

RESURRECTING THE TRACKING THEORIES

Fred Adams and Murray Clarke

Much of contemporary epistemology proceeds on the assumption that tracking theories of knowledge, such as those of Dretske and Nozick, are dead. The word on the street is that Kripke and others killed these theories with their counterexamples, and that epistemology must move in a new direction as a result. In this paper we defend the tracking theories against purportedly deadly objections. We detect life in the tracking theories, despite what we perceive to be a premature burial.

Historical Prelude

Tracking theories¹ of knowledge, though quite popular among naturalists, cognitive scientists, and other scientifically-minded philosophers, have been under attack almost from their inception. Specifically, we have in mind the attacks of Pappas and Swain, Ray Martin, Saul Kripke, Laurence Bonjour, and Keith DeRose.² We don't think any of these attacks on tracking theories work. In this paper, we intend to take a step towards resurrecting the tracking theories of knowledge. We will not be able to focus attention on all of the relevant attacks on tracking theories in a single paper.³ However, we will be able to devote attention to five prominent attacks and to explain why these attacks do not succeed. In so doing, we hope to quiet some of the premature celebration of the demise of tracking theories of knowledge that we have noticed in the literature and at conferences in recent years.

We will focus mainly on the theories of Dretske and Nozick. These are the ones that we think have the best chance of surviving attack. Yet *tracking theories* as a type consider knowledge to be a kind of non-accidentally true belief. The non-accidentality of the truth of the belief is sustained by sets of

¹We include among the tracking theorists Armstrong [1973], Dretske [1971; 1981], Goldman [1976], and Nozick [1981], but we will focus mainly on Dretske and Nozick, even though the work of Armstrong and Goldman was certainly very influential as well.

²Why just these attacks? In our experience, these are some of the most voiced complaints about tracking theories and ones thought to herald their demise. We also think that by answering these challenges in particular we may help those who are inclined to be sympathetic to tracking models to move the enterprise forward. We find accounts that herald the demise of tracking models and then try to move epistemology forward to be less promising. We will not be able to discuss all attacks on the theories. So, for instance, we won't be able to discuss rejection of tracking theories because they deny closure. We realize rejection that knowledge is closed is highly controversial, and although we agree that knowledge is not closed under deduction, to defend that would take more space than we have available. So we deal here only with issues other than closure. However, we agree with Dretske's latest defence of non-closure [Dretske 2004].

³We are tempted to do a Shope-style book chronicling the failure of every last one of them.

conditions that (variously articulated) make for the distinctiveness of the tracking theories. For Armstrong [1973], tracking was being in a belief state such that one's evidence and belief serve as a 'completely reliable sign to thing signified' [182]. For Dretske [1971] tracking was having a reason R upon which one's true belief was based, such that given R and the repeatable kind of circumstances C in which one finds one's self, it is circumstantially not possible to have a false belief that p, when that belief is based upon R. And later for Dretske [1981] tracking was having one's true belief that p caused or sustained by the information that p. Information was defined such that it could not be false. Thus, there is a conditional probability of one that a belief that p is true, when caused or sustained by the information that p. For Nozick [1981], tracking was having a nomically reliable method of belief formation M such that one's belief that p is based upon M and, when it is, one's belief co-varies with (tracks) the truth. Were p true, one would believe it. Were p false, one would not believe it. And for Goldman [1976; 1979], tracking amounts to having one's true beliefs formed by cognitively reliable belief forming processes (where reliability is measured by the propensity of producing true beliefs as outputs). Notice that what these so-called 'tracking' theories have in common is that they do not involve doxastic decision procedures about when one should (or should not) believe something. They are not couched in terms of the justification of beliefs (with the possible exception of Goldman who seems to equate tracking with a kind of 'external' justification). Instead, the 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 (screening off mere luck that the belief is true). Knowledge is a matter of having one's true beliefs caused or sustained by the right reasons, factors, or evidence, where these reasons, factors, or evidence, nomically track (co-vary with) the truth.

