TRACKING THEORIES OF KNOWLEDGE

Fred Adams*

SÍNTESE – As teorias epistemológicas do rastreamento sustentam que o conhecimento é uma relação real entre o agente cognitivo e seu ambiente. Os estados cognitivos de um agente epistêmico fazem o rastreamento da verdade das proposições que são objeto de conhecimento ao embasarem a crença em indicadores confiáveis da verdade (evidência, razões, ou métodos de formação de crença). A novidade nessa abordagem é que se dá pouca ênfase no tipo de justificação epistêmica voltada ao fornecimento de procedimentos de decisão doxástica ou regras de responsabilidade epistêmica. Este artigo oferece um pouco da história das teorias de rastreamento e, então, defende-as contra muitas objeções que se pretendem (equivocadamente) refutadoras dessas teorias.
PALAVRAS – CHAVE – Teorias de rastreamento. Nozick. Dretske. Conhecimento.

ABSTRACT – Tracking theories of knowledge maintain that knowledge is a real relation between cognitive agent and environment. Cognitive states of a knower track the truth of known propositions by basing belief on reliable indicators of truth (evidence, reasons, or belief forming methods). The novelty of this approach is that it places little emphasis on epistemic justification of a kind that aims at guiding epistemic agents by giving doxastic decision procedures or rules of epistemic responsibility. This paper gives some of the history of tracking theories, and then defends them against many of the objections most often judged (mistakenly) to refute tracking theories.
KEY WORDS – Tracking theories. Nozick. Dretske. Knowledge.

1 Introduction

Tracking theories of knowledge began with Alvin Goldman's (1967) "A Causal Theory of Knowing." Goldman launched a bold new approach to knowledge, in the wake of Edmund Gettier's (1963) elegant demonstration that, whatever else knowledge is, it isn't merely justified true belief. At least it was not this on the prevailing theories of justified belief *circa* 1963. Goldman agreed and headed off in a new direction, offering no account of justified belief or warranted belief in his 1967 paper.

Goldman's diagnosis of what went wrong in Gettier's examples was that the thing that made one's belief true in the examples was causally disconnected from the thing that made a person hold the belief. One could justifiably and truly believe (on good evidence) that p (= that Jones owned a Ford or Brown was in

* Department of Philosophy – University of Delaware.

Barcelona). Yet one may fail to know that p since one's belief that p could be true by accident because the thing that brought one to believe was not connected to the thing that made one's belief true. I may believe Jones owns a Ford because I rode in his new Ford and saw the bill of sale before the car was repossessed, and so on. Yet, what makes my belief true is the accident (of which I am unaware) that today Brown is indeed in Barcelona. My belief, though true and justified, is true by accident since Jones's Ford was repossessed and I knew nothing of the specific whereabouts of Brown.

Goldman's diagnosis, in the face of Gettier's examples, and many other cases of perception, testimony, and memory where we would agree that we do have knowledge, was that the accident that bars knowledge would be removed if one's nose were rubbed in the truth, as it were. That is, if there were an unbroken causal chain running from the thing that makes one's belief true to one's belief, then one's true belief would be knowledge. No specific role need be played by one's *justification* for the belief. At least, that was the main hope behind Goldman's causal theory.

It later turned out that beliefs causally connected with the truth could still be true by accident, and the new tracking theory approach to knowledge had to be revised further. Still, the bold new approach, once begun by Goldman, moved forward over several iterations. In what follows, I will recount some of that history.

Other epistemologists who helped develop the tracking theories were Armstrong (1983), and Dretske (1970, 1971, 1981), and, of course, the "tracking" name comes from Robert Nozick's (1981) account, which I'll discuss later.

2 Goldman

As I mentioned above, Goldman's "Causal Theory" led the way to the approach that is now known as the "tracking theories." It turned out that there were at least two problems with Goldman's original version of a causal theory of knowledge. Modifications moved Goldman's (1976) account virtually to the current versions of tracking theories. The first problem was that even causal connections (unbroken causal chains) between the facts that make a belief true and the belief itself, do not eliminate accidentally true beliefs (Gettier-type problems). The second problem is that a modification made by Goldman himself within the 1967 paper, caused an internal tension in his view that he may not have fully appre ciated at the time.[1]

Let's take up the second problem first. The original idea behind Goldman's theory was to have just three conditions for knowledge:

1. p is true,
2. S believes that p,

[1] I once asked Goldman about this tension and he replied "you are the first person ever to ask me that question."

3. There is an unbroken causal chain between the fact that makes (1) true and (2). However, in response to a "newspaper reporter" example, Goldman added a fourth condition:

4. S must be able to reconstruct all important links in the causal chain of (3). This is added to avoid a possible accidentally true belief case with an unbroken causal chain involving a newspaper reporter. A newspaper reporter may witness that p is true, write a report that p, the story may contain a typographical error that not-p, be read too quickly by S as stating that p, and cause S to truly believe that p. These accidents cancel one another and S truly believes that p, while *there is an unbroken causal chain* from the fact that p to S's belief that p. Causal chains are blind to epistemic accidents. Goldman added condition (4) to block such cases from being cases of knowledge. Clearly, S does not know that p in this case. It is just too accidental that S's belief is true. Since S is unaware of the canceling of accidents in the causal chain, S would not put in the canceling of the accidents upon "reconstructing" the links of the causal chain. Due to conditions (4), S would not know that p (as should be the case).

Condition (4) presents at least two problems for Goldman. First, Goldman was trying to chart a new course for theories of knowledge—moving away from a pivotal role for the notion of "justified" belief. But when he adds his fourth condition, he requires that in reconstructing the causal chain, the agent's inferences must "be warranted..must genuinely confirm p very highly, whether deductively or inductively." This makes knowledge not only a matter of being properly causally connected to the truth, but also of being justified ("warranted") in one's background beliefs or inferences. Since Goldman offered no account of what "warrants" or justifies a belief (at least not in this 1967 paper), knowledge could not be merely a causal matter, as the theory seemed to set out to show. An unexplained account of justified belief[2] lurks in the background as a necessary condition for knowing (at least in cases involving inference).

A second problem arises toward the end of the paper. Goldman gives a case where he strongly maintains that he knows Lincoln was born in 1809, but no longer remembers where he learned this or how he inferred its truth from relevant evidence. His newly inserted condition (4) blocks his knowing this. For he can no longer satisfy condition (4). He seems not to have noticed this problem. Either he really does know that Lincoln was born in 1809, and then condition (4) is not necessary for knowledge, or condition (4) is necessary and then he does not really know that Lincoln was born in 1809, despite his claim that "many things we know were originally learned in a way that we no longer remember. The range of our knowledge would be drastically reduced if these items were denied the status of knowledge." (83)

---

[2] Of course, Goldman (1979) filled this gap with his externalist reliabilist account of justified belief.

Now quite independent of these problems is that Goldman's causal theory still was subject to problems of accidental causal chains. Even with the addition of his condition (4), his view allows there to be an unbroken causal chain, and for one to be correct in reconstructing the major links in that chain, and still one's belief may be accidentally true. Perhaps the best example that illustrates all of this is one from Dretske (1971). Suppose Harold buys a fair lottery ticket with a chance of 10,000 to 1 of winning (let us say a new Porsche). Rockaford, a rich friend of Harold's decides to buy Harold a new Porsche, if Harold loses the lottery. You know these facts. Then later you see Harold riding in his new Porsche and believe, correctly, that the best explanation of how Harold got the Porsche is from Rockaford. However, you do not know how the lottery came out. Now all of Goldman's conditions are met. Harold got the Porsche from Rockaford, you believe this, the first event caused the second via an unbroken causal chain, and you correctly and with warrant, reconstruct the important links in the chain. Yet, you still don't know how Harold got his Porsche. You have not eliminated the very real possibility that he did not get it from Rockaford. For all you know, Harold might have won the Porsche in the lottery.

