositions, then, given closure, we lack knowledge of all sorts of ordinary truths, for example, that *N.N. will be at the party Saturday, that my car is in lot F*, and that *N.N. is right now the U.S. president*. In effect, lottery-like propositions wreak havoc on knowledge perhaps to an even greater extent than do anti-skeptical propositions that the skeptic claims are also unknowable, such as I am not a brain-in-a-vat (BIV), insofar as many philosophers think we can know the latter (p. 6, [25]).

Perhaps, then, we should simply flip modus tollens to modus ponens in a Moorean way—I do know I have hands, so I do know I'm not a BIV. But that move does not generalize well. By the same reasoning, I could use my knowledge that I cannot afford a safari to infer and come to know that my ticket is a loser. But I don't know that. So what's the salient difference between lottery propositions and skeptical ones—the difference that explains how the latter are knowable (assuming that they are) and the former are not? And where do lottery-like (but not explicitly lottery) propositions fit in the scheme?

Nozickean sensitivity was designed partly to cauterize the wound to our knowledge statuses created by (what Nozick took to be) our lack of knowledge with respect to anti-skeptical propositions. One knows one has hands because if it were false, one would not believe it, even if one does not know one is not a BIV. It is also true that one does *not* know one will not be able to afford a safari, because one might well believe it even if it were false. But lottery propositions, one might think, are rare, and therefore something like sensitivity suffices to stave off rampant skepticism. Hawthorne's assimilation of lottery-like propositions to explicit lottery propositions, both generating skeptical implications, is intended to render the thought naïve.

Indeed, Hawthorne discusses Nozickean sensitivity and its RNS variant in some detail (pp. 9–15, [25]). One major complaint is that RNS saves ordinary empirical knowledge only by denying closure—S knows they have hands but not that they're not a BIV. The concern on which Hawthorne focuses, however, is that RNS denies knowledge not only of propositions that entail explicit lottery propositions, such as *Safari*, but also of propositions that entail lottery-like propositions, such as "*My car is in Lot F*." If my car were stolen and no longer in Lot F, I would still believe it is in Lot F. The wound is cauterized, but the infection has spread further than we'd thought. The naïve thought is exposed as such. On

this point, Hawthorne discusses a case that has long been an irritant for sensitivity theory and suggests that even it is lottery-like (p. 19, [25]). The Heartbreaker is a difficult golf hole. It seems fairly obvious that I (can) know that not all 60 golfers will ace the Heartbreaker [29]. But if it were false—in the closest worlds where all 60 golfers do ace the Heartbreaker—I would believe it anyway.

To summarize the lottery puzzle and its specific application to sensitivity, Hawthorne argues that much ordinary knowledge is under threat insofar as it implies lottery-like propositions that we seem not to know. RNS staunches the bleeding, preserving ordinary, non-inductive knowledge, but only by (a) denying closure and (b) doing little, perhaps nothing at all, to preserve apparently easily achievable inductive knowledge (*that my car is in Lot F; that not all 60 golfers will ace the Heartbreaker*). After a review of Hawthorne's approach to the puzzle (2.2.2), I shall argue that RDS, like RNS, implies that (a) one does not know truths that entail explicit lottery propositions, but (b) one can know truths that entail anti-skeptical propositions, and that RDS generates these implications without closure violations (2.2.3). I then argue (2.2.4) that, while Hawthorne's assimilation of lottery-like propositions to explicit lottery propositions is insightful and challenging, there are differences between them—the sorts of differences that sensitivity and safety theorists have been exploiting for decades—and that RDS can exploit them, too, while preserving closure.

2. Hawthorne's Proposal

Hawthorne's own strategy is to say that we can know lottery-impacted propositions, such as p (I will not be able to afford to go on safari this year) and know that they entail explicit lottery propositions, such as q (My lottery ticket is a loser)—and thereby also know that q , which to my mind is a significant data point *against* his view—but when we explicitly entertain p and $p \rightarrow q$, our practical interests (such as the desire to go on safari) can raise the stakes regarding whether a proposition is true, in turn making knowledge more difficult to achieve and, in some cases, erasing the knowledge that we had had, namely that p . Hawthorne cashes out the phenomenon in terms of practical reasoning: “Insofar as it is unacceptable—and not merely because the content of the belief is irrelevant to the issues at hand—to use a belief that p in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time” (p. 176, [29]). (An example: Since I (know I) will not be able to afford to go on safari this year, I might as well sell this worthless lottery ticket for a penny. Bad reasoning implies a false premise.) It follows that knowledge can depend on practical interests. Normally and most of the time, being confronted with the BIV hypothesis has no practical impact, unlike lottery propositions and those practically impacted by them, and generates no unacceptable practical reasoning. Therefore, it doesn't erase one's knowledge that they have hands, from which one may infer and come to know they're not a BIV.

In the service of arguing that RDS treats the lottery problem better than both RNS and Hawthorne's interest-relative invariantism (aka sensitive moderate invariantism, aka subject-sensitive invariantism), I raise two concerns about this picture, one that, so far as I know, is mine alone; the other is fairly obvious. The first has to do with the widely accepted, standard Williamson [6]–Hawthorne [25] formulation of the closure principle, roughly:

W–H Closure: Necessarily (for all S and p, q), if S knows that p and competently deduces q from p [Hawthorne: while maintaining knowledge that p], then S knows that q .