The Martin Counterexample

Against Dretske

Dretske himself was so taken with an example of Ray Martin's [1975] that he began looking for a new framework for his account of knowledge. He turned to information theory [Dretske 1981] because he thought that it avoided the type of problem raised by Martin. He also thought that with information theory as a framework for knowledge, he could do much more.⁴ He could show the important interconnection between issues in epistemology and philosophy of mind. While we don't disagree that information theory is useful for achieving these ends, we also don't think that the Pappas and Swain and Martin examples are counterexamples to Dretske's original conclusive reasons approach. Indeed, we think that if the information based approach to knowledge is able to avoid the difficulties raised in these

⁴Personal communication with Dretske.

attacks, then the conclusive reasons approach is able to avoid them for the same types of reasons. We will now attempt to explain how tracking accounts avoid the Martin example and the Pappas and Swain examples.

In a set of papers, Ray Martin offered the same counterexample to the tracking theories of Dretske [1971] and Nozick [1981]. Against Dretske, the example was designed to be problematic for the notion of a conclusive reason. Dretske had offered an account of empirical knowledge according to which a person *S* knows that *p* when *S* believes that *p* on the basis of a conclusive reason *R*. *R* was assessed as conclusive, relative to a fixed set of circumstances *C*, when given *R* (in *C*) it is not possible to have *R* and not *p*.⁵

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.

Against Nozick

Before responding to the Martin example against Dretske, we will point out that Martin used the same example against Nozick. It is easy to see why the example fails against Nozick and this will help us to explain why it fails against Dretske (for similar reasons). So we begin there.

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

⁵Sosa and Luper have pointed out that Dretske's conclusive reason condition 'if *p* were false (in *C*) then *R* would not have indicated that *p* was true' does not support closure, but that another condition 'if *R* indicates that *p* is true, then *p* would be true' does support closure. The latter condition, however, is virtually identical to Armstrong's [1973] condition. We do not believe that Armstrong himself, or Sosa [2004], or Luper [1987] ever adequately responded to the criticisms that Armstrong himself presented to his own condition. However, we do find the tracking theories of Nozick and Dretske to adequately address those original counterexamples to the latter formulation of the connection between *R* and *p*, and we cannot discuss the larger issue of closure here. That would take a separate, longer treatment itself.

know that Gumshoe won. We agree that Zack does not know, but we disagree that method M tracks the truth.

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 that in close possible worlds Zack's method M tracks the truth. 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 Gumshoe wins, no matter what. His method is somewhat irrational. So 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).⁶

Now we return to the example as used against Dretske. For reasons very similar to why Martin's example fails against Nozick's account, it fails against Dretske's.⁷ 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. We claim that Martin changes the circumstances from C to C' after the first race. Then Martin claims, given that Tagalong lost (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'). 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).

We suspect that one sympathetic to Martin would claim that our reply makes what counts as the relevant circumstances slippery or unprincipled. Our response is that this is not so. Here is a structurally similar example, familiar from subsequent literature. Suppose that in circumstances C (lots of barns and fake barns all around) John says that Tracy knows there is a barn in the field by its look. The look of a barn R , is not a conclusive reason,

⁶Martin may claim that the world where Tagalong wins and Gumshoe loses is not a possible world or is a world too far away to be relevant. We disagree. Given the standard understanding of possible worlds and their long history of use in epistemology, this would be a desperate move on his part.

⁷Martin's example also fails to work against Dretske's [1981] information theoretic account of knowledge, and it fails for similar reasons to why it fails against his conclusive reasons account. However, we will not go into the details here beyond saying that for Zack to know that Gumshoe won he needs to receive the information that Gumshoe won. He does not receive that information from the payout alone. He receives only the information that Tagalong or Gumshoe won, but not which. To receive the information that Gumshoe won, Zack would have to receive the information (or know) that Tagalong lost. And Zack cannot, without changing channel conditions, receive the information that Gumshoe won (or that Tagalong lost) from the payout alone.

relative to C at the time John makes the claim. At that time (in C), it is possible to have R and not p (possible for it to look to Tracy as if there were a barn when there is not). But suppose that a few hours after John's utterance the fake barn police remove all of the fake barns to Sweden. One would be wrong to maintain that John's utterance was true after all because now it is not possible to have R and not p. We suggest that it is clear that Tracy did not know there was a barn. It would be a mistake to claim the status of R's being a conclusive reason does not *change* between the time of John's utterance and the time the barn police permanently banish the fake barns to Sweden. It was for this reason that Dretske stipulated that R's status as a conclusive reason is relative to a set of *fixed circumstances* C. To us this means there must be some local stability in one's environment for one to have a conclusive reason (or for the assessment of the truth value of subjunctives, generally).