Goldman (1976) later "abandon[ed] the requirement that a knower's belief that p be causally connected with the fact, or state of affairs, that p" (p. 120). Instead, he favored a *reliability* requirement[3] where "the causal processes or mechanisms responsible for a belief produce true beliefs, and would produce true beliefs or at least inhibit false beliefs in relevant counterfactual situations" (p. 120). To this he added that the relevant mechanisms enable a person to "discriminate or differentiate incompatible states of affairs" (p. 120). And importantly, a knower can discriminate the truth of p from relevant alternative states (p. 121). This is the paper where Goldman introduced the examples of the *barn façades*, which, when relevant, one must be able to discriminate from genuine barns if one is to know the barn he is looking at is a barn. We can see that if one must discriminate the barns from the local barn façades in order to know there is a barn, one would also have to be able to discriminate Harold's getting his Porsche from Rockaford vs. getting it via the lottery, if one were to know Harold got the Porsche from Rockaford. So Goldman's new account was designed to rule out just such cases of multiple causal chains that proved to be a problem for his earlier causal account of knowing.

Still in his 1976 paper, Goldman was looking to move epistemology away from what he called the "Cartesian" perspective that "combines a theory of knowledge with a theory of justification" (p. 143). He adds: "My theory requires no justification for external-world propositions... It requires only... that beliefs in the external world be suitably caused... If one wishes, one can so employ the term 'justification' that belief causation of *this* kind counts as justification... But this is entirely different from the sort of justification demanded by Cartesianism" (p. 143).

[3] Indeed, the wording in that follows in this quote from Goldman is surprisingly close the view Nozick (1981) would eventually develop.

Goldman maintained that "the traditional justified-true-belief account of knowledge is of no help in explaining our assessment of why one does know there is a barn present, when there are no barn façades nearby, and why one does not know there is a barn, when there are barn façades nearby that one cannot discriminate from barns. For one's justification for believing there is a barn "is the same in both cases" whether there are or are not indiscriminable barn façades nearby (p. 122). Further, Goldman candidly confesses: "My old causal analysis cannot handle this problem either" (p. 122).

On Goldman's (1976) new account, however, "S has perceptual knowledge if and only if not only does his perceptual mechanism produce true belief, but there are no relevant counterfactual situations in which the same belief would be produced via an equivalent percept and in which the belief would be false" (139). So if Henry *cannot* discriminate barns from barn façades (and there are barn façades near by), then Henry doesn't know there is a barn in the field. And if Sam *can* discriminate Judy from her identical twin Trudy, though others cannot, then Sam can know that it is Judy in front of him. For, while Henry is looking at a real barn, he would have the same type of percept if looking at a nearby barn façade. In contrast, Sam would not be caused to have the exact same type of percept he is having from Judy (were he viewing Trudy). So Sam knows it is Judy.

3 Armstrong

Armstrong (1973) offered a "reliable indicator" account of knowledge. When a reliable thermometer indicates that it is 70 degrees Fahrenheit in the room, it is 70 degrees Fahrenheit in the room. Using a thermometer as a model of a gauge or instrument that can pick up information about the environment when working reliably, Armstrong thought non-inferential perceptual knowledge would have the same basic properties. The senses are like gauges or instruments and they feed one's perceptual beliefs. That is, when working reliably one's perceptual beliefs would stand as reliable signs or indicators of one's environment. On this view, one's perceptual judgments, such as "It sounds to me as if there is a noise," "It feels to me as if it is getting hotter," "It looks to me as if there is something red and round" are reliable indications (in the right circumstances) of the presence of a noise, its getting hotter, or the presence of something red and round. That is, one's perceptual belief that p stands as a reliable indicator of p, on this view. Armstrong puts it this way: "My suggestion is that there must be a *law-like connection* between the state of affairs Bap [A's believing that p] and the state of affairs that makes 'p' true such that, given Bap, it must be the case that p" (1973, p. 166).

The departure from Goldman's 1967 is that if there is a *law* that relates the believer's belief to the state of affairs that p, then there is not merely a *causal chain* in place but a counterfactual relationship (a nomic relationship). For instance, Armstrong says that the counterfactuals that connect the knower to his

environment include[4] that "if p had not been the case, then it would not have been the case that A believed that p" (p. 169). So, a knower is not merely causally, but also nomically related to his environment (such that there will be relevant properties of the knower and of the thing known that figure in the counterfactual relationships that support knowledge). Exactly what the *laws* that support knowledge are is something Armstrong leaves to future science. Armstrong says that person A non-inferentially knows that p if, and only if,

(1) p is the case,

(2) There is some specification of A such that, if any person is so specified, then, if they further believe that p, then p is the case[5] (Armstrong, 1973, p.168).

After presenting his theory of non-inferential knowledge, Armstrong labored over three potentially damaging objections: Deutscher's, Murphy's, and Waller's. Deutscher suggested an example of a "veridical hallucinator," a person who no matter what the next stimulus in his environment, would be caused by it to believe that there is a sound. If, as luck has it, the next stimulus is a sound, it appears that conditions (1) and (2) are met, but we do not want to say the hallucinator knows that there is a sound. Had there been a light instead, he would have believed there was a sound. Armstrong puts the problem this example poses to his account in this way: "suppose...we go on specifying...A's properties in indefinitely greater and greater detail. Will we not inevitably arrive at a complex property H that A has, such that, given a believer who had just that complex property *and* believed that there was a sound in his environment, then there *must* (nomically must) have been a sound there?...It seems that some restricting condition must be placed upon H in order to rule out such a counter-case" (p. 172). The restriction he selects is that "H must be such that the situation has some real probability or at least possibility of being repeated" (p. 173). Armstrong was never fully happy with this restriction, saying that it was "not very precise" and admitted of "ambiguity," but it was all he gave to block Deutscher's example.[6]

Murphy's objection was along these lines. Suppose Gary is an incorrigible hypochondriac who always believes he is about to die. Suppose further that he develops an incurable fatal illness. Put all of this into H (the specification of A in Armstrong's condition (2)). Now when Gary believes that he is about to die, it is so. Yet

---

[4] Although, surprisingly, this condition does not make it into Armstrong's two formally stated conditions for perceptual knowledge below. This type of condition does make into Nozick's (1981) account, and because of this, Nozick's account easily deals with difficult putative counterexamples that Armstrong labors over in his book.

[5] There is a more formal statement of these conditions as well, but the formality does not really change anything. It merely makes clear that there is some set of properties H that must be in place for the knower to be nomically related to the thing known. In the formulation above, that is covered by the "There is some specification of A."

[6] Neither Dretske's nor Nozick's accounts would have any problem dismissing the three examples over which Armstrong struggles. It is a bit surprising, since Armstrong knew of Dretske's account and the counterfactual he mentions (above in the text) would stop the Deutscher example in its tracks, had Armstrong employed it.

Gary does not know that he is about to die (as he has not been informed of his fatal illness). Armstrong points out that it does not matter in Gary's case whether he believes he is about to die or not, condition (2) will be true. So Armstrong has to block this case by saying that properties of one's belief must be relevant to satisfying condition (2) in cases of genuine knowledge. In the example, so much is packed into H that it does not really matter what Gary believes, condition (2) will still be true.

Waller's objection was similar. Suppose there is a brain chemical T that is necessary for conscious thought. Suppose further that Ken mishears someone say he has "expletive –T" for brains, and comes to believe he has chemical T in his brain (which is true). Now when Ken believes that he has chemical T in his brain, of course it will be true. But Ken does not know this. Armstrong claims that Ken's belief is not a "reliable sign" of the presence of the brain chemical T, so he does not know. Armstrong requires yet another "block" for Waller's example.

So instead of modifying his account to take care of these three examples, Armstrong puts in place three restrictions to block them:

(a) H must be such that there is a real possibility of the situation covered by the law-like connection recurring.
(b) The properties mentioned are nomically relevant to the law-like connection.
(c) The relation of the [belief] to [the belief's truth maker] is that of *completely reliable sign to thing signified*[7] (p. 182).

