Knowledge as Fact-Tracking True Belief

Fred Adams
University of Delaware
Department of Philosophy
Department of Linguistics and Cognitive Sciences
U.S.A.
fa@udel.edu

John A. Barker
Southern Illinois University at Edwardsville
Department of Philosophy
Edwardsville, IL 62026
U.S.A.
johnabarker1@cs.com

Murray Clarke
Concordia University
Department of Philosophy
1455 De Maisonneuve Blvd.
West Montreal, Quebec
H3G 1M8
Canada
murray.clarke@concordia.ca

ABSTRACT

Drawing inspiration from Fred Dretske, L. S. Carrier, John A. Barker, and Robert Nozick, we develop a tracking analysis of knowing according to which a true belief constitutes knowledge if and only if it is based on reasons that are sensitive to the fact that makes it true, that is, reasons that wouldn't obtain if the belief weren't true. We show that our sensitivity analysis handles numerous Gettier-type cases and lottery problems, blocks pathways leading to skepticism, and validates the epistemic closure thesis that correct inferences from known premises yield knowledge of the conclusions. We discuss the plausible views of Ted Warfield and Branden Fitelson regarding cases of knowledge acquired via inference from false premises, and we show how our sensitivity analysis can account for such cases. We present arguments designed to discredit putative counterexamples to sensitivity analyses recently proffered by Tristan Haze, John Williams and Neil Sinhababu, which involve true statements made by untrustworthy informants and strange clocks that sometimes display the correct time while running backwards. Finally, we show that in virtue of employing the paradox-free subjunctive conditionals codified by Relevance Logic theorists instead of the paradox-laden subjunctive conditionals codified by Robert Stalnaker and David Lewis.

Tracking Theories of Knowledge

Tracking theorists hypothesize that in order to qualify as knowledge, a true belief must *track* the relevant fact—the belief must depend in an appropriate way upon the fact that accounts for its truth. "... tracking theories see knowledge as a real relation between a believer, the truth of the relevant belief, and the environmental conditions that nomically ensure the truth of the relevant belief..."[1] We develop a tracking account according to which a true belief constitutes knowledge if and only if it is based on reasons that are *sensitive* to the fact that makes it true, that is, reasons that wouldn't obtain if the belief weren't true.[2]

Our sensitive reasons account closely resembles Fred Dretske's well-known *conclusive reasons* account (Dretske 1971). Like Dretske and other externalists, we maintain that in order to qualify as knowledge, a belief need not be justifiable from the subject's own perspective. While we focus on knowledge itself, which we call *generic* knowledge, we also investigate several species of knowledge, including some that involve justifiable belief.

Sensitivity Analyses of Knowing

Although a few tracking analyses of knowing were published in the 1960's, sensitivity analyses *per se* first appeared in Dretske 1971, Carrier 1971, Barker 1972, and Nozick 1981. These analyses were widely regarded as successfully handling myriad Gettier-type cases and lottery problems; but they suffered from difficulties of their own. Drawing inspiration from these views, we develop a sensitive reasons analysis of knowing that appears to overcome these difficulties. Like Dretske's and L. S. Carrier's analyses, ours explicitly involves the subject's *reasons* for believing the relevant proposition. Like John Barker's

---

[1] Adams and Clarke 2005, p. 207.
[2] Sensitivity terminology was introduced in Barker 1972, p. 317, and tracking terminology was introduced in Nozick 1981, p. 211. In the latter work, the term 'sensitivity' was employed in a wider sense, but it is now commonly used in the epistemological literature in its original sense. See, e.g., Shope 1983 for discussion of analyses that employed the notion of sensitivity and related notions.

analysis, ours employs paradox-free subjunctive conditionals, validates the epistemic closure thesis that correct inferences from known premises yield knowledge of the conclusions, and accounts for correct inferences involving intermediate conclusions. And like Robert Nozick's analysis, ours is designed to accommodate knowledge-affording methods of belief formation and knowledge of necessary truths.

Sensitivity and Closure

Dretske and Nozick argued that closure should be rejected and that, in virtue of invalidating closure, their sensitivity analyses undermined skepticism. Dretske's rejection of closure, however, was motivated by the plausibility of his sensitivity requirement rather than by a desire to discredit skepticism:

> As a historical footnote, I wasn't led to deny closure because it represented a way around skepticism. I was led to it because it was a result of what I took to be a plausible condition on the evidence (justification, reasons) required for knowledge. If your reasons for believing P are such that you might have them when P is false, then they aren't good enough to know that P is true. You need something more. That is why you can't know you are going to lose a lottery just because your chances of losing are 99.99 percent. Even with those odds, you still might win (someone with those odds against him will win). That is why you can't learn—can't come to know—that P is true if all you have to go on is the word of a person who might lie about whether or not P is so. This is just another way of saying that knowledge requires reasons or evidence (in this case, testimony) you wouldn't have if what you end up believing were false. You can learn things from people, yes, but only from people who wouldn't say it unless it were true. (2005, p. 43-44)

The thesis that rejection of closure is mandated by acceptance of a sensitivity analysis has impeded adoption of such analyses, with many justificationists viewing such a rejection as a *reductio ad absurdum* of the analyses. We argue that our sensitivity analysis implies that closure is valid and that, ironically, justificationist analyses typically imply that it is invalid.

Inferential knowledge is acquired via closure only when the premises are known to be true. Ted Warfield and Branden Fitelson have argued convincingly that inferential knowledge is sometimes acquired when the

premises are false and therefore not known to be true.[3]   We show how our sensitivity analysis can accommodate cases in which knowledge is acquired via inference from false premises and therefore independently of closure.

Lotteries, Untrustworthy Informants, and Necessary Truths

As the above quotation from Dretske suggests, sensitivity analyses of knowing can easily handle lottery cases, an achievement that few competing analyses can replicate. It's quite plausible that if you believe that you'll lose a fair lottery just because your chances of losing are 99.99 percent, then even if the belief is true, reliable, safe, and justifiable, it doesn't qualify as knowledge— if it hadn't been the case that you'd lose, your reasons for believing you'd lose would have been the same. Moreover, sensitivity analyses can readily explain how information provided by deceivers and other untrustworthy sources of communication, such as dysfunctional clocks and measuring instruments, can give rise to true beliefs that fail to qualify as knowledge. We discuss several putative counterexamples to sensitivity accounts devised by Tristan Haze, John Williams, and Neil Sinhababu, and we argue that these cases involve true beliefs that don't constitute counterexamples to our analysis.[4]

While Dretske and Carrier restricted their sensitivity analyses to empirical truths, Nozick extended his to necessary truths.[5] Unlike Barker, however, Nozick employed subjunctive conditionals that are subject to paradoxes of implication. He struggled to cope with the counterintuitive properties of such conditionals, especially the infamous paradox that a necessarily false proposition implies every proposition. We argue that by employing paradox-free subjunctive conditionals, our sensitivity analysis can readily accommodate knowledge of necessary truths.

---

[3] Warfield 2005 and Fitelson 2010.

[4] Haze 2015 and Williams and Sinhababu 2015.

[5] Dretske, however, noted that in some cases conclusive reasons are necessarily or logically conclusive, that is, demonstrable on purely logical or definitional grounds; see Dretske 1971, p. 12.

Sensitive Reasons, Epistemic Basing, and Knowledge Efficacy

Our analysis of knowing centers on the notion of *sensitive reasons* for believing:

SensitiveReasonsAK (Sensitive Reasons Analysis of Knowing): S knows that p iff: (i) it is the case that p; (ii) S believes that p on the basis of reasons, R; and (iii) R's being the case is *sensitive* to p's being the case, i.e., if p weren't the case, then R wouldn't be the case.