We admit that there is some vagueness over how close the relevant alternatives (fake barns) can be before one loses knowledge that there is a barn nearby. This is a general difficulty about metrics for nearness of possible worlds, when testing counterfactuals. But we insist that it is not unprincipled or unruly whether the circumstances have changed when the fake barns are here in North America inhabited by Tracy vs. when they are banished to Sweden and ecologically bound there. Similarly, we insist that it is not unprincipled or unruly that the assessment of 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').⁸

The Pappas and Swain Examples

Pappas and Swain [1973] argued that Dretske's requirement that one can only know that p if one has a conclusive reason R for p was *too strong*. Through two examples, they argued that one could know that p with less than a conclusive reason R. We will show why neither example works.

⁸In correspondence, John Barker pressed the issue of what we would say if one insists that C' is an appropriate context to assess a conclusive reason (maybe not Zack's, but how about John's thinking that Gumshoe won based upon the payout R)? In that case we would maintain that, as for Nozick, the conclusive reason counterfactual should be assessed across close worlds. It is not true that Zack would not get a payout unless Gumshoe won. For there are close worlds where Tagalong won. This response involves reinterpreting the semantics for conclusive reasons in a way different than Dretske originally suggested. But we feel this is consistent with the spirit of the tracking approach in Dretske's account. Furthermore, when he originally wrote his article, Dretske did not have available the possible world semantics for counterfactuals or he might well have interpreted the conclusive reason counterfactual this way from the start.

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 and 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.

Hologram Cup Example

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. 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 and 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 is false' [64].

We agree that the conditional is false, and for that reason we do not see how George could possibly know that p. 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. This example is no different than the ones Goldman [1976] popularized with his barn facades. Virtually, the entire profession agrees that George cannot know that he is seeing a barn if indistinguishable fake barns are all around. Hence, since this is structurally similar, we maintain that it is not a case in which George knows he is seeing a cup, even if he is seeing a cup.

If, however, the hologram machine is in a remote location (as are Goldman's fake barns when isolated in Sweden, a place to which George will not go), then we see no reason why George does not know that there is a cup on the table. We also see no reason why the corresponding subjunctive would be false, *viz.*, that George would not be having an experience as of a cup unless there were a cup on the table, in these circumstances. Therefore, the first example is just not a counterexample to the conclusive reasons approach.

Generator Example

In their first example, it still may have been somewhat unclear just where the hologram machine was. So in their second example, they removed that ambiguity. They placed a backup generator in the basement of Marshall's friend Louise. 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.

What Marshall does not know is that Louise has a backup generator in the 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 the house of Louise, and his background knowledge from working at the power company). Pappas and Swain insist that 'surely we do not want to say that the fact that his friend 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].

That is *just* what we want to say. If Marshall is inside the house of Louise and is not looking out on the street lights across the whole city, seeing only the lights on in the house of Louise 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).

We suspect that Pappas and Swain put more into R than *just seeing the lights on in the house of Louise*. They seem to include Marshall's seeing the streetlights in the city on the way to the house of Louise. Since they underdescribe the example, let us suppose that Marshall enters the house of Louise and simultaneously sees her lights and, from her front room, the lights on across the city. Given this addition to R , it may seem to make it more likely that Marshall knows that the power company is supplying the power to the house of Louise. For not all of the other houses in the city visible from Louise's front room have back-up generators. We suspect that it was this type of reasoning that led Pappas and Swain to believe that Marshall would indeed know that p . However, we hold that the same reasoning would lead one to maintain that if these are stable circumstances in the local environment, then R would not be the case unless p (once the addition of the other lights to R is made). With that condition, Marshall would know that p , but Dretske's conclusive reason condition would also be *satisfied*.⁹

⁹In personal communication, Dretske favored this result when the example was further specified in this way.

Still we maintain that Marshall would not know that p . Joe may still be in the basement giving the back-up generator its annual test. 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. In our neighbourhoods accidents occur frequently near our houses, 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 (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). That subjunctive conditional is indeed false, but precisely because it is false, there is no way that Marshall can know that p , by R alone. Therefore, this example too is not a counterexample to the conclusive reasons approach.¹⁰

We conclude that neither of the Pappas and Swain examples show that the requirement of a conclusive reason is too strong. We will now turn to other attacks on the tracking models that have received considerable attention in recent years.