While it is beyond question that Armstrong's account moved the tracking theories forward and influenced Dretske (1981) and Nozick (1981), it is also clear that his account is not satisfactory. His restrictions to block counter-examples do not seem well-motivated and appear to be a patchwork rather than genuine insight into the nature of knowledge. However, his recognition that there is a genuine nomic relationship between a knower and the known lies at the very core of the tracking approach to knowledge. I am confident that Nozick's tracking account was formulated in full cognizance of the struggles Armstrong made in laboring over these three types of example. Nozick's account handles them with ease.[8]

4 Dretske

Dretske's (1971) *conclusive reasons* approach was similar to Goldman's (1976) and Armstrong's (1973) accounts. It requires a nomological connection between one's reason for believing and the truth of one's belief. So it is stronger than a causal connection. Yet it differed from Armstrong's account in an important respect. Armstrong wanted the reliable indicator relationship that a knower instantiates to relate the knower's belief with the truth of the belief. Dretske wanted the basis for the

[7] In this case, it seems to me that Ken's belief is a reliable sign, despite Armstrong's restriction (c). When Ken thinks chemical T is present, it is (and is nomically guaranteed to be there, to boot).

[8] I will not demonstrate that Nozick's account has no problem with these, but leave it to the reader to easily verify for one's self.

belief (reason state R) to serve this role. This makes it possible to have knowledge on Dretske's theory, but not on Armstrong's in special cases.[9]

For Dretske, a reason state R can be any sort of empirical basis for belief (most often a perception of some physical event or condition). His conditions for knowledge are that S knows that p when:

1) p is true,
2) S believes that p,
3) S's belief is based upon a conclusive reason R.

A conclusive reason R is defined as one in which, relative to one's circumstances C, it is "not possible that R & -p."[10] The relevant sense of "possibility" is *circumstantial*. It is physical possibility relative to one's circumstances. Taking Goldman's example of the twins Judy and Trudy, it is not impossible that someone (other than Judy) be in front of you and look just like Judy. Trudy would look just like Judy to you (since you cannot tell the twins apart). However, if the twins were separated at birth and one (Trudy) were in Irkutsk, and would never leave, and you are here, and would never go there, then it is not possible that you are having an experience as of Judy and Judy not be before you—relative to your aforementioned circumstances C.[11]

So if you do know what Judy looks like and there are no Judy-look-a-likes in your local environment, then you do know it is Judy in front of you. It is true. You believe it. Your reason state R for believing Judy is present is your perceptual experience as of Judy. And it would not be possible (in C) for someone to look just like Judy to you and not be Judy. So it is not possible for R and not-p. Hence, you know it is Judy (even though Trudy exists and would pose a threat to your knowledge, if she were to become a world traveler).

What makes R "conclusive" is nothing about subjective certainty or confidence in belief. It is the objective fact that, in C and given R, it is not possible that not-p. That nomic guarantee of the truth of one's belief is the nature of the conclusiveness in Dretske's account.

And, as we can see in the twin example, conclusiveness does not rule out all physical possibility of "R and not-p." Such a combination is still physically possible in other circumstances, and, when it becomes actual, it can rob one's knowledge. However, when it is physically blocked by the physical factors of one's local environment, it can secure knowledge, just as easily.

Of course, neither Dretske nor any of the tracking theorists require that one knows one has knowledge (or has a conclusive reason or knows one's circumstances C). No such requirement is placed upon knowing.[12]

---

[9] This would be the case in Kripke's famous "red barn" example, which we will consider later when looking at Nozick's account. Here I will just note that Dretske's theory allows knowledge of the barn which is red, but Armstrong's would not because one's barn beliefs are not reliable signs of barns.

[10] Alternatively, Dretske uses the wording "R would not be the case unless p." This too is relativized to circumstances C.

[11] Let's assume no one else is a dead ringer for Judy in your surroundings, either.

[12] DeRose (1995) makes much of this in his objections to tracking theories, as we shall see.

5 Nozick

Nozick's (1981) tracking account of knowledge is very similar to the forgoing accounts.[13] Similar to Dretske, he departs from Armstrong in requiring that the "tracking relations" go through a belief-forming method M,[14] rather than just correlate one's belief with the truth. As we shall see, this becomes important when considering objections by Kripke and others, objections which lose track of the method M used to form beliefs. Amplifying on Dretske's account, Nozick uses the "possible worlds" analysis of counterfactuals to analyze the conditional sentences in his account. As we shall see, this too saves Nozick from a purported counterexample by Martin (1975).

Hence, Nozick requires of a knower that:

1) p is true,
2) S believes that p,
3) If p were not true and S used belief-forming method M, then S would not believe p,
4) If p were true and S used belief forming method M, then S would believe that p.

He calls conditions (3) and (4) "tracking conditions." The name has stuck to his and the other accounts of knowledge above. In part, he motivates the tracking conditions by showing how they handle Goldman's barn façade examples and Harman's deceased dictator example. The barn façade example is handled by condition (3) because, if we allow the casually inspected look of a barn to be method M, it is not true that if there were not a barn, one would not believe there were. A barn façade would look like a barn from a distance and bring one to believe it was a barn. So even though the other conditions may be true, one would not know there were a barn in a barn façade-infested environment.

In the Harman example, the dictator dies. The newspaper prints the story, and fearing a coup, the next day prints a retraction. You read the first paper, but not the second, and then go to sleep. While you sleep the retraction is printed and delivered. Do you lose knowledge that the dictator dies? Harman says yes. Nozick agrees. Condition (4) explains this loss of knowledge. While the other conditions are satisfied, it is not true of you that, if the dictator were dead you would believe it. He is dead and you do believe it, but it is not true of you that you would believe it, for you are not unlike your neighbors who gave up believing when they read the retraction (in the paper on your front doorstep). You would give up believing the dictator is dead too, if you were to read the paper.[15]

---

[13] Indeed, to his credit, in footnote #53 of the chapter on knowledge, Nozick (1981) acknowledges just how close is his account to Dretske's, and acknowledges his debt to Goldman.

[14] See Nozick (1981) for finer details about methods than I'm going into here.

[15] Nozick's conditions easily handle the examples Armstrong struggled over, as well.

That concludes my brief review of the tracking theories of knowledge. In what follows, I will defend tracking theories against some recent examples and objections that are believed to show that the tracking theories fail. I cannot address all such objections here, but I can discuss those that are most prominent in the literature and in conversation at several recent conferences on epistemology.[16]

6 The Martin Example

Ray Martin's (1975 and 1983) offered the same counterexample against the accounts of both Dretske (1971) and Nozick (1981). Martin's purported counterexample went as follows. Zack goes to the racetrack and places bets on Gumshoe to win in the first race and Tagalong to win in the second. Zack is absent for both races and does not hear the outcomes prior to going to the ticket window. He presents his ticket stub and receives a payoff. He presumes it is for Gumshoe winning in the first race. He is correct, but has heard nothing about which horse (Gumshoe or Tagalong) won. Martin claims that the event of Zack's receiving the payoff should constitute a reason R (in Dretske's broad sense of 'reason'). He further claims that Gumshoe's winning in the first race constitutes proposition p. Then, Martin claims R should be "conclusive for p" (on Dretske's definition) because Tagalong's losing in the second race made it impossible that Zack receive a payoff unless Gumshoe won. So Martin claims that if Zack believes that p on the basis of R, Zack satisfies Dretske's conditions for knowing. However, it seems clear that Zack does not know that Gumshoe won, since his evidence is not sufficient to rule out Tagalong's winning. Thus, according to Martin, Dretske's conditions for knowledge are satisfied, but Zack does not have knowledge.

Martin believed that the same example worked against both Dretske and Nozick. First, I'll explain why it does not work against Nozick and then I'll explain why it does not work against Dretske. Martin maintained that the method M of belief formation Zack used to believe Gumshoe won was the following. Zack shields himself from any information about the results of either of the two races. And then after the conclusion of the second race Zack presents his ticket to an automated cashier. Martin claims that using M Zack tracks the truth, but does not know that Gumshoe won. It is true that Zack does not know, but it is false that method M tracks the truth, on Nozick's account of "tracking."