The following plausible thesis helps clarify the familiar notion of a subject's believing something on the basis of reasons:

EpistemicBasing: S believes that p on the basis of a reason, R, iff: either (i) R is among S's reasons for believing that p, and R consists of one or more experiential states of S; or (ii) S's believing R to be the case is among S's reasons for believing that p, and S knows that R is the case.

Clause (i) specifies that beliefs can be based on reasons that consist of experiential states of the subject, such as sensory appearings and intellectual seemings. Clause (ii), which is plausible in its own right, is supported by a special case of the following plausible principle:

KnowledgeEfficacy: If S does something because S believes that p, then if S knows that p, S does it because p; in other words, if S's believing that p is among S's reasons for doing something, then if S's belief qualifies as knowledge, the fact that p is among S's reasons for doing it.[6]

To get from Athens to Larissa, Sophia took A1 because she believed it to be the route to Larissa on the basis of MapQuest's saying so, and her belief qualified as knowledge. KnowledgeEfficacy has the plausible consequence that she took A1 because it was the route to Larissa. Clause (ii) receives strong support from a special case of KnowledgeEfficacy that involves an act of believing: if Sophia believed that A1 was the route to Larissa because she

---

[6] See Barker and Adams 2010 and 2012 for discussion of EpistemicBasing and KnowledgeEfficacy.

believed that MapQuest said so, and she knew that MapQuest said so, then she believed that A1 was the route to Larissa because MapQuest said so. KnowledgeEfficacy is a highly plausible principle that helps explain why knowledge is more valuable than true belief that doesn't constitute knowledge. According to KnowledgeEfficacy, possession of knowledge enables one to do things for reasons that are facts. This doesn't hold for beliefs that don't constitute knowledge, even for beliefs that are true, reliable, safe, and justifiable. Sophia believed that she would lose a fair lottery because her chances of losing were 99.99 percent, and she bought a cheap car because she believed she would lose. Her belief was true, reliable, safe, and justifiable, but she didn't know that she would lose, and she didn't buy a cheap car because she would lose.[7]

Sensitive Reasons and Conclusive Reasons

Our sensitive reasons analysis of knowing is very similar to Dretske's *conclusive reasons* analysis. On Dretske's view, a subject's reasons for believing that p constitute *conclusive* reasons if and only if they wouldn't obtain unless p were the case.[8] We will follow the usual practice of equating subjunctive conditionals involving 'unless' with corresponding subjunctive conditionals involving 'if-then'. Accordingly, conclusive reasons can be construed as sensitive reasons: a subject's reasons, R, for believing that p are *conclusive* if and only if they are sensitive, i.e. if and only if they are such that if p weren't the case, then R wouldn't be the case.

Dretske formulated his analysis of knowing in terms of the subject's having conclusive reasons for believing the relevant proposition:

> S has conclusive reasons, R, for believing P iff: (A) R is a conclusive reason for P ..., (B) S believes, without doubt, reservation, or question, that P is the case and he believes this on the basis of R, (C) (i) S knows that R is the case or (ii) R is some experiential state of S (about which it may not make sense to suppose that S knows that R is the case; at least it no longer makes much sense to ask how he

---

[7] Some theorists would contend that Sophia did know that she would lose the lottery (see, e.g., Lycan 2006, Note 23). Such theorists, however, would need to reject KnowledgeEfficacy or to accept the implausible proposition that Sophia bought a cheap car because she would lose.

[8] See, e.g., Dretske 1971, pp. 1 ff.

> knows). With only minor embellishments, to be mentioned in a moment, I believe that S's having conclusive reasons for believing P is both a necessary and a sufficient condition for his knowing that P is the case. The appearance of the word 'know' in this characterization (in (Ci)) does not render it circular as a characterization of knowledge since it can be eliminated by recursive application of the three conditions until (Cii) is reached. (1971, pp. 12-13)[9]

Thus, Dretske's analysis in effect incorporates both SensitiveReasonsAK and EpistemicBasing. (Expressing the latter as a separate thesis makes it clear that opponents of sensitivity analyses can accept it without accepting SensitiveReasonsAK.)

Knowledge via Experience

Clause (i) of EpistemicBasing enables SensitiveReasonsAK to account for acquisition of knowledge by means of experience. During a guided tour of the London zoo accompanied by his mother, Sophia, Jimmy sees a zebra-looking animal, X, and believes it to be a zebra on the basis of its appearing to him that X is a zebra. X is in fact a zebra, and he wouldn't have this experience if X weren't a zebra. Since his belief is based on sensitive reasons, SensitiveReasonsAK has the plausible consequence that the belief qualifies as knowledge. Unbeknownst to Jimmy, Charles Darwin recounted that a zebroid, i.e., a zebra hybrid, was once on display at the London zoo.[10] Nevertheless, there are no zebroids anywhere in Jimmy's environment, and this is fortunate, for such hybrids might look enough like zebras to fool him. Sophia also sees X and comes to know that X is a zebra on the basis of a sensitive reason consisting of its appearing to her that X is a zebra. She is a zoologist who is quite familiar with the visible differences between zebras and zebroids, and if X had been a zebroid, she wouldn't have been fooled. As we'll discuss later, Sophia is in a position to acquire what we'll call *contrastive* knowledge that X is a zebra *rather than* a zebroid.

---

[9] Dretske 1971, pp. 12-13. The "minor embellishments" referred to in this passage aren't relevant to present concerns.
[10] See, e.g., http://www.ebooklibrary.org/articles/zonkey

Generic Knowledge and Incontrovertible Knowledge

A skeptic might argue that Jimmy's belief that X is a zebra fails to qualify as genuine knowledge. What Jimmy sees could be a zonkey—a zebra/donkey hybrid—that appears to him to be a zebra. A zonkey that looked like a zebra was recently displayed in the Taigan zoo in Crimea.[11] Since Jimmy's visual experience wouldn't rule out X's being a zonkey, Jimmy doesn't acquire knowledge that X is a zebra. This skeptical argument is fallacious. If what Jimmy sees weren't a zebra, it might be a giraffe, a wildebeest, or some other non-zebra; but it wouldn't be a non-zebra that appears to him to be a zebra. Hence, even though what he sees *could* be a zonkey that appears to him to be a zebra, it *would* not be one. Granted, Jimmy lacks what we'll call *incontrovertible* knowledge:

IncontrovertibleKnowledge: S has incontrovertible knowledge that p iff: (i) S knows that p on the basis of reasons, R, and (ii) for every q that is such that q could be the case if p weren't the case, R wouldn't be the case if q were the case.

Although incontrovertible knowledge is an important species of knowledge, it should not be confused with knowledge itself, or *generic* knowledge.

Generic Knowledge and Contrastive Knowledge

While visiting the Taigan zoo, Jimmy and Sophia see a zebra-looking animal, Y, and believe it to be a zebra on the basis of its appearing to them that Y is a zebra. Jimmy's having this experience, however, isn't a sensitive reason for his belief. Although Y is a zebra, there is a zonkey nearby, and if what he sees weren't a zebra, it might be a zonkey that appears to him to be a zebra. Since his belief that Y is a zebra isn't based on a sensitive reason, SensitiveReasonsAK does not have the implausible consequence that the belief constitutes knowledge. Sophia also sees Y and believes it to be a zebra on the basis of its appearing to her that Y is a zebra. If what she sees weren't a zebra,

---

[11] See, e.g., http://abcnews.go.com/International/introducing-latest-zonkey-zebradonkey-mix/story?id=24944710 .

it might be a zonkey, but it wouldn't appear to her to be a zebra—in virtue of her expertise, she wouldn't be fooled if it were a zonkey. Since Sophia believes that Y is a zebra on a basis of a sensitive reason, SensitiveReasonsAK has the plausible consequence that her belief constitutes knowledge.