One of the virtues of W–H Closure is that it has the air of precision and appears to be water-tight. Another virtue is that it is true. Always nice. However, in my view, and perhaps it is only mine, W–H Closure also surreptitiously plays into Hawthorne's hands, since it is stated in terms both of (a) S in fact competently deducing q and of (b) S maintaining knowledge that p . (b) positions Hawthorne to say that, when considering q

raises practical stakes with regard to whether p is true, S may lose knowledge that p , which is part of his strategy for saving closure when one does not know that q . (a) Precludes interpretations of closure that do not depend on q explicitly coming to mind for S , which in turn allows Hawthorne to say that, made explicit, some q possibilities indeed raise stakes, which again is part of his strategy for saving closure. However, before pragmatic encroachment ever became a “thing,” it would not have been crazy to think of closure as saying that if S knows that p and knows that p entails q , that supplies everything needed for S to know that q , via deduction, straightaway, regardless of whether one performs the deduction. (I surmise that this sort of closure was always the basis of the widely held criticism of RNS.)

The second concern is simply that once we enter the territory of pragmatic encroachment on epistemic matters—that is, where practical stakes can partly determine one’s epistemic position—there are bound to be lingering questions. For example, when exactly is a premise unacceptable in practical reasoning, and does lack of knowledge imply unacceptability? (Feldman [30] probes this.) Is it true that knowledge can come and go so easily, simply because of practical concerns? On the one hand, one nice implication of Hawthorne’s approach is that, generally and typically, we know that skeptical hypotheses are false because their falsity (I’m *not* a BIV; the animal in the pen is *not* a cleverly-disguised mule) is inferred from premises that have no practical consequences (I have hands; the animal is a zebra). But they could. (“Hey buddy, I bet you a million dollars that you don’t have hands!” “Sure I do. See?” “Prove it. You could be a BIV and you’d still think that.” “Aw, there goes my safari, my dream home, *and* my new carbon-plate Nikes.”) In such a case, Hawthorne can consistently say that one does not know that one has hands, a proposition that implies that one is not a BIV, or that one does not know that the animal is a zebra, a proposition that implies that it is not a cleverly disguised mule. Okay. But are our knowledge statuses that fickle? Is there an identifiable line between the BIV possibility having and not having knowledge-precluding practical implications? (The concern implied by this question is not that there may be gray areas where it is indeterminate, or at least underdetermined, whether S knows that p . I think that may be true, as will become apparent. The concern is that, if we can understand the source of indeterminacy by reference to “traditional,” more obviously epistemic factors, then pragmatic encroachment needlessly complicates theory.) I do not claim that Hawthorne owes us answers. I merely suggest that the picture, in trying to generate all the right-seeming implications, holds knowledge hostage to factors that nobody would have dreamed of 30 years ago and yet remain nebulous.

Perhaps the previous comment sounds benighted. A more enlightened way to think about the introduction of pragmatic encroachment into the theory of knowledge is as a paradigm shift and, as such, presented simply as an improved, all things considered more satisfying understanding of knowledge, one that takes practical interests to be essential not only to understanding the point of knowledge but also to understanding its nature. As a new paradigm, it brings its own unique puzzles. No paradigm is perfect. Okay, but paradigm shifts are justified essentially by reference to accommodating recalcitrant anomalies. What if RDS accounts for the anomalies while retaining the old paradigm of invariantist, reasons-based epistemology, without pragmatic encroachment?

The second concern and its associated matters just mentioned are also related to my first concern about how to characterize closure. Here is a less precise, but intuitively compelling, pre-pragmatic-encroachment-days characterization:

Closure-EP: If S knows that p and that p entails q , then S is in a strong enough epistemic position to know that q [full stop].

Hawthorne claims that, because they are straightforwardly deducible from ordinary “known” propositions, one often knows lottery propositions, so long as they are not explicitly entertained in practical reasoning. Presumably this is because he accepts something like Closure-EP. But because Hawthorne is more interested in explaining how explicit reasoning involving lottery propositions can undermine knowledge one might otherwise have had, he prefers the W–H Closure formulation. I suggest that any position that accommodates Closure-EP without falling into skepticism, without implying that we know lottery propositions, and without the counterpoint line of practical interests woven into the traditional epistemological melody, is worth serious consideration. RDS is just such a view.

2.3. RDS on Safari! (vs. RDS on Skepticism)

To see this, let us first treat Hawthorne’s lead-off lottery case with RDS. What is S’s reason R for believing that p?

p: I will not be able to afford a safari this year. [It is not the case that, at time t a few months in the future, I will have enough money to go on safari.]

R: It is extremely unlikely that I will come up with that kind of money.