Nozick was very clear that methods only track the truth, on his account, if they do so not only in the actual world but also in close possible worlds. It is true that in the actual world, given that Tagalong lost, the automated teller would not have paid off unless Gumshoe won. It is not true in close possible worlds. For in some close possible world Gumshoe loses, Tagalong wins and the automated teller pays off. Martin stipulates that Zack's method is still to believe that Gum-

[16] I gave some of these replies at the Rutgers Epistemology Conference as an invited discussant in April of 2001, and in (Adams, 2004). Also see (Adams & Clarke, 2005) for a more elaborate version of some of these objections and replies.

shoe wins, no matter what. His method is somewhat irrational. Nonetheless, it is false that Zack's method M "tracks the truth" in Nozick's sense of tracking. Even if Gumshoe were to lose and Zack were to receive a pay out in a close possible world, Zack would still believe Gumshoe won (violating Nozick's third condition that if p were not true and one used method M then one would not believe that p). Thus, Martin's is not a counterexample to Nozick.

Martin's example also fails against Dretske.[17] Dretske specifies that conclusive reasons hold relative to a specific set of circumstances C. So R (a payout) is a conclusive reason for p (that Gumshoe won) when, given C, it is not possible to have R and not-p. In the Martin example, Zack buys one ticket prior to the race giving him Gumshoe in the first and Tagalong in the second. The circumstances C become fixed at the time of purchase. So a payout by a teller (R, with no more information) says only that either Gumshoe won or Tagalong won. It turns out that Martin actually *changes the circumstances from C to C' after the first race*. Then Martin claims, given that Tagalong lost (in C'), a payout (R) says that Gumshoe won (that p) because given C' it is not possible to have R and not p (not possible to receive payment unless Gumshoe won). But R's status as a conclusive reason is tied to circumstances C (not C') because Zack purchased the ticket in C (not C'), and Zack has acquired no new information about the outcome of the first race. R is not a conclusive reason relative to C, and that is why Zack does not know that Gumshoe won on the basis of R. The most R says, relative to C, is that either Gumshoe won or Tagalong won (but not which).

Now Martin may claim that what counts as the relevant "circumstances" is fast and loose on Dretske's theory. But that is not so. Consider the twin case above. Whether Trudy is a world traveler or not is something that must go into Tom's circumstances, when considering whether Tom knows it is Judy in front of him. If Trudy is "stuck" in Irkutsk, so to speak, then Tom has no trouble knowing it is Judy he is seeing. If, however, Trudy has become a world traveler, then she might be anywhere, even in Tom's vicinity. Then it truly is possible that, given someone looking like Judy in front of Tom, it might be Trudy (even though it actually is Judy), and Tom will not know it is Judy. These considerations (and factors) about Trudy and her travels are objective and relevant. And when Zack buys his ticket (and which races have been run or not run) is also an objective matter of his circumstances. There is nothing fast and loose about one's circumstances or the objective factors that determine them.

True, there is some vagueness over how close the relevant alternatives (Trudy) can be before Tom loses knowledge that Judy is in front of him. This is a general difficulty about metrics for nearness of possible worlds, when testing counterfactuals. But it is not unprincipled or unruly whether the circumstances have changed. Similarly, it is not unprincipled or unruly that the assessment of

[17] Yet, Dretske was so moved by Martin's example that he abandoned his "conclusive reasons" view and took up his new "information theoretic" account of knowledge. I talk about this new view below.

whether R is a conclusive reason in Martin's example is fixed by circumstances C (or C'). It is clear from the example itself, that the relevant circumstances are C (not C') because Zack purchases the ticket *prior to the races*. The whole pretext for Zack to believe anything at all about Gumshoe at the payout is that he purchased a ticket *on both horses* in context C. Hence the payout R is indicating something about the *races for which Zack holds a ticket* (not just about the outcome of the second race). Thus, the matter of whether payout R is a conclusive reason is to be assessed relative to the ticket purchase and *to context C* (not C').

7 The Pappas & Swain Examples

Dretske maintains that S has a conclusive reason R for p when, in circumstances C, it is not possible to have R and not-p. In a roughly equivalent formulation, Dretske (1971) formulates a conclusive reason this way: "In circumstances C, R would not be the case unless p were the case." Pappas & Swain correctly point out that "like most subjunctives, this one is vague on the question of what precisely we are allowed to imagine in considering whether R would be the case if p were not" (61). So they offer two examples with different things held fixed in the two examples to fit two different interpretations of the conclusive reason subjunctive. In either case, they maintain, S can know that p with less than a conclusive reason.

In their first example, we are to suppose that George is having visual experiences as of a cup on a table. George has all the relevant concepts to know cups and tables. He is paying attention, the lighting is good and his visual system is functioning properly. In such a case, George knows that p (there is a cup on the table) on the basis of R (a visual experience as of a cup on the table and George's background knowledge of cups, tables, and so on). Now imagine one significant change. Let it be physically possible (in C) that instead of there being a cup on the table, there is an indistinguishable hologram of a cup projected to the spot on the table where the cup would have been. In such a case, the subjunctive "George would not be having the visual experiences as of a cup being on the table in these circumstances unless there were a cup on the table" is false. Pappas & Swain conclude "We thus have a case in which S does know that p (there is a cup on the table) on the basis of R (his visual experiences plus background knowledge) in the relevant circumstances, but the corresponding conditional [R would not be the case unless p_{FA}] is false" (64).

Now it is true that the conditional – R would not be the case unless p – is false. But it is also true, because of this, that George could not possibly know that p. In this case, as in the case with the twin's Judy and Trudy, we must ask, where is the hologram machine? If it is anywhere near George, anywhere near the places he frequents, then, given his visual experience as of a cup on the table, it is pure accident that there is actually a cup on the table (rather than an indistinguishable hologram). George could no more know that there is a cup on the table when, in his circumstances, there might be a hologram there instead, than Tom could know it is Judy before him, when it might be Trudy (our world traveler) instead. Hence, this is not a case in which George knows he is seeing a cup, even if he is seeing a cup.

Of course, were the hologram machine in a remote location (stuck in Irkutsk, along with Trudy, say), then George would be able to know there is a cup on the table here in Ohio. However, then the conclusive reason conditional – R would not be the case unless p (relative to C) – would be *true*, not false.

In their second example, Pappas & Swain remove all doubts about the potential ambiguity of circumstances C. They place a back-up generator *in the basement* of Marshall's friend Louise (not off in Irkutsk). In the example, the generator is turned off (though it works, and it is physically possible for it to be turned on). Marshall works for the local power company. He has excellent knowledge of the generators, the power plant, the circuits of the city, and so on. He sees that the generators have just been checked over and are in excellent working order. In the evening he strolls down through streets of the town to the house of his friend Louise. He sees the street lights on along the way. When he gets to the house of Louise, he sees that her lights are on as well. Marshall comes to believe that p (the downtown generators are powering the lights in the house of Louise). He is correct.

Marshall does not know about the back-up generator in Louise's basement that comes on smoothly and without a hitch (let's say with no dimming of lights), when the power goes off to the house. Marshall enters the house and believes that p on the basis of R (his seeing the lights along the way, his seeing the lights at Louise's house, and his background knowledge from working at the power company). Pappas & Swain insist that "surely we do not want to say that the fact that his friend [Louise] has a generator in his basement *prevents* S [Marshall] from having knowledge that the company's generators are causing the lights to be on" (66).