The Kripke Examples¹¹

The Red Barn

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 colours 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 that it was 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. Nozick is very clear that his tracking conditions must be relativized to methods. For example, otherwise only God could satisfy the condition that

¹⁰Nor is it a counterexample to Nozick, as should be clear.

¹¹Kripke gave these examples at a session of the APA in the 1980s, but the paper was never published. It has been widely circulated and almost everyone we meet thinks the examples work, but they do not work—as we shall now begin to show for a subset of them.

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.

Let's look at a related example¹² that has the same structure and makes the point clearly. 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. 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 Beetle carries the information 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 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.¹³ So contrary to Kripke, we maintain that 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.

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 faithfully carries both pieces of information. Since Peg is using the reddish barnish look to form the belief that there is a red barn, she satisfies Nozick's tracking condition that 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 barn, she would not be using the same red barn look method. This leads us to believe 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 that we are not claiming that there are no cases where one's knowledge that something is a red barn would come apart from one's knowledge that something is a barn. We suspect that one could employ different methods with this result. 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 pitch darkness in the

¹²This example was based on one by Goldman and appeared in Dretske's [1975] review of Armstrong.

¹³It is true that if she used the 'looks like a Saab' method, she might believe a non-Volkswagen to be a Volkswagen. But this is irrelevant to whether she knows it is a Volkswagen when using the VW-Beetle method.

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.

Someone may object that we 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 we suspect people will balk at the notion of a red-barn method of belief formation. Our reply is that in the wild, 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 Deceased Dictator

The strongman dictator of a totalitarian country dies. The state-run newspaper prints that the dictator is dead (p). Later, fearing a coup, the newspaper prints a retraction saying that there was a mistake. They show pictures of the dictator alive and well. Ken returns from a long trip out of the country. He reads the first newspaper, but then goes to sleep due to jetlag. The rest of the country read and believe the retraction. Ken, like others, would believe the retraction if he were to read the second newspaper. So while Ken sleeps he loses his knowledge that the dictator is dead because he fails Nozick's condition that if the dictator were dead, Ken would believe it. Ken would not believe it, if he read the retraction.¹⁴

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), Ken knows $p \& q$ but not p (as we've seen above, Ken fails to know p). Supposedly, Ken knows $p \& q$ because his belief that $p \& q$ does satisfy the condition that if $p \& q$, Ken would believe $p \& q$. This is true because for Ken to believe that he had read an uncontradicted report of the dictator's death, he would have to have read an uncontradicted report. So Ken would have not to have read the second newspaper (keeping q true).

In reply, we will suggest that this example fails because Ken knows neither p nor $p \& q$, on Nozick's own conditions. To see why, we maintain that Ken will either be savvy and not believe something merely because it appears in the state controlled newspaper or gullible, and he will believe everything that appears there. So suppose Ken is savvy. Then we maintain that Ken does not know $p \& q$. He does not because he fails the fourth condition. It is not true of Ken that if $p \& q$ then he would believe it. A savvy Ken rejects many things that he reads in the state run newspaper. So, suppose that Ken is gullible. Then he believes everything that appears in the state run newspaper. In that case Ken fails condition three, *viz.*, that if $p \& q$ were not true, Ken would not believe $p \& q$. A gullible Ken believes falsehoods

¹⁴This example turns on Nozick's fourth condition only. So we will focus attention on that condition in presenting the example.

that are printed in the newspaper. In that case, Ken cannot (without independent confirmation) learn anything from a state run newspaper that is highly unreliable. So on Nozick's own conditions, we maintain that the dictator example is not a case where Ken knows $p \& q$ but does not know p (as intended by Kripke).

For good measure, we will discuss one more of Kripke's purported counterexamples to tracking theories. We will not attempt to exhaustively cover them here, but we maintain that each fails. We hope to have demonstrated that we could go on and demonstrate the ineffectiveness of the rest.

The Sloppy Scientist

Steve is looking for a cure for a particular disease. Steve selects a drug D for several trials of treatment. The trials are effective in curing the disease and Steve comes to believe that drug D will cure this disease (p). Unfortunately, Steve is a sloppy scientist. He forgets to run a control using a placebo in a counterbalancing exercise. As it turns out, had he used the controls, the placebos would have been negative.