In both the London zoo case and the Taigan zoo case, Sophia is in a position to acquire what we'll call *contrastive* knowledge. Suppose that in each case Sophia believes that the animal she sees is a zebra *rather than* a zonkey. Since she is very familiar with the visible differences between zebras and zonkeys, the animal's appearing to her to be a zebra constitutes not only a sensitive reason for her believing that it's a zebra, but also a sensitive reason for her believing that it's not a zonkey. Consequently, her belief qualifies as *contrastive* knowledge, that is, as knowledge that the animal is a zebra *rather than* a zonkey. Her visual experience functions as a *differentiator* that enables her to distinguish between the two possibilities, the animal's being a zebra and its being a zonkey. Contrastive knowledge can be explicated as follows:

ContrastiveKnowledge: S knows that x is A rather than B iff x's being A entails x's not being B, and S believes that x is A and not B on the basis of a *contrastively sensitive* reason, R, i.e., (i) if it weren't the case that x is A, then R wouldn't be the case, (ii) if it were to be the case that x is B, then R wouldn't be the case, and (iii) R doesn't entail either x's being A or x's not being B.[12]

On our view even generic knowledge can be said to be implicitly contrastive. For example, in the London zoo case, Jimmy is in a position to acquire contrastive knowledge that X is a zebra rather than not a zebra, for its appearing to him that X is a zebra can qualify as a contrastively sensitive reason—if it weren't the case that X is a zebra, then he wouldn't have this experience, and if it were the case that X isn't a zebra, then he wouldn't have the experience. Owing to her expertise, however, Sophia is in a position to acquire explicitly contrastive knowledge. She can acquire contrastive knowledge that X is a zebra rather than a zonkey, for its appearing to her that X is a zebra can qualify as a contrastively sensitive reason—if it weren't the case that X is a zebra, then she wouldn't have this experience, and if it were the case that X is a zonkey, then she wouldn't have the experience.

---

[12] See Barker and Adams 2010 for a more precise explication of contrastive knowledge and for discussion of the function of Clause (iii).

Knowledge via Closure

Clause (ii) of EpistemicBasing enables SensitiveReasonsAK to imply the plausible thesis:

KnowledgeClosure: If S knows that p and S knows that p implies q, then if S correctly infers that q from the premises that p and that p implies q, then S knows that q.[13]

In the London zoo case, Jimmy knows that X is a zebra on the basis of its appearing to him to be one. Although he doesn't know what wildebeests look like, he knows that X's being a zebra implies its not being a wildebeest, and he infers that X isn't a wildebeest. His believing X to be a zebra is among his reasons for believing his conclusion, and his believing X to be a zebra qualifies as knowledge. Hence, Clause (ii) of EpistemicBasing implies that he believes that X isn't a wildebeest on the basis of a reason consisting of X's being a zebra. X's being a zebra constitutes a sensitive reason for his belief that X isn't a wildebeest, for if X were a wildebeest, then it wouldn't be a zebra. SensitiveReasonsAK has the consequence that Jimmy knows that X is not a wildebeest. Thus, EpistemicBasing and SensitiveReasonsAK together imply KnowledgeClosure.

Although Jimmy doesn't know what zonkeys look like, he knows that X's being a zebra implies its not being a zonkey, and he infers that X isn't a zonkey. Even though X wouldn't look like a zebra to him if it were a wildebeest, it might look like a zebra to him if it were a zonkey. Nevertheless, EpistemicBasing and SensitiveReasonsAK imply that Jimmy knows that X isn't a zonkey. Since his believing X to be a zebra is among his reasons for believing his conclusion and since he knows that X is a zebra, EpistemicBasing implies that he believes that X isn't a zonkey on the basis of a reason consisting of X's being a zebra. X's being a zebra constitutes a sensitive reason for his belief that

---

[13] See Barker and Adams 2010 for an extensive treatment of closure and for references to relevant literature. We subscribe to the view proffered in Barker 1972, pp. 320 ff. that knowledge-affording correct inferences containing intermediate conclusions involve conjunctions of premise-propositions and conclusion-propositions in ways that ensure that intermediate conclusions qualify as known, and we plan to develop this view in detail in future papers.

X isn't a zonkey, for if X were a zonkey, then it wouldn't be a zebra. SensitiveReasonsAK has the consequence that Jimmy knows that X is not a zonkey.

Some theorists might argue that this consequence is unacceptable, and that KnowledgeClosure must therefore be rejected as false. Unlike Sophia, Jimmy cannot tell the difference between zebras and zonkeys. If what Jimmy sees were a zonkey, it might appear to him to be a zebra and this might result in his believing it to be a zebra. This argument, however, fails to distinguish between knowledge itself, which we've called *generic* knowledge, and the species of knowledge we've called *contrastive* knowledge. Jimmy knows that X isn't a zonkey on the basis of a reason consisting of X's being a zebra, something that wouldn't be the case if X were a zonkey. Nevertheless, unlike Sophia, he is in no position to acquire *contrastive* knowledge that X is a zebra *rather than* a zonkey, for X's appearing to him to be a zebra cannot serve as a differentiator enabling him to distinguish between the two alternatives, X's being a zebra and X's being a zonkey.

Justifiable Knowledge and Tenable Knowledge

Akira, Jimmy's tour guide at the London zoo, is a retired philosophy professor with an eccentric sense of humor. She tells the tour group that X is a donkey that has been disguised to look like a zebra. She adds that when the zoo's zebra died unexpectedly, zoo officials drew inspiration from zookeepers in Gaza and arranged for the disguised donkey to function as a temporary replacement.[14] Jimmy, who doesn't hear what Akira says, continues to believe that X is a zebra on the basis of its appearing to him to be a zebra. As SensitiveReasonsAK implies, his belief still constitutes generic knowledge - there are no zebra look-alikes anywhere in the vicinity, and what he sees wouldn't appear to him to be a zebra if it weren't one.

Jimmy's belief also constitutes what we'll call *justifiable* knowledge:

JustifiableKnowledge: S has *justifiable* knowledge that p iff: (i) S knows that p on the basis of reasons, R, and (ii) S *justifiably* believes that p on the basis of R,

---

[14] See, e.g., http://www.nbcnews.com/id/33278616/ns/technology_and_science-science/t/gaza-zoo-replaces-zebras-painted-donkeys/#.V3_kf_krLX4

i.e., S's believing p on the basis of R is more reasonable from S's own perspective than withholding judgment regarding p or disbelieving p.[15]

Possessing a justifiable belief can facilitate metacognitive self-monitoring, as it is an *internal* state of the subject, that is, a state that is of its very nature normally accessible to the subject via reflection. Moreover, Jimmy's belief constitutes what we'll call *tenable* knowledge, a species of knowledge that involves *indefeasibly justifiable* belief. A belief that is indefeasibly justifiable is one whose status as justifiable is immune to destruction by acquisition of additional justifiable true beliefs. Tenable knowledge can be defined as follows:

TenableKnowledge: S has *tenable* knowledge that p iff: (i) S knows that p on the basis of reasons, R; (ii) S's believing that p on the basis of R is *indefeasibly justifiable*, i.e., for every true proposition q1 there is a true proposition q2 that is such that if: (a) S justifiably believed both q1 and q2, (b) S continued to believe that p solely on the basis of R, and (c) p remained true, then S would still justifiably believe p on the basis of R.[16]

Akira's saying that X was a disguised donkey constitutes what can be called a *menacer* of Jimmy's justifiable belief that X is a zebra, for if he were to learn that she said this, his belief might no longer be justifiable. But there is a true proposition, viz. she was joking, that constitutes what can be called a *protector* of his justifiable belief—if he justifiably believed she was joking, then the status of his justifiable belief that X is a zebra would be unaffected, and he could continue to believe that X is a zebra solely on the same basis as before. Thus, one has tenable knowledge that p if and only if one knows that p on the basis of reasons that are such that for every menacer of one's justifiably believing p on the basis of these reasons, there is a protector that can nullify the adverse

---

[15] The notion of reasonableness employed herein concerns attaining the goal of believing p if and only p were to be the case. The notion of justifiable belief employed herein, which relates closely to notions employed in Chisholm 1966, is an internalist notion that should not to be confused with the externalist notion introduced in Goldman 1979.