On the contrary, that *just is* what I want to say. If Marshall is inside Louise's house, and is not looking out on the street lights across the whole city, seeing only the lights on in the house will not tell Marshall that the downtown generators are supplying power to the house of Louise. Why not? Because the back-up generator in Louise's basement may well be on (maybe Joe is giving it its annual test to make sure it still works—overriding the power supplied from the power company at that very moment). Squirrels perish each year in transformers and from chewing through power lines (to their demise). Each year huge limbs fall across power lines breaking the connection lines to houses. Accidents occur frequently knocking out power. There is no reason to think it is different near the house of Louise. The lights behind her house, not visible from her front room, may be out and the power to the houses across the street may be on a different circuit. There is no end to the things that have non-zero probabilities that may be happening to cause the basement generator to kick in. If the basement generator kicks in smoothly and silently (from the room where Marshall and Louise are located), then it is indeed false that R would not be the case (Marshall would not see the lights on in the house of Louise and the other street lights viewable from inside the house) unless p (the downtown generators were supplying power to the house of Louise). But precisely because that subjunctive conditional is false, Marshall is prevented from knowing that p, by R alone. Therefore, this example too is not a counterexample to the conclusive reasons approach (nor to tracking theories generally).

8 The Kripke Examples

Famously, in an unpublished lecture given at the American Philosophical Association in the early 1980s,[18] Saul Kripke presented a host of purported counterexamples to Nozick's tracking theory of knowledge. There were two examples, more than any others, that have "stuck" and can be found sprinkled throughout the epistemological literature purported as showing definitively that Nozick's account does not work. I will explain why those examples do not show anything of the kind.[19]

The first is the "red barn" example. Peg is looking at a red barn. As it turns out, she is in an environment where red barns cannot be faked though barns of other colors can be faked. Peg has perfect eyesight and is observing the barn under standard lighted viewing conditions. According to Nozick's conditions (truth, belief, and the two tracking conditions: if p were not true, S would not believe that p, and if p were true, S would believe that p), Peg knows there is a red barn, but Kripke claims that it is a consequence of Nozick's theory that Peg does not know there is a barn. Peg does not know there is a barn because she fails' Nozick's tracking condition that if there were not a barn, Peg would not believe there were. She fails this condition because she would believe of a white fake barn, for instance, that it is a barn.

This only appears to be a case where one knows that there is a red barn but not that there is a barn if one loses sight of the method of knowing, M. Nozick is very clear that his tracking conditions must be relativized to methods. For example, otherwise only God could satisfy the condition that if p were true, one would believe it. When we specify the method of belief formation that yields Peg's knowledge that there is a red barn, we see that the same method yields her knowledge that there is a barn. Here's how.

Let's look at a related example, that, as you will see, is like Goldman's (1976) dachshund-wolf example. Suppose Pam never confuses Volkswagen Beetles with Saabs. However, for some strange reason Pam thinks Saabs are the bigger luxury version of the Volkswagen and that Beetles are the less expensive economy version. So she does think that some non-Volkswagens (namely, Saabs) are Volkswagens. But when Pam sees a Beetle under ideal viewing conditions she knows that it is a Volkswagen Beetle. Let's say that her method of knowing is to look at the shape and size of the Volkswagen Beetle. The visual experience of the shape and size of the Volkswagen Beetle reliably indicates both that the car is a Beetle and that it is a Volkswagen. So Pam not only tracks that the car is a Beetle but

[18] A session that was so crowded that Fred Dretske and I were unable to get into the room, by the way. We had to wait for the manuscript to be circulated unofficially, and for word of the examples to spread to know what Kripke said.

[19] I first responded to these two examples in a commentary on a paper by David Cole "Counterfactual Analyses of Knowledge and Relevant Alternatives" at the Central APA in Chicago, IL April 1989 (I later sent a copy of that commentary to Nozick). For replies to more of Kripke's examples, see (Adams & Clarke, 2005).

also that it is a Volkswagen. Not only does she know that it is a Beetle, but she also knows that it is a Volkswagen. Pam satisfies Nozick's tracking condition that if it were not a Beetle and she used the VW-Beetle method, she would not believe that it was a Beetle. It is also true that if it were not a Volkswagen and she used the VW-Beetle method, she would not believe it was a Volkswagen. So contrary to Kripke, this is a case where *Pam knows both* that it is a Beetle and that it is a Volkswagen because she is using the same reliable method to detect both. Of course this only works in environments where VW-Beetles are not faked and where they are visually distinguishable from Saabs.

Returning to the red barn example, one can see that it has the same structure. Peg is using the red barn look to detect the information both that there is something red and a barn. Since red barns cannot be faked (as per Kripke's stipulation), the look of a red barn reliably indicates both a barn and one that is red. Since Peg is using the reddish barnish look to form her belief that there is a red barn, she satisfies Nozick's tracking condition that it if there were not a red barn, she would not believe there were. And if there were not a red barn present and Peg were to believe there were a red barn, she would not be using the same red barn look method. This means that when Peg believes there is a red barn, she believes there is a barn in part employing the red barn method. This method insures that Peg knows of this structure that it is both red and a barn. Hence, Peg indeed does know of the red barn that it is a barn, contrary to Kripke's claim. Essentially, Kripke's example only appears to be problematic if one violates Nozick's strictures on method.

It should be noted there still may be cases where one's knowledge that something is a red barn would come apart from one's knowledge that something is a barn. This may happen, were one to employ different methods than Peg's. For instance, if Peg sees the red barn in good lighting in the morning and believes it is red barn, she should know. But if Peg becomes disoriented and sees the red barn in complete darkness in the evening she may believe, but not know, that it is the same barn, and therefore, may not know of the same barn that it is a barn (because she can no longer see its color).

Someone may object that this is slicing methods too thinly or without principle merely to rescue Nozick. Maybe there will be some sympathy for the notion of a VW-Beetle method of belief formation, but some may balk at the notion of a red-barn method of belief formation. However, consider that even the Stickleback fish has to solve the red dot detection method of rivals during mating season. If fish can solve this epistemic detection problem, yielding knowledge of their rivals, then surely we can solve it and there are real live models of this abstract method of knowing

The second of Kripke's examples that I'll discuss here is the modified deceased dictator example. I gave the Harman example above when explaining the motivation for Nozick's condition (4). Kripke's modification of this original example by Harman is to add the proposition (q) that "I have read an uncontradicted report of the dictator's death." Kripke maintains that on Nozick's conditions (truth, belief,

and the two tracking conditions), you know p (the dictator is dead) & q (you have not read a contradicted report of his death), but not p (the simple proposition that he is dead). As we saw above, Nozick maintains that you don't know p alone because you fail to satisfy condition (4).

Supposedly, you know p&q because your belief that p&q does satisfy the condition that if p&q, you would believe p&q. This is true because for you to believe that you had read an uncontradicted report of the dictator's death, you would have to have read the first newspaper but not the second contradicting newspaper (keeping q true). And, of course, you satisfy condition (3) because if not p&q, you would not believe anything about the dictator's death to start. When p&q is false either there is nothing in the paper about the dictator because he's not dead, or he is dead but you have read a contradicted report and drop your belief, or you simply would not know what to believe. While a bit complicated, you would satisfy all four of Nozick's conditions for knowing p&q, according to Kripke. Naturally, it would be bad news to be able to know p&q but not p by the same method (reading the newspaper).

In reply, I maintain that you know neither p nor p&q, on Nozick's own conditions. To see why, I suggest that you will either be savvy and not believe something merely because it appears in the state-controlled newspaper or gullible, and will believe everything that appears there. So suppose you are savvy. Then you do not know p&q because your method of distrust causes you to fail the fourth condition. It is not automatic that, for this conjunctive truth, you would believe it. A savvy you rejects many things that you read in the state-run newspaper. So, suppose that you are gullible. Then you believe everything that appears in the state-run newspaper. In that case you fail condition three, viz. that if p&q were not true, you would not believe p&q. Gullibly, you believe falsehoods that are printed in the newspaper. In that case, you cannot (without independent confirmation) learn anything from a state-run newspaper that is highly unreliable. So, on Nozick's own conditions, the dictator example is not a case where you know p&q, but not know p (as intended by Kripke).