If p were false, R would still be the case. R is true but inconclusive. S does not know that p. This may strike you as an undesirable implication. One reason might be that you think S does know that p. There are, however, reasonable grounds to deny this (see second paragraph, II.B.1), even if the intuition that S knows is strong among some philosophers. (How strong is it? Dretske did not share it. See Dretske (pp. 5–6, [14]) for his treatment of a similar lottery puzzle.) And even Hawthorne says such knowledge disappears when the stakes go up. A second, stronger concern one might have about denying that one knows that one will not be able to afford a safari is that the skeptical threat might spread all the way *through* propositions that entail lottery-like propositions *to* ordinary propositions that entail anti-skeptical ones. Can RDS at least cut off that threat without denying closure?

BIV. Suppose that I know that p: I have hands. I know that p entails q: I am not a BIV. What is my reason for believing that p?

R_p: It looks and feels like I have hands.

If p were false, R_p would not be the case. RDS delivers the result that I know that I have hands. Do I know that q: I am not a BIV? Hawthorne’s view says yes, since it is not salient to any practical concerns. (Hawthorne can, with consistency, say no, if in weird cases it is salient because of certain practical interests, but then neither would I know that I have hands. He says feature. I say bug.) RNS says I do not know that q because if I were a BIV, I would believe, via my experience-based method, that I am not a BIV. Since RNS implies that I know that p but not q, it implies closure violations. (Nozick says feature. I say bug.) What does RDS say about knowledge that q?

Well, what is my reason—R_q—for (believing that (hereafter read)) q? Given that I deduce it from known premises, my reason R_q is simply p (and that $p \rightarrow q$).

R_q: p (and $p \rightarrow q$)

If q were false—if I were a BIV—p would not be the case, so p is actually true and conclusive for q, even if I would nonetheless *believe* I am not a BIV if it were false. Recall that RDS does not require that one not have one’s actual reasons if p were false, only that those reasons would not be the case—they would not be true. And recall that counterfactuals are not transitive. If q were false, R_q would not be the case, and if R_q (viz. p) were false, R_p would not be the case, even though, if q were false, R_p would still be the case.

A natural question at this point: What is the RDS explanation for the pull of the skeptic’s claim that S does not know they’re not a BIV, even though RDS says S does know this? Well, here Nozick was on to something. S’s *belief* is not sensitive. The fact that S

would believe they're not a BIV (even if they were a BIV) very strongly suggests some kind of epistemic predicament for S. (DeRose [31] says something similar, but because he is not thinking in terms of reasons-based sensitivity, he believes that the skeptical intuition holds sway in skeptical contexts, in which it is false that "S knows they're not a BIV." He says feature. I say bug.) In sum, RDS treats the BIV case in a straightforward and elegant way while distinguishing it from lottery cases, as desired. S knows they're not a BIV if S's reasons are conclusive. (Note that if S believes they're not a BIV *solely on the basis* that S does not appear to themselves as a BIV, S does not know because S's reason is insensitive.) Interest-relative invariantism and contextualism, by contrast, introduce *prima facie* non-epistemic factors into the question, namely, S's practical interests or features of conversational context. These seem like helpful insights, especially insofar as they permit closure maintenance and explain shifting intuitions about whether S knows, in certain cases. But they also complicate epistemological theory, in my view needlessly, and generate their own questionable implications.

I pause for just a moment to repeat the crucial features that give RDS a huge advantage over RNS. RDS changes the focus from belief sensitivity and method-relative belief sensitivity to reasons sensitivity. RDS does not require that if p were false, S would not have R (a belief or experience), but rather that if p were false, R (a proposition expressing S's actual reasons) would not be the case. And finally, RDS presents reasons as structured, so that R can be a conclusive reason for p and R* a conclusive reason for R, even if R* (for example, a basic experience) is not a conclusive reason for, i.e., does not track or isn't sensitive to, p.⁵ Dretske's examples of non-closure all hinge on experience not tracking anti-skeptical propositions, but the resources required to overturn this result are available in his own description of conclusive reasons.⁶

2.4. Other (Lottery-like) Cases

It would be pointless to try to cover every possible test case. Still, there are potentially epistemologically salient factors that distinguish not only *Safari* from *BIV*, but both of them from Hawthorne's lottery-like cases that prompt skeptical concerns. *BIV* pairs an ordinary, intuitively known proposition with a radical, global skeptical hypothesis. What are the implications of RDS with regard to a case that pairs an ordinary, intuitively known proposition with a less radical, local skeptical hypothesis? For example, what about cleverly disguised mules? Given that it is possible that the zebra-looking animal in the pen is a cleverly disguised mule (CDM), the proposition that it is not a CDM can be treated as lottery-like.

Cleverly Disguised Mules? [15] What is my reason R_p for believing that p?

p: The animal in the pen is a zebra.

R_p : The animal in the pen looks just like a zebra.

If p were false, R_p would not be the case. What is my reason for believing that q?

q: It is not a CDM.

R_q : It is a zebra. ($R_q = p$.)