[16] Cf. Barker 1976. We think Barker's account and similar accounts (e.g., Klein 1981) should be construed as explications of tenable knowledge rather than of generic knowledge. The clause '(c) p remained true' is designed to handle a counterexample to Barker's account devised by Robert Shope (see Shope 1983, p. 73).

effects of the menacer without supplying one with new or supplemented reasons for believing p.

Indefeasibly Justifiable Beliefs and Generic Knowledge

Indefeasibly justifiable beliefs are more valuable than mere justifiable beliefs, for they are guaranteed to be true and guaranteed to be based on sensitive reasons. Assume that one's believing p on the basis of R is indefeasibly justifiable but p is false. Then there is a true q_1, viz. not-p, that is such that if one were to acquire a justified true belief that q_1, the justificational status of one's believing p would be destroyed. Assume instead that one's believing p on the basis of R is indefeasibly justifiable but R isn't a sensitive reason for believing p. Let q_1 be the proposition that R might be the case even if p weren't the case, i.e., it's false that if p weren't the case, then R wouldn't be the case. If one justifiably believed that R might be the case even if p weren't the case, then the justificational status of one's believing p would be destroyed—one's believing p on the basis of R would no longer be justifiable, for withholding judgment regarding p would be at least as reasonable from one's own perspective as believing p, and there is no true proposition q_2 that would nullify the adverse effect of q_1 without supplying one with new or additional reasons for believing p.

Thus, given that one possesses an indefeasibly justifiable belief, it follows that one's belief is based on sensitive reasons, and SensitiveReasonsAK implies that it qualifies as knowledge. The strategy of construing knowledge as some kind of indefeasibly justifiable belief has been very popular, but most theorists who have pursued this strategy have rejected tracking analyses. In our opinion, the above result indicates that these theorists haven't been completely "off track," for their main mistake has been confusing tenable knowledge with knowledge itself.[17]

---

[17] We conjecture that many of the numerous indefeasibly justifiable true belief accounts of knowledge imply that knowing involves believing on the basis of sensitive reasons.

JustificationClosure and Skepticism

Upon learning what Akira said, Jimmy googles "zebra donkey Gaza" and discovers that she is telling the truth about the disguised donkey incident. Nevertheless, he stubbornly continues to believe that X is a zebra. SensitiveReasonsAK implies that the belief still qualifies as generic knowledge, despite no longer qualifying as justifiable knowledge or as tenable knowledge. Knowing that X's being a zebra implies its not being a disguised donkey, Jimmy infers that X isn't a disguised donkey. KnowledgeClosure implies that his conclusion belief constitutes generic knowledge, but does not imply that it constitutes justifiable knowledge.

Suppose that a skeptic claims that unjustifiable beliefs cannot qualify as knowledge, and proffers the following argument: Jimmy's conclusion belief is merely pseudo-knowledge. Edmund Gettier was correct in endorsing the thesis:

JustificationClosure: "... for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q". (1963, p. 121)

Since Jimmy's conclusion belief isn't justifiable, JustificationClosure implies that his premise belief that X is a zebra isn't justifiable either. Consequently, besides not knowing that X isn't a disguised donkey, he doesn't even know that X is a zebra. Hence, SensitiveReasonsAK must be rejected as unsatisfactory.

This argument is fallacious. JustificationClosure is a stronger thesis than Gettier needed for devising his counterexamples to the justified true belief analysis of knowing. Consider the following weaker thesis:

JustificationClosure*: For at least some proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q.

Suppose that Smith, who justifiably believes p, that Jones owns a Ford, employs deduction to acquire a justifiable belief that q, that either Jones owns a Ford or Brown is in Barcelona. (Cases of this type are allowed by JustificationClosure*.) Suppose also that p is false but q is true because Brown happens to be in Barcelona. Smith justifiably believes the true proposition, q,

but his belief fails to constitute knowledge. Thus, the justified true belief analysis of knowing is subject to counterexamples, but the strong version of JustificationClosure isn't needed to establish that this is so—the weaker and more plausible thesis, JustificationClosure*, is adequate.

This result is fortunate for Gettier—as the following considerations indicate, JustificationClosure seems to be invalid. Harry and Sally, whose favorite painter is Johannes Vermeer, justifiably believe p, that the painting they are viewing in a billionaire's mansion is a Vermeer, on the basis of its looking like a Vermeer to them. Harry deduces q, that the painting is not a fake like those created by the famous art forger, Henricus van Meegeren, whose fakes fooled even the best art experts.[18] Sally, however, withholds judgment regarding q, and when Harry criticizes her for being too cautious, she resists changing her mind because she knows that neither of them is capable of distinguishing between a Vermeer and a van Meegeren-like fake. Unlike JustificationClosure*, JustificationClosure has the consequence that Harry *justifiably* believes q, that is, that his believing q is more reasonable from his own perspective than withholding judgment regarding it or disbelieving it. Clearly, for both Harry and Sally, withholding judgment regarding q would be at least as reasonable from their own perspectives as believing it or disbelieving it. Hence, JustificationClosure is invalid.

Defenders of justificationist analyses have often faulted sensitivity analyses for invalidating theses like KnowledgeClosure. As we have seen, our sensitivity analysis validates KnowledgeClosure. Ironically, however, justificationist analyses are apparently incapable of validating KnowledgeClosure without the help of JustificationClosure, which the above considerations show to be invalid.

Knowledge via Falsehood

During a visit to the Marah Land zoo in Gaza, Alima sees a zebra-looking animal, Z, and believes it to be a zebra on the basis of its appearing to her to be a zebra. Her having this experience, however, isn't a sensitive reason for her belief, for Z is a donkey cleverly disguised to look like a zebra. Alima's father, Malik, also sees Z and believes it to be a zebra on the basis its appearing to him to be a zebra, an experience that doesn't qualify as a sensitive reason for his

---

[18] See, e.g., https://en.wikipedia.org/wiki/Han_van_Meegeren.

belief. Having heard a rumor that there was a zonkey nearby, Malik draws the conclusion that Z isn't a zonkey from the premise that Z is a zebra. Being a zoologist who is very familiar with the visible differences between zebras and zonkeys, he is capable of visually distinguishing between zebras and zonkeys. When he infers that Z isn't a zonkey, Z's appearing to him to be a zebra becomes one of his reasons for believing that Z isn't a zonkey, and becomes part of the basis of this belief.[19] This experience constitutes a sensitive reason for his belief, for if Z were a zonkey, it wouldn't appear to him to be a zebra. SensitiveReasonsAK and EpistemicBasing have the plausible consequence that Malik knows that Z isn't a zonkey. Thus, he acquires a knowledge-qualifying true belief via an inference from the false premise that Z is a zebra.