9 DeRose's Abominable Conjunction

Many people follow DeRose (1995) in abandoning the tracking theories because both Dretske (1970) and Nozick (1981) reject closure. Both maintain that it is possible to know that p and know that p entails q and yet fail to know that q. This leads to what DeRose calls the "abominable conjunction." You might know that you are reading this paper right now and know that if you are reading this paper right now then you are not a brain in a vat, but not know that you are not a brain in a vat. DeRose and many others reject the theories on the grounds of their "intuitively bizarre results." They embrace the conjunction that it is possible that you know you are reading this paper right now but you don't know you are not a brain in a vat.

I do not here want to discuss the merits of accepting or rejecting the closure principle. It would take us too far afield and there are excellent ways to access the debate (DeRose & Warfield, 1999). Instead, let me point out that at the same time that philosophers reject tracking theories on the grounds of the abominable conjunction, they acknowledge that a dog or infant might be subject to the same abominable conjunction. An infant may know its mother is holding it, but not know it is not a brain in a vat. My dog Raven seems to know when I'm home, but I doubt seriously whether she knows she is not a brain in a vat. If knowledge of any kind can have this result, then it is not sufficient to reject the tracking theories because the results don't match one's favored intuition. Sometimes intuitions have to go (absolute space and time, Euclid's parallel postulate, flat earth, etc.). Furthermore, some abominable conjunctions turn out to be true – conjunctions such as Moore's Paradox: "It is raining but I don't believe it."

True, examples of the infant and my dog Raven are not examples of failures of closure, for they are not cases of known implication. Nonetheless, they are examples of abominable conjunctions. Perhaps the conjunctions seem *even more abominable* when one adds in failure of closure. Dretske (1970) maintains that one can know that there is a zebra in the cage at the zoo and know that, if it is a zebra, then it is not a painted mule, while failing to know that it is not a painted mule. And, of course, Nozick (1981) maintains that one can know that he is reading this paper now and know that if he is, he is not a brain in a vat, while failing to know that one is not a brain in a vat. Both Dretske and Nozick claim that their theories of knowledge have these consequences and explain why. Basically, the idea is that one tracks the truth about the zebra (or one's reading) but not the truth about the painted mule (or the vat) because to determine whether one tracks the latter one must look in environments of precisely the kind where there are painted mules in zebra cages or environments where people are unwittingly captured and envatted. In such environments, tracking fails, closure fails, and if closure fails, the abominable conjunctions may seem all the more abominable. But that they happen at all, as above, is not a reason by itself to reject tracking theories. Or so I claim. What one needs to show is not that tracking theories are committed to them, but that worse than abominable, they are false. That remains to be shown.[20]

10 Bonjour's Clairvoyants

Bonjour (1980) presented a series of examples and modifications revolving around the example of a clairvoyant who reliably believes the truth. The point of the examples was to challenge externalist accounts of knowledge and justification. Bonjour claimed that knowledge requires internally justified (epistemically

[20] See (Steup & Sosa, 2005) for a recent exchange between Dretske and Hawthorne, where Dretske continues to defend non-closure.

responsible) belief and that in his examples, subjects did not know because they were epistemically irresponsible in their beliefs. Bonjour's primary targets were clearly Armstrong (1973) and Goldman (1976, 1979), but the tradition has taken his attack to extend to tracking theories generally.

In a typical Bonjour example, Sally the clairvoyant has a special faculty or power that tracks the truth about p (the whereabouts of the president). Sally clairvoyantly believes that the president is in New York. The president is in New York. Sally is aware of a massive media blitz orchestrated by the secret service saying that the president is at that time in D.C. The cover up is designed to thwart an assassination attempt on the life of the president. Sally has no reason to think that she has clairvoyant powers. Indeed, she has good reason to think that there are no clairvoyants. Sally also has no good reason to believe that the media blitz is a cover up or hoax. Bonjour claims that Sally would be epistemically irresponsible to continue to believe that the president is in New York in the face of the overwhelming evidence to the contrary. Let us suppose that, because of her clairvoyance, Sally is unable to shake her belief. Bonjour claims that despite the fact that Sally is tracking the truth about the whereabouts of the president, she does not know that he is in New York because she is being epistemically irresponsible in her belief.

I disagree. For reasons central to the heart of tracking theories, I would maintain that Sally knows the president is in New York. Her belief is true. It is not an accident that her belief is true. There is no close world in which were she to apply her clairvoyance, it would lead her astray. Her belief is produced by a nomically reliable primary cognitive process. What more could one want from a cognitive process but that it be nomically guaranteed to produce truth?[21] What Sally doesn't have is good reasons to believe. She doesn't have evidence that she has a special power. What she should do (if she were able) is resist such thoughts, unsupported as they are by any evidence we would normally take to be a sound basis for belief. Normally if one has no reason to believe that one has such a capacity, it would be rational to suspend belief. And Sally would suspend belief too, if it were not for the overpowering clairvoyant impulse.

Bonjour needs to *show* that it is impossible for a belief such as Sally's to be knowledge. Bonjour is pitting his intuition that such cases cannot be knowledge against the intuitions of the tracking theorists that such cases can be knowledge. However, to *show* that Sally doesn't know the president is in New York, Bonjour has to move beyond the level of pre-theoretic intuition. Otherwise, there is a corresponding equally plausible intuition on the side of the tracking theorists. Consider my dog Raven. Raven knows that a squirrel ran up the tree. Her belief is based solely on her untutored cognitive processes. She has no secondary checks

[21] Of course, Bonjour would answer that one wants "epistemically responsibly formed beliefs." But that is what we are denying is advantageous in such cases. Adding responsibility does nothing for the agent with respect to truth or likelihood of truth in one's epistemic environment that is safe for knowledge. For more on this see Adams, 1986, 2003, and 2004.

in place on whether her cognitive processes are reliable or whether her relying on them is justified. That is, she has no mechanism for confirming whether her cognitive mechanisms are reliable. Yet nearly everyone agrees that Rave can know the squirrel ran up the tree.[22]

11 Dretske's Information-Based Account of Knowledge

Dretske was very taken with Martin's counter-example to Dretske's "conclusive reasons" account. Though he may have believed there were ways around it, he sought a different approach for his version of a tracking theory. He turned to information theory. In this section, I will just briefly give the basics of his information-theoretic account of knowledge and then discuss a small sample of objections that have been raised to it (and explain why I think they fail).

To adapt information theory to a format friendly to a theory of knowledge, several matters need to be resolved. For example, to know that Bush was elected president involves information being generated by the event of his election. It also involves transmission of that information to a prospective knower S. S must detect physical events that carry that transmitted information, and those events must cause or sustain S's belief that Bush was elected.

Let's begin with generation of information. An event's occurrence generates information. How much is generated is a function of how likely was the event's occurrence. The more likely an event, the less information it generates--while the less likely the event, the more information it generates. Different ways of classifying events may result in different amounts of information generated. And there are many different ways of trying to measure or quantify amounts of information. Dretske follows the communication industry standard (Weaver & Shannon, 1949) of measuring information in bits (*bi*nary digi*ts*), representing the number of binary partitions necessary to reduce a collection of equally probable outcomes to one (e.g., beginning with 8, a three-step reduction to 4, to 2, to 1 = 3 bits). The amount of information generated at a source s by the reduction of n equally likely possibilities to one is represented: I(s)=log n (base 2). Here I(s) represents the average amount of information generated at a source by a reduction of equally likely events. If the range of possible events at the source s_1, s_2, ..s_n, are not all equally likely, then the amount of information generated by the occurrence of s_i is: $I(s_i) = \log 1/p(s_i)$ (where p = probability). So, for example, suppose ten persons apply for a job and nine are from inside the company, one from outside. If s_1 is the selection for the job of someone outside the company, then $I(s_1)= \log 1/.1 = 3.33$ bits of information. For contrast, selection of someone from inside the company, s_2 would generate 1/.9 = .15 bits of information.

[22] Of course, some will claim that this is cheap and dirty animal knowledge, not fancy scientific human knowledge. That's fine with me. Let Sally and Raven have only the cheap and dirty variety of knowledge. That's knowledge enough.