If q were false, R_q would not be the case. This can all be true even though, if q were false (if it were a CDM), R_p (my experience) would still be the case. This seems to me to have been the correct result all along, Dretske [15] notwithstanding. I know both p and q, unless, of course, if p were false and the animal in the pen were not a zebra, it *would* be a CDM. There have been at least three zoos accused of displaying fake zebras already. If these occurrences become less abnormal and their fakes really can fool people, we may have to revisit knowledge of p and q.⁷ At any rate, RDS delivers a desirable result in this case, a case with a clear skeptical threat, like the BIV scenario, but also one with an implication that can be treated as lottery-like. RDS implies knowledge without denying closure.

Safari pairs an explicit lottery conclusion with a “lottery-impacted” premise. In addition, it involves beliefs about things unobserved. How does RDS respond to other cases of lottery-like propositions about things unobserved, known by induction?⁸ These being the most complicated cases, they require a somewhat extended discussion. What about, for example, *Parked Car*?

p: My car is in Lot F.

$p \rightarrow q$

q: It is not the case that my car has been stolen and is no longer in Lot F.

What is my reason for *p*?

R_p: I left my car in Lot F an hour ago.

If *p* were false, maybe *R_p* would still be the case, maybe not. If there were a fairly high incidence of car thefts in the neighborhood, then *R_p* might easily be the case if *p* were false, but I also would not know that my car is in Lot F (though I may know that it’s likely that it is).

But the high theft-incidence scenario does not address the question at issue, which presumes that it is very unlikely that the car is stolen. Assimilating lottery-like propositions such as *q* above, to explicit lottery propositions, Hawthorne challenges us to explain how they are distinct, again on pain of the threat of skepticism. He also claims that lack of knowledge of explicit lottery propositions does not “depend crucially upon the fact that the epistemic subject... has merely statistical reasons for believing that [they] will lose the lottery” (p. 9, [25]). The reasons cited on behalf of RDS in *Parked Car*—*R_p* for *p*, and *p* for *q*—are not statistical. But whether they are conclusive does sometimes depend on whether the falsity of *q* is merely statistically very, very unlikely, on the one hand, or out-of-the-ordinary (abnormal) on the other. (Whether one’s reasons are conclusive depends not on how good one thinks they are, or whether one takes one’s reasons to be merely statistical, but on their relations to objective facts.) The more abnormal it is that not-*q*—for example, that the car *was* stolen—the greater the impetus to interpret the sensitivity conditional “If *p* were false...” in a backtracking way. That is, as we imagine the event of the car’s being stolen to be increasingly bizarre, even holding the statistical likelihood constant, the reading of the conditional—If *S*’s car were not in Lot F, (it is because) *S* would not have parked it there—becomes increasingly plausible.⁹ In that case, *R_p* is conclusive—it would not be the case if *p* were false—so *S* knows both *p* and, by deduction, *q*. Since I don’t have a view on which of the two possibilities better describes a “typical” parking lot in a “typical” environment, I will sketch extreme cases as partial defense of the claims just presented.

Lottery Theft. Once every year, a car is selected to be stolen by a lottery that is operated by thieves. The otherwise nearly infallible, neighborhood-patrolling, local police and parking lot security agents across the city turn a blind eye to this one theft to appease the bandits, who otherwise would do something far worse than steal a car. Chances of any given car being stolen are 1-in-100,000 (# of cars) $\times$ 365 (# of days in a year) = 1 in 36,500,000.

Abnormal Theft. Lot F is in an extremely safe lot in an extremely safe neighborhood where crime is virtually nonexistent. Anyone wanting access to the neighborhood must enter through a gate and present identification with a chip indicating that they have no criminal record. The lot is observed both via surveillance cameras that are watched diligently by four security employees and by vigilant security guards surrounding the lot. The security company has issued an insurance policy to the lot owners, requiring payment of five times the value of any stolen car, plus a promise to fire every employee on duty during a theft. The security measures are so strict that the estimated incidence of theft per car parked per night is approximately 1 in 36,500,000.

Given the facts of the cases, the following upshots are at least defensible. One might refuse to attribute knowledge in the *Lottery Theft* version since, knowing the facts, they see

that the scenario in which the car has been stolen is something that could easily happen. In the *Abnormal Theft* version, one might at least be more open to attributing knowledge because the scenario in which the car is stolen involves an extraordinary event. (In saying “at least defensible,” “one might refuse,” and “one might be more open,” I seek only a foothold, with an aim to making the case stronger as we go.)

The likelihood is the same in both cases, so what even *could* justify describing *Abnormal Theft* as a case of knowledge but not *Lottery Theft*? In an actual lottery, when q is false and S wins the lottery, we seek no explanation for the fact. It is simply unlikely, but “someone has to win.” This point applies to *Lottery Theft*. Where q ’s being false is abnormal, we want to know what happened. It is not just unlikely—it cries out for explanation. (Did two guards outside in the lot fall asleep at the same time, with a simultaneous camera glitch? Compare with Smith’s [37] conception of justification as normic (rather than statistical) support. Our question is not justification *per se*, but whether one’s reasons are conclusive.) Given that the falsity of q is *abnormal*, then, the greater the unlikelyhood that not- q , the more plausible the backtracking reading of the relevant counterfactual conditional: If the car were not in Lot F, (it is because) it would be false that I parked it there. If p were false, R_p would not be the case; hence, R_p is conclusive. So, I know my car is in Lot F, from which I infer that it’s not stolen, which I know because if it had been, it would not be in Lot F.