Ted Warfield has argued that inferential knowledge can be acquired from false premises, and has devised several cases to show that this is so. Branden Fitelson has strengthened Warfield's argument by modifying one of the cases to make the falsity of the relevant premise an indispensable aspect of the case. The above case can be strengthened along these lines by adding the supposition that if Z had been a zebra, the zonkey rumor wouldn't have been started and Malik wouldn't have inferred that Z wasn't a zonkey.

We'll focus on Warfield's fifth case (which Fitelson also focused on):

Example #5. I have a 7pm meeting and extreme confidence in the accuracy of my fancy watch. Having lost track of the time and wanting to arrive on time for the meeting, I look carefully at my watch. I reason: 'It is exactly 2:58pm; therefore I am not late for my 7pm meeting'. Again I know my conclusion, but as it happens it's exactly 2:56pm, not 2:58pm.

We'll refer to the relevant individual as *Ted*. Ted believes that the time is exactly 2:58pm on the basis of his watch's saying so, and he infers that he isn't late for his 7pm meeting. His premise belief is false, for the time is exactly 2:56pm. When Ted infers that he isn't late for the meeting, the watch's saying that it is exactly 2:58pm becomes one of his reasons for believing that he isn't late, and becomes part of the basis of this belief. The watch is keeping approximately correct time, and it wouldn't say that it is exactly 2:58pm if Ted were late. Consequently, his watch's saying that it is exactly 2:58pm qualifies as

---

[19] We do not maintain that reasons for believing the premises invariably constitute reasons upon which the conclusion belief is based.

a sensitive reason for his belief. SensitiveReasonsAK and EpistemicBasing have the plausible consequence that Ted knows that he isn't late. Thus, our analysis of knowing can account for the intuition that Ted's inference provides him with knowledge of his conclusion even though his premise is false. Since the epistemic status of his belief that he isn't late derives from his having a sensitive reason for believing this rather than from his false premise, it seems more appropriate to say that he acquires the knowledge *via* a falsehood than to say that he acquires it *from* a falsehood.

The soundness of this argument depends upon two suppositions that do not explicitly appear in Warfield's description of the case: (1) When Ted infers that he isn't late for the meeting, the watch's saying that the time is exactly 2:58pm becomes one of his reasons for believing that he isn't late, and becomes part of the basis of this belief; and (2) The watch is keeping approximately correct time, and it wouldn't say that it is exactly 2:58pm if Ted were late. Since each one of these suppositions seems quite plausible, we feel confident that the argument is sound. Nevertheless, while most theorists would probably accept (1), at least a few may reject (2). Accordingly, we proffer the following considerations to enhance the argument.

Suppose that in addition to his 7pm meeting, Ted also has a 2:59pm meeting. Ted believes that the time is exactly 2:58 pm on the basis of his watch's saying so, and he infers that he isn't late for his 2:59 pm meeting. His premise belief is false, for the time is exactly 2:56 pm. When he infers that he isn't late for his 2:59pm meeting, the watch's saying that the time is exactly 2:58pm becomes one of his reasons for believing that he isn't late for the meeting, and becomes part of the basis of this belief. It's implausible that the belief qualifies as knowledge, and it's implausible that it is based on a sensitive reason. Even though the watch is keeping approximately correct time, it might be the case that the watch says that it is exactly 2:58pm even if he were late for his 2:59pm meeting; in other words, it is false the watch wouldn't say this if he were late. Consequently, SensitiveReasonsAK doesn't have the implausible consequence that Ted knows that he isn't late for the meeting.

The case for accepting the argument as sound can be further enhanced by consideration of contrastive knowledge. Branden, a friend of Ted's, hears someone say that the time is 7pm, and he comes to believe that it is approximately 3pm on the basis of his watch's saying that it is 3pm. He then infers that the time is approximately 3pm rather than 7pm on the basis of what his watch says. His premise belief is true, for the time is approximately 3pm.

The watch is keeping approximately correct time, and it wouldn't say that the time is 3pm if it weren't approximately 3pm, and it wouldn't say that it is 3pm if it were approximately 7pm. Consequently, his watch's saying that it is 3pm qualifies as a sensitive reason for his conclusion belief. SensitiveReasonsAK and EpistemicBasing have the plausible consequence that Branden knows that the time is approximately 3pm rather than 7pm.

Tax-Lawyer Cases

As we remarked previously, sensitivity theories typically have little difficulty explaining how information provided by deceivers and other undependable sources, such as dysfunctional clocks and measuring instruments, can produce true beliefs that fail to qualify as knowledge. We examine two recent attempts to discredit sensitivity theories by means of cleverly constructed cases involving flawed communications.

Tristan Haze has devised a pair of cases involving true statements about tax law made by his neighbor. We'll refer to the neighbor as *Norman* and to Haze himself as *Tristan*. In the first case, Norman is a trustworthy tax-lawyer; Tristan, however, distrusts lawyers and harbors a delusional belief that Norman is a divine oracle:

> I have a deep-seated, counterfactually robust delusional belief that my neighbour is a divine oracle. He is actually a very reliable and truthful tax-lawyer. There is a point about tax law he has always wanted to tell me, p. One day, he tells me that p, and I believe him, because I believe he is a divine oracle. I would never believe him if I knew he was a lawyer, being very distrustful of lawyers. (Haze 2015, p. 310)

Haze argued that Tristan's belief fails to qualify as knowledge. We disagree with this verdict, but we readily acknowledge that the epistemic status of the belief is far from ideal. Given that Norman is a reliable and truthful tax-lawyer who tells Tristan something true about tax law, viz. p, it's plausible that Norman wouldn't say that p if it weren't true. Hence, Tristan's belief is based on sensitive reasons, and SensitiveReasonsAK has the plausible consequence that it constitutes *generic* knowledge. It's also plausible that the belief constitutes *justifiable* knowledge, for in spite of his delusions, his believing p is more reasonable *from his own perspective* than withholding judgment regarding it or

disbelieving it. But his belief clearly fails to constitute *tenable* knowledge, for there is a true proposition, viz. that Norman is a lawyer and not a divine oracle, that could destroy his belief's status as justifiable—the proposition is a menacer of his justifiable belief for which there is no protector. If Tristan justifiably believed the proposition, he would no longer justifiably believe p on the basis of Norman's saying that p, and there is no other true proposition that would protect his justifiable belief without supplying him with new or additional reasons for believing p. Thus, Tristan's belief fails to qualify as tenable knowledge, and this seems to account for Haze's intuition that it doesn't qualify as knowledge.

In Haze's second case Norman is a strange, deceitful tax-lawyer who asserts a true proposition about tax law along with five false ones:

> My neighbour is a tax lawyer. Here, unlike in the previous counterexample, I have no delusional belief. It is my neighbour who is the strange one: for years, he has intently nurtured an eccentric plan to get me to believe the truth about whether p, where p is a true proposition of tax law, along with five false propositions about tax law. His intention to do this is very counterfactually robust. He moves in next door and slowly wins my trust. One day, he begins to regale me with points of tax law. He asserts six propositions: p and five false ones. I believe them all. (Haze 2015, p. 310)

Tristan believes that p on the basis of reasons, R, that include Norman's saying that p. In view of the fact that Norman is an extremely untrustworthy informant, it is initially plausible that R doesn't qualify as a sensitive reason for the belief—if p were false, it still might be that case that R, for it still might be the case that Norman asserts that p anyway. To flesh out the case, let's assume that p is the proposition that many billionaires are not legally obligated to pay any income taxes, and let's assume that the five lies that Norman utters are equally disturbing propositions about tax law intended to shock Tristan. Now if p were false (say, because the relevant "loopholes" in the tax laws were recently closed), what would Norman say? It seems more likely that he would tell another lie by saying that p than that he would tell the truth by saying that not-p. Consequently, SensitiveReasonsAK has the plausible consequence that Tristan's belief that p fails to qualify as knowledge.