Next, let's consider information flow or transmission. For information at a receiving point r to be about a sending point s, there must be dependence between the events at r upon those at s. Suppose at s there are eight candidates equally likely to be selected. A selection of Susan generates 3 bits of information. Suppose at r there are eight equally likely names that may be put on the employment forms in the employment office. A selection of "Susan" generates 3 bits of information. But there would also be 3 bits generated if, mistakenly, the name "Tony" were placed on the employment forms. Clearly, though this amount of information is the same it is not the information that Susan was selected. We want the information at r to be about the events that transpired at s. Letting "$I_s(r)$" represent this information, $I_s(r) = I(r) - \text{noise}$. Noise is the amount of information generated at r that is independent of what happens at s (not about s), and when "Tony" is placed on the forms, but Susan was selected, the noise = 3 bits. Thus, no information about s arrives at r.

Now for our purposes, the import of these formulae for calculating amounts of information is not so much the absolute values of information generated or transmitted by an event, but the conditions necessary for transmission. For most events it would be difficult or impossible to determine the exact probabilities and ranges of possibilities closed off by an event's occurrence. What is important is whether one receives at r as much information as is necessary to know what happened at s (under a relevant specification). For a signal or message to carry the information that Bush was elected, it must carry as much information as was generated by Bush's election. We know this is more information than that a Republican ran for office, and more than that someone was elected. Calculating exactly how much information is generated by Bush's election is not as important as determining under what conditions the information that does arrive carries the information that Bush was elected. This is what Dretske calls the informational content of a signal.

Informational content: A signal r carries the information that s is F = The conditional probability of s's being F, given r (and k), is 1 (but, given k alone, less than 1). K is a variable that takes into account how what one already knows may influence the informational value of a signal. If one knew nothing, k would go to zero. If I know that Vice President Cheney is from Texas or Wyoming, and I learn that he is not from Texas, I thereby have the information that he is from Wyoming. If you hear that he is not from Texas, but don't already know Wyoming is the only other possibility, you do not thereby receive the information that he is from Wyoming.

This account of the informational content of a signal has important virtues. If a signal carries the information that Bush was elected, then since the conditional probability that Bush was elected given the signal is 1, then *Bush was elected*. Hence, the account gives information a connection to truth. Clearly it will also be the case that the signal carries as much information about s, $I_s(r)$, as was generated by the fact that Bush was elected. Noise about the fact *that Bush was elected* is zero. Hence, the account gives us a way to understand transmission or flow of information of a specific propositional (factual) content from source to receiver – not just amounts of information.

Finally, we can give an information-theoretic account of knowledge: K knows that s is F = K's belief that s is F is caused (or causally sustained) by the information that s is F. Dretske says that this is intended to account for perceptual knowledge only, that is, perceptually knowing of something s that it is F. Knowing, by my current visual experiences of my computer, that it is on, would count as perceptual knowledge. And so would my knowing that the coast is clear, by hearing three knocks, on a pre-arranged signal. However, here my knowledge also involves knowing the pre-arranged signal. I know that the coast is clear by hearing the three knocks. While brief, this gives us a good basis for understanding Dretske's move to the information-based version of his tracking theory of knowledge.

Foley (1987) says that any informational account that relies upon causation of a belief by information will be susceptible to well-known problems of causal deviance. Ironically, this is exactly one of the things an informational account was designed to avoid with the notion of information, because an information channel must screen off causal deviance. So what is Foley's example and does it work? Foley focuses on Dretske's example of three quick knocks at the door causing a spy to believe that the courier has arrived and carrying that information, as well. Foley modifies the case so that it involves a wayward causal chain, as follows. The spy suddenly goes deaf. Then come three knocks carrying the information about the courier's arrival. The knocks cause the spy's partner to trip, causing a box to fall on the spy's head, in turn jarring the spy's brain in such a way that he suddenly comes to believe that the courier has arrived.

Now it is quite clear that an information-theoretic account needs the *proximate* cause of the belief (or sustaining cause) to carry the information that p. In Foley's example, that is not the case. Even if the three knocks contained the information that p and for some strange reason the spy's partner would not have tripped unless the knocks did contain that information, the rest of the story doesn't preserve information. The communication channel has been broken by the time the box falls on the spy's head and his brain is jarred. The conditional probability that the courier has arrived, given the knocks is 1, but the conditional probability that the courier has arrived given that the events in the jarred brain have occurred is not 1. This is because the spy's brain, since jarred sufficiently hard by the box, might be in that state (seeming to hear three knocks) even if the courier had not arrived and there had not been three knocks at the door. Let me explain.

For the purpose of addressing Foley's example, it is important to discuss the matter of an information channel. An information channel condition is any fixed condition (other than conditions existing at the source s or receiver r) which, by variation of its value, would be able to introduce noise between source and receiver. In order for the information that a is F to flow from source s to receiver r, there must be as much information about a's being F arriving at r as is generated by its occurrence at s. That will be possible only if the channel conditions that permit information to flow remain fixed. They must not themselves vary,

thereby generating information or non-redundant information. Provided that the channel conditions remain fixed, they do not generate information in the form of noise on the channel, with respect to the information that a is F.

Let's consider an example. Suppose Al has a metal detector that emits a tone when metal is within ten inches of its detection surface. For its tone to carry the information that there is metal present, the detector depends on several channel conditions. The power supply must be adequate and charged. The magnetic field that detects the metal must be in place. The wires that activate the tone must be well functioning (no shorts or breaks in the wire). And so on. Were any of these conditions not to remain fixed (a short circuit, say), the detector may emit a tone when there is no metal being detected (even if there is metal present). Therefore, the detector's tone will carry the information that there is metal present only when all channel conditions are *fixed*. In virtue of these fixed conditions, information can flow from the source (detection of metal in the magnetic field of the surface) to the receiver (Al hears the emitted tone).

Notice that the tone carries information about the presence of the metal, but it does not also carry information about the channel conditions of the metal detector. We can use the tone to tell us about the channel conditions, as we will now see. But we cannot use it to tell us about its channel conditions *and* about presence of metal *at the same time*. If we know in advance that there is metal (or is none), we may check to see whether the detector emits a tone (or does not). Thereby, we can use old information about what we already know (there is/is not metal present) to gain information about the channel conditions of the detector. Now, because we *already know* whether or not there is metal present, we can test to see whether or not the detector is working properly. We can then tell that it is working properly and its channel conditions indeed are fixed (or tell that it is broken because its channel conditions are variable). But if we do not already know in advance about the presence of the metal, the tone carries no (new) information about the channel conditions themselves.

This is not to say that channel conditions last forever. Metal detectors break or wear out. But the point is that for the detector to be a source of information about the presence or absence of metal, its channel conditions must be fixed. When are they fixed? When they generate no (or no new) information.

It is now easy to see that Foley's example is one where noise is introduced into the communication channel because the channel conditions of the man's jarred brain are not fixed. The internal workings of such a brain introduce variability and noise to the system. The same would be true if there were a short in our metal detector (when there was metal present). In neither case (hearing at tone / seeming to hear three knocks) would the relevant piece of information be carried by the tone (or the auditory experience as of three knocks).

Similar remarks apply to an example by Plantinga. Suppose K suffers from a brain lesion that causes K to believe a variety of mostly false propositions. It also causes K to believe that he has a brain lesion, but K has no evidence for this be-

lief. Nonetheless, referring to Dretske's analysis above, Plantinga maintains that the "probability on k & *K is suffering from a brain lesion* is 1" (Plantinga, 1993, p. 195). Notice that this is not a case of perceptual knowledge. There is no signal that K perceives and which carries the information that he has a brain lesion. But even setting this aside, it is clear that the brain lesion is not a fixed channel condition. It introduces noise to K's cognitive system. Indeed, Plantinga grants that the lesion causes K to believe mostly false propositions (a sure sign of noise). Therefore, despite Plantinga's claim, it simply is false that the conditional probability that K has a brain lesion, given that K believes he has a lesion, is 1. K's beliefs do not track the truth, as Nozick would say. *K's suffering from a lesion* guarantees that he has a lesion, but not his *belief that he has a lesion*, not even if his belief is caused by the lesion (as for a tone from a metal detector with a short caused by its detection of metal). Therefore, the example fails.