Appeals to backtracking raise hackles. (I remind the reader, though, that the RDS resolutions of the BIV and CDM cases do not depend on backtracking. RDS outperforms RNS in gaining traction against the spread of the skeptical threat without denying closure.) What can be said in favor of the basic idea from the previous paragraph, namely, that the facts that a not- q event is abnormal is not sufficient to warrant backtracking, and neither is the fact that not- q is unlikely, but when these facts co-occur, backtracking readings of counterfactuals can be justified. That they are justified is an integral part of the strategy for distinguishing explicit lottery propositions from some lottery-like propositions, especially those involving induction.¹⁰

As (merely) an analogy, compare three versions of a newspaper misprint example: *Spurs*.

p : The Spurs (intentionally ambiguous, depending on your sports preferences, if any) won last night.

q : The newspaper’s report that the Spurs won last night is not a misprint.

R_p : The newspaper says the Spurs won last night.

(1) *Typical Newspaper*. It makes errors. They are abnormal, at least in the sense of involving a mistake or glitch in the reporting process. They are not random occurrences built into the reporting process. They are very unlikely, perhaps not nearly as unlikely as winning the lottery, but still unlikely, say 1 in 20,000. (2) *Less Reliable Newspaper*. It makes lots of errors. 1 out of 50 scores is misreported. But the errors are abnormal, as described under (1). The difference between the two papers is akin to the difference between a very reliable car and a less reliable one. (3) *Lottery Newspaper*. The reporting process enjoys elite-level reliability, and nary a mistake is made, with one exception. The editor rolls a pair of dice four times, and if each comes up 12 (a 1 in 20,736 occurrence), they print an incorrect winner for the relevant game. The errors are normal. They are built into the reporting process.

Do I know that p , the Spurs won last night, on the basis of R_p ? Perhaps a better question, given that there are no doubt conflicting intuitions about all of this, is whether there are epistemically salient differences between my reason R_p in the three cases, depending on which of the papers I read. In the first case, it is fairly uncontroversial that a third-person attributor would say I know, and that R_p would not be the case if the Spurs had not won.

(RDS implies that I know that p and can use that knowledge to infer and come to know that there is no misprint, even if, if there were a misprint, I might still have my reason R_p . R_p tracks p , and p tracks q , but R_p does not track q . RNS stumbles here because it ignores reasons and their recursive structure and simply asks whether I would believe that there is no misprint, even if this were false.) In the second case, a third-person attributor would at least be hesitant to attribute knowledge, citing the lesser reliability of the paper (though 98% correct is considered reliable in general)—“they might have reported that the Spurs won [R_p might still be the case—it’s not the case that R would not be the case] even if the Spurs had lost.” And the same goes for the *Lottery Newspaper*, though the thought that R_p might still be the case if p were false would have a different basis—“It’s built into their process that they print errors—unlikely, yes, but that’s how they operate.”

We come now to the question about backtracking. Hawthorne discusses the newspaper misprint example (p. 4, [25]) as yet another lottery-style case. (In that respect, the fact that RDS nicely resolves the *Typical Newspaper* variant is already a plus.) It is surely not a case of induction, and so far as I know nobody has presented it as an analogy to inductive knowledge (such as knowing that my car is now in Lot F). Therefore, the question of backtracking *per se* does not arise. Still, the newspaper case does have a *structure* that is similar to that of inductive knowledge. We tend to do an analog of backtracking automatically when the falsity of p involves something abnormal and is statistically unlikely, and not otherwise (normal but unlikely; not *very* unlikely but abnormal). For the *Typical Newspaper*: The newspaper would not have printed it if it were not true, and therefore (from the third-person perspective) S ’s reasons R_p would not have been the case. We do not think this for the *Less Reliable* and *Lottery Newspapers*: If it were false that the Spurs won last night, the newspaper might have said they did anyway [it’s not the case that R would not be the case]. This differential inclination to backtrack (or “analog-backtrack”) maps directly onto our (perhaps graded) intuitions about whether one knows. (Here, finally, is an example of the potential indeterminacy, or perhaps just underdetermination, that I mentioned in II.B.2.)

The *obvious* objection to the analogy is that, in the *Spurs* case, the fact that p , the Spurs won, causes the newspaper report, the source of R_p , whereas in the *Parked Car* case, just the opposite is true (the car’s being in Lot F does not cause me to park it there), which, in turn, is why it is true that *if it were false that the Spurs won, the newspaper would not have reported that they did*, and this has nothing to do with backtracking in inductive cases. Yes, but the example is meant only as an analogy. And, yes, there is this huge disanalogy. However, what makes the analogy worth thinking about is that the causal relation between the Spurs winning and the newspaper report forges a robust connection between the events, which is why we think if the Spurs had lost, the report would have been different. What happens when there is a built-in weakness in the otherwise robust connection, as in the *Lottery Newspaper* version of the *Spurs* case? If they had lost, we do not assume that R_p would not be the case. We know it might be. (A friend says, “Keep your betting receipt. You know they randomly report mistakes, right?”) Similarly, *Lottery Theft*, if the car were not in Lot F, we do not assume that (it is because) it was not parked there, precisely because it might simply—might *easily*—have been stolen. (A friend says, “You know about the lottery theft system, right?”) No matter how good the security is otherwise, if it is built into the security apparatus that one car will be stolen at random, there is a readily identifiable weakness in the connection between p and R_p . When we repair the connection, as in *Typical Newspaper*, we think that if p were false, it would not have been reported. (“Sorry about your bet. How much did you lose?”) In *Abnormal Theft*, there is a robust security connection, with no built-in weaknesses, between my parking in Lot F and my car’s being there, which makes it at least not unnatural to consider the idea that if my car were not there, it would be because