What are we to make of Haze's stipulation that Norman's intention to carry out his eccentric plan "...is very counterfactually robust?" (p. 310) We take this to mean that if p were false, Norman would still attempt to get Tristan to

believe that p along with the five lies, for this would be compatible with Norman's untrustworthy character. On this interpretation, however, R would still be the case if p were false, and therefore R wouldn't constitute a sensitive reason for the belief. But perhaps the stipulation means that if p were false, Norman would attempt to get Tristan to believe that not-p along with the five lies. On this interpretation, R wouldn't be the case if p were false, and therefore R would qualify as a sensitive reason for the belief. But this stipulation would make the case incoherent, for it would be incompatible with Norman's untrustworthy character. Norman is a deceitful person who cannot be depended upon to tell the truth. To paraphrase Dretske, you can learn things from lawyers, yes, but only from lawyers who wouldn't say it unless it were true.[20]

Backward Clock Cases

John Williams and Neil Sinhababu have devised a case involving a strange clock designed by a cult to run backwards during certain "cursed" hours:

> You habitually nap between 4 pm and 5 pm. Your method of ascertaining the time you wake is to look at your clock, one you know has always worked perfectly reliably. Unbeknownst to you, your clock is a special model designed by a cult that regards the hour starting from 4 pm today as cursed, and wants clocks not to run forwards during that hour. So your clock is designed to run perfectly reliably backwards during that hour. At 4 pm the hands of the clock jumped to 5 pm, and it has been running reliably backwards since then. This clock is analogue so its hands sweep its face continuously, but it has no second hand so you cannot tell that it is running backwards from a quick glance. Awaking, you look at the clock at exactly 4:30 pm and observe that its hands point to 4:30 pm. Accordingly, you form the belief that it is 4:30 pm. (Williams and Sinhababu, 2015, p. 48)

We'll refer to the hapless napper as 'John'. We think that the vaguely described situation involving John can be interpreted in at least two ways. On the first interpretation, which we'll call the *Untrustworthy Backward-Clock* case, the cult's dishonest clockmakers have designed their special clocks to deceive viewers during the cursed hour. At 4:30 pm, John looks at his clock and

---

[20] See Adams and Clarke 2016 for additional discussion of Haze's cases.

believes that the time is 4:30 pm on the basis of the clock's saying that it's 4:30 pm. Since the clock was designed to fool viewers during the cursed hour, the clock isn't trustworthy—if hadn't been the case the time was 4:30 pm, it might nonetheless have been the case that the clock said it was 4:30 pm anyway. But why does John's clock say it's 4:30 pm at 4:30 pm? Well, deceivers need not always say what isn't so—they can be misinformed, clever enough to sometimes tell the truth, etc. And instruments they devise to dupe their victims need not always say what isn't so—they can be poorly constructed, cleverly designed to sometimes tell the truth, etc. Perhaps John's clock was the work of a neophyte who failed to replicate the design of the cult's other clocks, which run backwards at a slow speed to display an incorrect time all during the cursed hour. Whatever its design, John's clock is an imperfect deceiver, for even though it would fool viewers at all other times between 4:00 pm and 5:00 pm, it happens to induce a true belief at 4:30 pm. Nevertheless, his untrustworthy clock doesn't enable him to *learn* that it's 4:30. To paraphrase Dretske, you can learn what the time is from clocks, yes, but only from clocks that wouldn't say what the time is unless it were true.

On the second interpretation, which we'll call the *Proprietary Backward-Clock* case, the cult's clocks are meant to be viewed only by cult members, who know that the clocks run backwards during the cursed hour and can readily use them to ascertain the correct time even during these periods. The clockmakers, who don't desire to deceive anyone, teach fellow cult members how to determine what the special clocks say during the cursed hour. For instance, members "in the know" understand that when the hands of a special clock are in what would normally be the 4:35 pm position, the clock is really saying that the time is 4:25 pm, and when the hands are in what would normally be the 4:25 pm position, the clock is really saying that the time is 4:35 pm. Since John doesn't know what his clock is saying during the cursed hour, his true belief that the time is 4:30 pm fails to qualify as knowledge. He believes that it's 4:30 pm because he believes that the clock says so, but he is in no position to tell that this is what the clock says. Even though the clock says that it is 4:30 pm, and even though a cult member would know this, he doesn't. Unlike cult members, he doesn't know how to "tell time" during the cursed hour by means of this clock. John's belief isn't based on the clock's saying that it's 4:30 pm, and therefore isn't based on a sensitive reason.

The Proprietary Backward-Clock case can be instructively compared to the following one, which we'll call the *Bolivian Backward-Clock* case. While

vacationing in Bolivia, Gertrude suffers a stroke, loses consciousness, and is taken to a hospital in La Paz. At 11:55 am, she briefly regains consciousness and, upon looking out the window, sees the Bolivian Congress Building, which features a clock that (unbeknownst to her) runs backwards.[21] Although she cannot read the numerals on the clock's face, at 11:55 am the clock's hands appear to her to be in the 12:05 pm position, and she believes that the time is 12:05 pm. At 12:05 pm the next day, she briefly regains consciousness, sees the clock's hands in what appears to her to be the 11:55 am position, and believes that the time is 11:55 am.   At noon the next day, she briefly regains consciousness, sees the clock's hands in what appears to her to be the 12:00 pm position, and believes that the time is 12:00 pm. Unlike the residents of La Paz, she doesn't know how to "tell time" by means of the Bolivian Congress Building's clock. Consequently, even though her belief that the time is 12:00 pm is true, it fails to qualify as knowledge because she doesn't understand what the clock says about the time. Gertrude's belief isn't based on the clock's saying that it's 12:00 pm, and therefore isn't based on a sensitive reason.[22]

Methods of Arriving at Belief

According to Nozick's initial sensitivity analysis of knowing, S knows that p iff:

(1) p is true.

(2) S believes that p.

(3) If p weren't true, S wouldn't believe that p.

(4) If p were true, S would believe that p.[23]

His later sensitivity analysis incorporated "methods" of arriving at beliefs, which he argued were needed to handle cases like the following one:

---

[21] See, e.g., https://www.theguardian.com/world/2014/jun/25/bolivia-clocks-time-run-different
[22] See Adams, Barker and Clarke 2016 for additional discussion of the backward clock cases.
[23] Cf. Nozick 1981, p. 172 ff.

A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly we must restate our conditions to take explicit account of the ways and methods of arriving at belief. Let us define a technical locution, S knows, via method (or way of believing) M, that p:

(1) p is true.

(2) S believes, via method or way of coming to believe M, that p.

(3) If p weren't true and S were to use M to arrive at a belief whether (or not) p, then S wouldn't believe, via M, that p.