However, Plantinga could fix the example so that there is no noise on the channel. So we should consider what an information-based account of knowledge should say to such a case. Indeed, there is a case by Lehrer that will do (1990, pp. 163-4), his Truetemp case. Suppose that quite without his being aware, Mr. Truetemp has a high-tech, belief-producing thermometric chip implanted in his brain. The chip measures the surrounding temperature and then directly enters a belief that it is that temperature into Truetemp's belief box, if you will. Lehrer and Plantinga would agree that Truetemp does not know the temperature is 56 degrees Farenheit, when it is 56 degrees Farenheit, and Truetemp believes that it is. As before, there is no signal that Truetemp perceives and which carries the information. The information is mainlined into Truetemp's belief box, as it were. Suppose Truetemp is cognitively unable to withhold belief despite having no evidence that his beliefs are correct.

Now, unlike in Plantinga's example, there do seem to be fixed channel conditions at work here. The thermometric-cognition mechanism is a fixed condition (let's suppose that it has been in place for years). Since fixed, that it is working reliably generates no (or no new) information. It takes variation in the source (ambient temperature) and faithfully delivers information about it to the receiver (the belief box). It is an information delivery system no less faithful than the delivery systems of our senses. The probability that ambient temperature is 56 degrees Farenheit, given Truetemp's belief that it is, is 1. The main difference from our senses is that it skips the step of causing a conscious sensation that then causes a reliable belief about what is sensed. The device in Truetemp skips the conscious sensory step, but its delivery of information is, we are presuming, as stable and faithful as any normal sensory perception by standard perceivers under standard conditions.

It would certainly be possible for an information-based theorist to bite the bullet and insist that this is a case of knowledge, albeit a very unusual kind of case indeed. All that Plantinga or Lehrer have to fight this stand is a conflicting intuition and a conflicting theory. Since one probably should place little stock in the conflict of intuition, in the end all they may have is a conflicting theory. Which

theory is correct will have to be determined by looking at overall consistency and explanatory power. Saying that this is indeed a case of knowledge may be true and prove not to be so strange, in the long run. It may come to grow on one. Only time will tell.

Lastly, Keith Lehrer (1990, p. 184) argues convincingly that his theory has as a consequence that in, Kripke's red barn case, one does not know that the red barn is a barn, and Lehrer seems to maintain that this is a problem for informational accounts like Dretske's. Of course, above I have defended Nozick's account against this example of Kripke's. The informational account, too, is not susceptible to it. One can know both that something is a barn and red because it is a *channel condition,* due to the nature of Kripke's stipulation in the example, that red barns cannot be faked. In such an environment, something's looking like a red barn carries the full information that something is a barn and is red. If *that look* causes and sustains the belief that the barn is red, then one knows that it is a red barn on the informational account. The conditional probability of its being both red and a barn, given its red-barn appearance, is 1. So if one knows both, as I believe to be the case, then Lehrer's account of knowledge would be false, as he maintains strongly that one would *not* know both on his theory.

12 Conclusion

This concludes a brief survey of some of the best and most prominent attempts to kill off the tracking theories of knowledge. None of them succeeds. True, this essay does not contain an exhaustive list of attacks on these theories, but it does contain a representative sample of the kinds of objections to tracking theories most often cited in the literature or orally at conferences. And it does show why these specific objections do not work. Tracking theories remain viable accounts of knowledge.[23]

References

Adams, F. 1986. The Function of Epistemic Justification, *Canadian Journal of Philosophy* 16: 465-492.

Adams, F. 2003. Epistemology. In R*eflections on Philosophy*, eds. McHenry & Yagisawa, NY: Longman's Press: 81-101.

Adams, F. 2004. Knowledge. In *the Blackwell Guide to the Philosophy of Information and Computing*, ed. Luciano Floridi, Oxford: Blackwell Publishing: 228-236.

Adams, F. & Clarke, M. 2005. Resurrecting the Tracking Theories, *Australasian Journal of Philosophy*, 83, 207-221.

Armstrong, D. 1973. *Belief, Truth, and Knowledge*, Cambridge: Cambridge University Press.

Bonjour, L. 1980. Externalist Theories of Empirical Knowledge, *Midwest Studies* 5: 53-73.

[23] I would like to thank Gary Fuller for helpful conversation on some of these ideas going back as far as 1989. Thanks too to Murray Clarke, who helped me think through many of these ideas for our joint paper (Adams & Clarke, 2005). I'd also like to thank conference discussants Robert Audi, Ben Caplan, Phil Hanson, Igal Kvart, Danny Goldstick, Jay Rosenberg, Bill Seager, Ernie Sosa, and Crispin Wright for questions and comments. And special thanks to Claudio de Almeida and to John Barker and Fred Dretske for conversation and comments on earlier versions of the paper.

DeRose, K. 1995. Solving the Skeptical Problem, *Philosophical Review* 104: 1-52.
DeRose, K. & Warfield, T. 1999. *Skepticism,* Oxford: Oxford University Press.
Dretske, F. 1970. Epistemic Operators, *Journal of Philosophy* 67: 1007-1023.
Dretske, F. 1971. Conclusive Reasons, *Australasian Journal of Philosophy* 49: 1-22.
Dretske, F. 1975. Review of Armstrong's *Belief, Truth, and Knowledge*, *Journal of Philosophy 72*: 793-802.
Dretske, F. 1981. *Knowledge and the Flow of Information,* Cambridge, MA: MIT/Bradford.
Dretske, F. 2004. The Case Against Closure. In *Contemporary Debates in Epistemology*, eds. Ernest Sosa, & Matthias Steup, Oxford: Blackwell publishers.
Foley, R. 1987. Dretske's 'Information-Theoretic' Account of knowledge, *Synthese*, 70, 150-184.
Gettier, E. 1963. Is Justified True Belief Knowledge?, *Analysis*, 23, 121-123.
Goldman, A. 1967. A Causal Theory of Knowing, *Journal of Philosophy*, 64, 355-372. (Reprinted in G. Pappas & M. Swain (eds.). *Essays on Knowledge and Justification*, Ithaca: Cornell University Press, 1978. Page references to this version.)
Goldman, A. 1976. Discrimination and Perceptual Knowledge, *Journal of Philosophy* 73: 771-791. (Reprinted in Pappas & Swain (eds.) *Essays on Knowledge and Justification*, Ithaca: Cornell University Press, 1978. Page references to this version.)
Goldman, A. 1979. What is Justified Belief? In *Justification and Knowledge*, ed. George Pappas, Dordrecht: Ridel: 1-23.
Lehrer, K. 1990. *Theory of Knowledge.* Boulder: Westview Press.
Luper-Foy, S. (ed.) 1987. *The Possibility of Knowledge: Nozick and His Critics,* Totowa, NJ: Rowman and Littlefield.
Martin, R. 1975. Empirically Conclusive Reasons and Scepticism, *Philosophical Studies* 28: 215-217.
Martin, R. 1983. Tracking Nozick's Sceptic: A Better Method, *Analysis* 43: 28-33.
Nozick, R. 1981. *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
Pappas, G. & Swain, M. 1973. Some Conclusive Reasons Against 'Conclusive Reasons,' *Australasian Journal of Philosophy* 51: 72-76. (Reprinted in Pappas & Swain (eds.). *Essays on Knowledge and Justification*, Ithaca: Cornell University Press, 1978. Page references to this version.)
Plantinga, A. 1993. *Warrant: The Current Debate.* Oxford: Oxford University Press.
Sosa, E., forthcoming. Relevant Alternatives, Contextualism Included, *Philosophical Studies*.
Steup, M. & Sosa, E. 2005. *Contemporary Debates in Epistemology*. Oxford: Blackwell.
Weaver, W. & Shannon, C. 1949. *The Mathematical Theory of Communication*. Urbana: The University of Illinois Press.