I did not put it there. (“They have excellent security. If the car isn’t there I would presume that he didn’t park it there.”)

Little of the immediately preceding is original thinking. It is simply a way of presenting ideas that have been floating around modal epistemology for decades. It is not my ambition to convince the reader that backtracking readings of sensitivity conditionals are correct in cases of inductive knowledge. I am content with a smaller victory, namely, that the following line of thought is at least coherent, and, if true, explanatory. There is a salient difference between cases where not-*q* is merely statistically unlikely and cases where not-*q* is abnormal; the salient difference makes attributions of knowledge that *p* and, by inference, that *q* more plausible in the latter cases; this is because if *p* were false, *R* would not be the case; and finally, the backtracking readings of sensitivity counterfactuals that support the immediately preceding claim sometimes enjoy some degree of warrant.

In summary, the line of thought positions us to take full advantage of RDS. In lottery cases, if *p* (I cannot afford a safari) were false, *R* (It is very unlikely that I will be able to afford it) would be the case, so *S* knows neither *p* nor *q* (My lottery ticket is a loser). But this does not bring full-blown skepticism in its wake. If *p* (the animal in the pen is a zebra [or the Spurs won last night, or, sometimes, my car is in Lot F]) were false, *R* (the animal looks like a zebra [or (*Typical Newspaper*) the paper said the Spurs won, or, sometimes, I parked my car in Lot F]) would not be the case. So, one can know *p* on the basis of *R*. And one can therefore know *q* on the basis of (their knowledge that) *p*—their reason for believing *q*. And, most important of all, one can know both *p* and *q*, even if *R* does not track *q*.

3. A Few Final Remarks

Above, I intimated some potentially problematic features of rival theoretical frameworks. I conclude on a similar theme.

RDS is agile and elegant, and therefore, powerful (especially in its ability to handle the problems plaguing RNS [16–18]). Besides the other reservations about safety noted above (introduction to §II), safety’s neo-Moorean strategy often strikes philosophers as lacking explanatory force. *Both p (I have hands) and q (I’m not a BIV) are safe from error! “And?”* This is not to say that safety theorists haven’t said a lot about why being safe from error is important, but even so, their support for Moore’s controversial proof brings many people up short. For those people, that the belief that *q* is safe, not at risk of being false, is simply a fact. Perhaps a newly discovered one, but still just a fact nonetheless. Why does it matter? Lots of beliefs can be safe and not even in the ballpark of knowledge (Ref. [38], for starters). How is the idea of safety connected to our understanding of knowledge? If you told the person on the street that *knowledge is true belief that is safe from error* (and then told them what that means), I suspect they’d say, “Okay, that’s true. But what is knowledge?” You might protest, starting to sound like an Introduction to Epistemology professor: *Well, the safety idea is more significant than you might think. There was this groundbreaking paper in 1963, only three pages, showing that the justified true belief account of knowledge is incorrect. “The what now? Oh, wait ‘justified belief,’ that sounds like it has to do with knowledge.” But wait, you see, safety. . .* You get the picture, and if you don’t, then maybe you’re in too deep. Tell the same person that *to know is to have a true belief with conclusive reasons*, and you’re likely to get: “Sounds about right. What are conclusive reasons?” *They’re reasons that are sensitive to the truth. They track the truth. “Okay.”*

At this point, I will be accused of soap-boxing, but I beg for just another minute on the makeshift pedestal. Try this out on the street: *Whether you know something depends on how important it is for you. “It does? How come? I mean, I can see why if something were really important I would want to double check just to be absolutely certain. Is that what you mean?” No, I’m saying if it were really important, then you didn’t know in the first place. Look,*

there's this lottery puzzle. . . "If I win the lottery I'm going on safari! I know it's a longshot but heck, you never know!"

Well, some of my colleagues think that "S knows that p" can be true in one conversational context but false in another, depending on what is salient in that context. "Sheesh, where'd they come up with that?" Because knowledge is closed under known entailment, and we need a theoretical framework that fights skepticism. "If you say so. What about that conclusive reasons idea? Does it close out entailment or whatever and, uh, attack. . ." Fight " . . . fight skepticism?" Well sure, but. . . "Cool. Later skater."

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflicts of interest.