(4) If p were true and S were to use M to arrive at a belief whether (or not) p, then S would believe, via M, that p.[24]

For present purposes, methods of arriving at belief can be construed as procedures that yield reason-based beliefs. To illustrate, we'll use the case involving the grandmother, whom we'll call Granny, and her grandson, whom we'll call Bobby. It's plausible that Granny knows Bobby is well via a method consisting of observing him, which yields the belief that he is well based on R_1, viz. his appearing to her to be well. If he weren't well and she were to observe him, R_1 wouldn't obtain—he wouldn't appear to her to be well and she wouldn't believe he is well on the basis of R_1. Nevertheless, if he weren't well, a different method, one involving hearing others report on his health, would come into play, and she would believe he is well on the basis of R_2, viz. hearing them say he is well. SensitiveReasonsAK has the plausible consequence that Granny knows that Bobby is well on the basis of sensitive reasons, R_1, even though, were he not well, she would nonetheless believe he is well on the basis of non-sensitive reasons, R_2. Thus, SensitiveReasonsAK resembles Nozick's sensitivity analysis in significant ways.

---

[24] Nozick 1981, p. 179. In our opinion, (4) constitutes a valuable feature of certain types of knowledge but doesn't constitute a requirement for knowledge.

Paradox-Free Conditionals and Necessary Truths

   A major difference between SensitiveReasonsAK and Nozick's analysis involves the subjunctive conditionals that each employs. We follow Barker's lead in employing *relevant* conditionals, which are demonstrably free from the infamous "paradoxes of implication." Relevant conditionals, which have been formally codified in System R of Relevant Implication, have been discussed at length in numerous publications.[25] Such implausible theses as the following one are invalid for relevant conditionals:

Thesis A. Given that it is the case that p and the case that q, it follows that if it were to be the case that p, *then* it would be the case that q.

   For present purposes, we'll take it for granted that subjunctive conditionals of the form 'If it were to be the case that p, *then* it would be the case that q', which we'll call *implicative* subjunctive conditionals, are true only if the antecedent, together with existing conditions and laws of logic, laws of nature, etc., implies the consequent.[26] Suppose that Bobby prays to the sun god, Granny doesn't, and the sun rises. Thesis A has the implausible consequences that if it were to be the case that Bobby prays to the sun god, *then* it would be the case that the sun rises, and if it were to be the case that Granny doesn't pray to the sun god, *then* it would be the case that the sun rises. With implicative subjunctive conditionals construed as relevant conditionals, paradoxical theses like Thesis A are provably invalid.[27]
   Although Thesis A is implausible, the following thesis is plausible:

Thesis A*. Given that it is the case that p and the case that q, it follows that if it were to be the case that p, *still* it would be the case that q.

   Subjunctive conditionals of the form 'if it were to be the case that p, *still* it would be the case that q', which we'll call *contraimplicative* subjunctive conditionals, are true only if the consequent is true and antecedent wouldn't

---

[25] See, e.g., Belnap 1967, Barker 1969, Bacon 1971, Anderson and Belnap 1975, Anderson, Belnap, and Dunn 1992, and Mares 2004.
[26] Cf., e.g., Ramsey 1931, Goodman 1955, and Barker 1969.
[27] The logical matrices discussed in Barker 1972, p. 322 are useful for this purpose.

imply its falsity.[28] Thesis A* has the plausible consequences that if it were to be the case that Bobby prays to the sun god, *still* it would be the case that the sun rises, and if it were to be the case that Granny doesn't pray to the sun god, *still* it would be the case that the sun rises; in other words, regardless of whether Bobby or Granny does or doesn't pray to the sun god, *still* it would be the case that the sun rises.

With implicative subjunctive conditionals construed as relevant conditionals, the following counterintuitive thesis is demonstrably invalid:

Thesis B. Given that it is necessarily the case that p, it follows that if it weren't the case that p, *then* it would be the case that q.

Given that it is necessarily the case that 7+5=12, Thesis B has the implausible consequences that if it weren't the case that 7+5=12, *then* it would be the case that the sun rises; and if it weren't the case that 7+5=12, *then* it would be the case that the sun doesn't rise. The sum of 7 and 4 isn't relevant to the movements of the sun, and System R of Relevant Implication provides a way of proving that theses like Thesis B are invalid.

We'll call the conditionals that Nozick employed *S&L subjunctive conditionals*, as the theory of such conditionals was pioneered by Robert Stalnaker and David Lewis.[29] Their goal was to devise a "unified" account that in effect represented both implicative and contraimplicative subjunctive conditionals as conditionals of a single type.[30] Their theory was designed to validate theses like the following one:

Thesis A**. Given that it is the case that p and the case that q, it follows that if it were to be the case that p, [*then-or-still*] it would be the case that q.[31]

---

[28] Such conditionals are called 'semifactuals' in Goodman 1955 and 'counterconditionals' in Barker 1969 and 1973. See, e.g., Barker 1969 and 1973 and Bennett 2003 for extensive discussion of conditionals involving 'then' and 'still', and for references to relevant literature.

[29] See, e.g., Nozick 1982, p. 173, Stalnaker 1968, and Lewis 1973.

[30] Their account, however, isn't all-inclusive, as it doesn't represent conditionals of certain types; see Barker 1973 for discussion.

[31] We enclose 'then-or-still' in brackets to indicate that an S&L subjunctive conditional is expressible in the form 'If it were to be the case that p, it would be the case that q', which includes neither 'then' nor 'still'. In everyday discourse, however, a subjunctive

Thesis A** is acceptable, for it is entailed by the acceptable Thesis A*. The fact that Thesis A** follows not only from Thesis A* but also from Thesis A indicates that S&L subjunctive conditionals are much weaker than implicative subjunctive conditionals.[32] Unfortunately, S&L subjunctive conditionals are subject to paradoxical theses like the following one:

Thesis B*. Given that it is necessarily the case that p, it follows that if it weren't the case that p, [*then-or-still*] it would be the case that q.

As we'll now show, the fact that S&L subjunctive conditionals exhibit such counterintuitive properties as those reflected in theses like B*, and the fact that they are much weaker than implicative subjunctive conditionals, had deleterious

---

conditional that lacks 'then' and 'still' is usually interpreted either as an implicative subjunctive conditional or as a contraimplicative subjunctive conditional, depending on the conversational context, the subject matter, the intentions of the speaker, etc. (The default is to interpret such conditionals as implicative ones, for they constitute the majority.) Nevertheless, it may sometimes be useful to interpret such conditionals as S&L subjunctive conditionals. To illustrate, suppose that Robert is working on a draft of a true-false political opinion questionnaire containing the items: "(a) If the Syrian army were on the verge of losing the war, then it would use chemical weapons" and "(b) If the Syrian army were on the verge of losing the war, it would still use chemical weapons." He decides to replace these two items with the single item, "(c) If the Syrian army were on the verge of losing the war, it would use chemical weapons," thinking that respondents who answered *True* would all share the same opinion and those who answered *False* would all share the same opinion (cf. Stalnaker 1968, p. 100). (It seems plausible, however, that (c) wouldn't normally be construed as an S&L subjunctive conditional—(c) would be ambiguous, and the respondents' answers wouldn't accurately reflect their opinions.)