Notes

- ¹ I am not making Williamson's [6] more challenging point that the analysis of knowledge is a degenerating research program. Even putative illuminating *but merely necessary* conditions on knowledge look baroque.
- ² I'm uncertain whether even Dretske noticed the potential power of his ideas about conclusive reasons, given that he defended non-closure [15], which is evidence that he would not have endorsed the development of his account that I present here. But one might still wonder why the consensus is that Dretske's view is basically a notational variant on Nozick's [4]. Perhaps, too (as above), because Dretske rejects knowledge closure, which invites readers to see the differences between their versions of sensitivity as insubstantial. Add to this the fact that Nozick spilled a lot of ink and generated a lot of buzz, and one explanation is that his view became the received gold standard (or easy target, depending on your allegiances), leaving further development (though certainly not discussion) of Dretske behind.
- ³ Need something other than methods be held constant, such as features of the environment? See ref. [26].
- ⁴ In still one more otherwise potentially fatal kick to a dead horse, Haicheng Zhao (pp. 389–391, [28]) argues that both Nozickean Sensitivity and RNS fail to account for lack of knowledge in a clever variant of the basic lottery. Skipping the details, "Kattie" does not know that her lottery ticket is a loser, but her belief is Nozickean-sensitive anyway. Because winners are notified when they buy their tickets, if it were false that Kattie's ticket is a loser, she would not believe it. Zhao also argues that while method-relative RNS might be invoked to solve the problem, RNS fails to account for knowledge that it is likely that one's ticket will lose. RDS has no trouble with either case. If Kattie's ticket were a winner, her reasons for believing it is a loser—the long odds—would still be the case. Kattie does not know because her belief is RDS-insensitive. If it were false that Kattie's ticket is likely to lose, her reason—the long odds—would not be the case. Kattie does know because her belief is RDS-sensitive.
- ⁵ Roush [32] and Topey [33] both incorporate recursion within a sensitivity framework to maintain closure and resolve other issues with sensitivity. In their accounts (whose differences need not detain us here), an otherwise insensitive belief can still be knowledge so long as it is reliably inferred from a sensitive belief. My version of reasons-based sensitivity, RDS, while not Nozickean in its conception of what needs to be sensitive, does speak to Nozick's ambition to present sensitivity itself as a perfectly general requirement for knowledge.
- ⁶ For reasons set out in [16], if you believe, as I (think I) do, that knowing *that one's belief that p is not false* is a more substantial cognitive achievement than knowing that *p*—that higher-level knowledge requires more than first-order knowledge—it implies very special closure violations. If you don't like that, then see *ibid.* for an alternative. Both presentations of higher-level knowledge found there—with and without closure violations—are available through RDS.
- ⁷ A contemporary real-life example. Until a few years ago, deepfakes were not easy to produce. They were possible and actual, but unusual. I suspect that then, most people would accept that one can attain knowledge from a photograph along with knowledge that what's in the photograph is not fake. ("It wouldn't look like a barn if it weren't a picture of a barn. It's a picture of a barn, and therefore not a fake picture of a barn.") We are barreling rapidly toward a deepfake-barns world where if something were not a picture of an *x* it would (or at least might) look like an *x*, assuming that the picture's creator wanted to portray *x*-ness. In our brave new world, knowledge from pictures will become increasingly rare.
- ⁸ For a fuller treatment of RDS and inductive knowledge, see ref. [17].

- ⁹ Two notes here. (a) Wallbridge (pp. 121–122, [34]) appeals to Bennett [35] in defense of backtracking readings of some counterfactuals and rightly points out that Lewis [36] does not insist they are incorrect in every case, and indeed that they are required in order to avoid miracles in counterfactual situations. The imagined scenario I will describe in the text is not quite a *miracle*. But it is weird, unusual, and not merely a statistical aberration. (b) Hawthorne himself presents a quantum desk case (pp. 4–5, [25]), where because there is some possibility that the desk has become a mere desk façade, once I know this and entertain the possibility, “I will be far less inclined to think that I know by perception whether or not this is a desk façade” (5). I think this is 100% false. I presume you agree. Why is it false? The answer is contained in Hawthorne’s description of the case—the quantum event leading to the desk façade would be “bizarre.” If the piece of furniture I’m looking at were not a desk, it would not be a desk-façade. That would be bizarre. It would be something else, and it would not look like a desk. I couple these insights in what follows.
- ¹⁰ Hawthorne (p. 19, [25]) takes Vogel’s (p. 165, [29]) Heartbreaker case, where there is strong inclination to think S does know that not all 60 golfers will ace the difficult Heartbreaker hole, to be, in fact, just a disguised lottery. Wallbridge [34] appeals to backtracking to resolve cases like the Heartbreaker on behalf of sensitivity. However, like Nozick, his application of the strategy requires backtracking such that S would have used some other belief-forming method if p were false. But is it true, for example, that to know that the sun will rise tomorrow, I would have to have learned of the imminent catastrophe if the sun were not to rise tomorrow? That is too strong a requirement for easily obtainable, ordinary inductive knowledge. RDS offers a better picture. I treat the Heartbreaker in detail in my 2025b [17].

References

- Gettier, E. Is Justified True Belief Knowledge? *Analysis* **1963**, *23*, 121–123. [CrossRef]
- Goldman, A.I. A Causal Theory of Knowing. *J. Philos.* **1967**, *64*, 357–372. [CrossRef]
- Goldman, A.I. Discrimination and Perceptual Knowledge. *J. Philos.* **1976**, *73*, 771–791. [CrossRef]
- Nozick, R. *Philosophical Explanations*; Harvard UP: Cambridge, MA, USA, 1981.
- Sosa, E. How to Defeat Opposition to Moore. *Philos. Perspect.* **1999**, *13*, 141–153. [CrossRef]
- Williamson, T. *Knowledge and Its Limits*; Oxford UP: Oxford, UK, 2000.
- Pritchard, D. *Epistemic Luck*; Oxford UP: Oxford, UK, 2005.
- Goldman, A.H. An Explanatory Analysis of Knowledge. *Am. Philos. Q.* **1984**, *21*, 101–108.
- Jenkins, C.S. Knowledge and Explanation. *Can. J. Philos.* **2005**, *36*, 137–164. [CrossRef]
- Bogardus, T.; Perrin, W. Knowledge is Believing Something Because it’s True. *Episteme* **2022**, *19*, 178–196. [CrossRef]
- Clark, M. Knowledge and Grounds: A Comment on Mr. Gettier’s Paper. *Analysis* **1963**, *24*, 46–48. [CrossRef]
- Lehrer, K.; Paxson, T. Knowledge: Undefeated Justified True Belief. *J. Philos.* **1969**, *66*, 225–237. [CrossRef]
- Harman, G. *Thought*; Princeton UP: Princeton, NJ, USA, 1973.
- Dretske, F. Conclusive Reasons. *Australas. J. Philos.* **1971**, *49*, 1–22. [CrossRef]
- Dretske, F. Epistemic Operators. *J. Philos.* **1970**, *67*, 1007–1023. [CrossRef]
- Becker, K. Dretskean Sensitivity and Higher-Level Knowledge. *Episteme* **2025**, *23*, 237–250. [CrossRef]
- Becker, K. Dretskean Sensitivity and Induction. *Synthese* **2025**, *206*, 290. [CrossRef]
- Becker, K. Why Dretskean Conclusive Reasons Are Compatible with Closure. 2026; (Manuscript under review).
- Vogel, J. Reliabilism Levelled. *J. Philos.* **2000**, *97*, 602–623. [CrossRef] [PubMed]
- Wallbridge, K. Sensitivity and Higher-Order Knowledge. *Pac. Philos. Q.* **2016**, *99*, 164–173.
- Vogel, J. Tracking, Closure, and Inductive Knowledge. In *The Possibility of Knowledge: Nozick and His Critics*; Rowman & Littlefield: Lanham, MD, USA, 1987; pp. 197–215.
- Melchior, G. Sensitivity and Inductive Knowledge Revisited. *Dialectica* **2021**, *75*, 441–462. [CrossRef]
- Kripke, S. Nozick on Knowledge. In *Philosophical Troubles: Collected Papers, Vol. 1*; Oxford UP: Oxford, UK, 2011; pp. 161–224.
- Forbes, G. Nozick on Scepticism. *Philos. Q.* **1984**, *34*, 43–52. [CrossRef]
- Hawthorne, J. *Knowledge and Lotteries*; Oxford UP: Oxford, UK, 2004.
- Mortini, D. The Explanationist and the Modalist. *Episteme* **2024**, *21*, 371–386. [CrossRef]
- Hirvela, J. Global Safety: How to Deal with Necessary Truths. *Synthese* **2019**, *196*, 1167–1186. [CrossRef]
- Zhao, H. Solving the Lottery Problem: From Modal Accounts to Explanationism. *Can. J. Philos.* **2025**, *55*, 387–404. [CrossRef]
- Vogel, J. The New Relevant Alternatives Theory. *Philos. Perspect.* **1999**, *13*, 155–180. [CrossRef]
- Feldman, R. Knowledge and Lotteries. *Philos. Phenomenol. Res.* **2007**, *75*, 211–226. [CrossRef]
- DeRose, K. Solving the Skeptical Problem. *Philos. Rev.* **1995**, *104*, 1–52. [CrossRef]
- Roush, S. *Tracking Truth: Knowledge, Evidence and Science*; Oxford UP: Oxford, UK, 2006.
- Topey, B. Saving Sensitivity. *Philos. Q.* **2022**, *72*, 177–196. [CrossRef]
- Wallbridge, K. Sensitivity, Induction, and Miracles. *Australas. J. Philos.* **2018**, *96*, 118–126. [CrossRef]
- Bennett, J. *A Philosophical Guide to Conditionals*; Oxford University Press: Oxford, UK, 2003.

36. Lewis, D. Counterfactual Dependence and Time's Arrow. *Noûs* **1979**, *13*, 455–476. [[CrossRef](#)]
37. Smith, M. *Between Probability and Certainty: What Justifies Belief*; Oxford UP: Oxford, UK, 2016.
38. Becker, K. Reliabilism and Safety. *Metaphilosophy* **2006**, *37*, 691–704. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.