[32] Implicative subjunctive conditionals are useful for explicating many other types of subjunctive conditionals, such as those involving 'unless', 'provided that', and 'on condition that'. In contrast, S&L subjunctive conditionals are so weak that they are virtually useless for this task, and many of the familiar principles that hold for implicative subjunctive conditionals, such as contraposition, don't hold for them. (See, e.g., Barker 1973 for discussion.) It seems likely that if Nozick had emulated Dretske's use of 'unless' instead of opting for S&L subjunctive conditionals, he wouldn't have encountered so many problems, and his critics wouldn't have opted to use contrapositive forms of his S&L subjunctive conditionals to devise safety accounts. (See, e.g., Sosa 1999 for the use of such contrapositives.)

consequences for both of Nozick's analyses. Nozick himself was well aware of these consequences, and he struggled to cope with their damaging impact on his analyses. In fact, he refrained from committing himself to acceptance of such conditionals (p. 174), and stated that "Our purposes require, for the most part, no more than an intuitive understanding of subjunctives". (p. 680)

In virtue of Thesis A**, the fourth requirement of Nozick's initial analysis, viz. "If p were true, S would believe that p," fails to make an independent contribution to the analysis, as it is implied by the first two requirements, "p is true" and "S believes that p." Nozick offered some tentative suggestions for addressing this problem, but acknowledged that they weren't apt to be satisfactory: "If, as is likely, these explanations do not help, please use your own intuitive understanding of the subjunctives 3 and 4". (p. 176) As he recognized, Thesis B* made it especially difficult for his analyses to account for knowledge of necessary truths. Bobby believes that 7+5=12 because Granny says so; but his belief doesn't qualify as knowledge because she's just guessing. Given that 7+5=12 is necessarily the case, however, the third requirement of his revised analysis is automatically satisfied. In virtue of Thesis B*, if 7+5=12 weren't the case, it would follow that if Bobby used the same method to arrive at a belief regarding whether or not 7+5=12, he wouldn't believe via this method that 7+5=12. As a result, the counterintuitive consequence that Bobby's belief constitutes knowledge can be blocked only if the fourth requirement can be shown to be unsatisfied. In virtue of Thesis A**, however, this is a daunting task. Given that the first two requirements are satisfied, the properties of S&L subjunctive conditionals entail that the fourth requirement is satisfied, and, as we have noted earlier, Nozick himself wasn't confident that the suggestions he offered to cope with such a problem were acceptable.

With implicative conditionals construed as relevant conditionals, SensitiveReasonsAK can readily accommodate knowledge of necessary truths. Bobby believes that 7+4=11 on the basis of his teacher's saying so, and she wouldn't say this if it weren't so. Since his belief is based on sensitive reasons, SensitiveReasonsAK has the plausible consequence that he knows that 7+4=11. Knowing this, Bobby adds 1 to both sides of the equation, "perceives" that 7+5=12, and acquires a knowledge-qualifying belief that 7+5=12 based on an intellectual experience he wouldn't have if it weren't the case that 7+5=12.

Conclusion

We have developed and defended a tracking account that construes knowledge as true belief based on reasons that are sensitive to the relevant facts. Our account handles difficulties that have afflicted previous accounts, addresses several recent challenges to tracking theories, and accommodates knowledge of necessary truths.[33]

References

ADAMS, F. & CLARKE, M. "Resurrecting the Tracking Theories", *Australasian Journal of Philosophy,* 83, pp. 207-221, 2005.

_____ "Two Non-Counterexamples to Tracking Theories of Knowledge", *Logos & Episteme,* 7, pp. 67-73, 2016.

ADAMS, F., BARKER, J. A., & CLARKE, M. "Beat the (Backward) Clock", *Logos & Episteme,* 7, pp. 353-361, 2016.

ANDERSON, A. R. & BELNAP, N. D. JR. *Entailment: The Logic of Relevance and Necessity*, vol. I, Princeton: Princeton University Press, 1975.

_____ and DUNN, J. M. *Entailment: The Logic of Relevance and Necessity*, vol. II, Princeton: Princeton University Press, 1992.

BACON, J. "The Subjunctive Conditional as Relevant Implication", *Philosophia,* 1, pp. 61-80, 1971.

BARKER, J. A. *A Formal Analysis of Conditionals*, Carbondale, III: Southern Illinois University Monographs, 1969.

_____ "Knowledge and Causation", *Southern Journal of Philosophy*, 10, pp. 313-324, 1971.

_____ "Hypotheticals: Conditionals and Theticals", *The Philosophical Quarterly*, 23, pp. 335-345, 1973.

_____ "What You Don't Know Won't Hurt You?" *American Philosophical Quarterly*, 13, pp. 303-308, 1976.

BARKER, J. A. & ADAMS, F. "Epistemic Closure and Skepticism", *Logos & Episteme* I, 2, pp. 221-246, 2010.

_____ "Conclusive Reasons, Knowledge, and Action", *Philosophical Issues*, 22, pp. 35-52, 2012.

BELNAP, N. "Intensional Models for First Degree Formulas", *Journal of Symbolic Logic*, 32, 1-22, 1967.

BENNETT, J. *A Philosophical Guide to Conditionals*, Oxford: Oxford University Press, 2003.

CARRIER, L. S. "An Analysis of Empirical Knowledge", *Southern Journal of Philosophy*, 9, pp. 3-11, 1971.

CHISHOLM, R. *Theory of Knowledge*, 1st Edition, Englewood Cliffs, N. J.: Prentice Hall, 1966.

DRETSKE, F. "Conclusive Reasons", *Australasian Journal of Philosophy*, 49, pp. 1–22, 1971.

_____ "Reply to Hawthorne" in Steup, M. and Sosa, E. (eds.), *Contemporary Debates in Epistemology*, Malden, MA: Blackwell, pp. 43–46, 2005.

FITELSON, B. "Strengthening the Case for Knowledge from Falsehood", *Analysis*, 70, pp. 666-669, 2010.

GETTIER, E. "Is Justified True Belief Knowledge?" *Analysis*, 23, pp. 121–123, 1963.

GOLDMAN, A. "What is Justified Belief?" in Sosa, E. and Kim, J. (eds.), *Epistemology: An Anthology*, Blackwell, pp. 340-353, 1979.

GOODMAN, N. *Fact, Fiction and Forecast*, New York: Bobbs-Merrill, 1955.

HAZE, T. "Two New Counterexamples to Truth-Tracking Theories of Knowledge", *Logos & Episteme*, 3, pp. 309-11, 2015.

KLEIN, P. *Certainty: A Refutation of Scepticism*, Minneapolis: University of Minnesota Press, 1981.

LEWIS, D. *Counterfactuals*, Cambridge, MA: Harvard University Press, 1973.

LYCAN, W. "On the Gettier Problem Problem", in Hetherington, S. (ed.) *Epistemology Futures*, Oxford University Press, pp. 148-168, 2006.

MARES, E. *Relevant Logic: A Philosophical Interpretation*, Cambridge: Cambridge University Press, 2004.

NOZICK, R. *Philosophical Explanations*, Cambridge, MA: Harvard University Press, 1981.

RAMSEY, F. "General Propositions and Causality" in Ramsey, F., *Foundations of Mathematics and other Logical Essays,* Oxford: Routledge and Kegan Paul, pp. 237-257, 1931.

SHOPE, R. *The Analysis of Knowing: A Decade of Research*, Princeton: Princeton University Press, 1983.

SOSA, E. "How to Defeat Opposition to Moore", *Philosophical Perspectives*, 13, pp. 141-153, 1999.

STALNAKER, R. "A Theory of Conditionals", *Studies in Logical Theory, American Philosophical Quarterly Monograph*, 2, pp. 98-112, 1968.

WARFIELD, T. "Knowledge from Falsehood", *Philosophical Perspectives*, *19, Epistemology*: pp. 405-416, 2005.

WILLIAMS, J. & SINHABABU, N. "The Backward Clock, Truth-Tracking, and Safety", *Journal of Philosophy*, 112 (1), pp. 46-55, 2015.