Kripke maintains that this is a counterexample to Nozick because it satisfies Nozick's conditions for knowing, but that Steve does not know that p . Steve does not know because, claims Kripke, Steve's procedure is irrational. Steve is ignoring standard control procedures for sound experimental design. Despite this, Kripke insists that Steve satisfies Nozick's tracking conditions. Specifically, if p were not true Steve would not believe that p , and if p were true, Steve would believe that p (given his sloppiness).

Kripke's example is exploiting the intuition that one cannot learn something from an unreliable method or procedure. And Kripke is suggesting that due to sloppiness, Steve's experimental procedure is unreliable. Kripke is also suggesting that Steve's procedure is somewhat irrational and that one cannot know something if one is being somewhat irrational. This latter claim concerning irrationality is reminiscent of Bonjour's attack on externalist theories and we will say more about this it when we discuss Bonjour. Here we will focus our comments on the suggestion that Steve does not know that p because Steve's procedure is unreliable because Steve is sloppy. We maintain that being sloppy does not preclude latching onto knowledge or a reliable method.

There are at least two kinds of sloppiness. One kind would exist if Steve were sloppy in screening patients for the disease or in the experimental set up where he administered the drug. Call this primary process sloppiness. If Steve didn't control for variables and background conditions in the administering of the drug itself, then we would agree that Steve does not know that p . But in this kind of case, the sloppiness would cause Steve to fail Nozick's conditions. There could be false positives or false negatives in the administering of drug D and collecting the results. Please note that in Kripke's example, Steve was not sloppy in this primary way.

Another kind of sloppiness is secondary process sloppiness. Here Steve fails to provide an independent check on his results. We see this as a failure of confirmation not of information. This is not to say that Steve should not be banned from the lab, but it is to say that it matters where one lay the blame. Steve is at fault for not attempting to verify his information. His primary process procedure of administering the drug and controlling for variables, is impeccable. Where he is sloppy is in his neglect of secondary process controls. We maintain that, perhaps surprisingly, Steve does know that p because he satisfies Nozick's tracking conditions. Steve does receive the information that drug D cures the disease. What he does not do is gather sound confirmation that he has received that very information.

We suspect that dissenters will say that scientific knowledge only comes into existence after the secondary process controls are in place. Isn't that what double blind experiments are all about? But we maintain that this is what confirmation is all about. It is part and parcel of attempting to confirm a hypothesis that says that a particular theory or claim is true. It is like cross-examining a witness. Prior to the cross-examination, a witness may be telling the truth and be perfectly reliable. But the jury won't know that or suspect it until *after the cross-examination*. This does not mean that the witness was not reliable until after the cross-examination. But that is the kind of thing Kripke is suggesting with this example.

DeRose's Abominable Conjunction

We have found that 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.

We will not try to recount all of the reasons to reject that knowledge is closed or to defend that it is. There are excellent ways to access the debate [DeRose and Warfield 1999]. We will 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. Duchess the Labrador may know its master is home, but not know it is not a brain in a vat. We maintain that 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 favoured intuition. We find that sometimes intuitions have to go (absolute space and time, Euclid's parallel postulate, flat earth, etc.). Furthermore, we cannot

resist pointing out that abominable or not some such conjunctions turn out to be true—conjunctions such as Moore’s Paradox: ‘It is raining but I don’t believe it’.

True our examples of the infant and Duchess the Labrador 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 at least, we claim that it is not. 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.

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 externalists’ accounts of knowledge and justification. Bonjour claimed that knowledge requires internally justified (epistemically 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 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.

We disagree. We maintain that for reasons similar to the ones we gave in response to Kripke's example of the sloppy scientist, 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?¹⁵

What Sally doesn't have is what Steve the sloppy scientist doesn't have, *viz.*, secondary process controls that would confirm that she knows (or possesses the capacity to know). 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. We maintain that 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 Duchess the Labrador Retriever. Duchess knows that a squirrel ran up the tree. Her belief is based solely on primary cognitive processes. She has no secondary process controls in place. That is, she has no mechanism for confirming whether her cognitive mechanisms are reliable. Yet everyone agrees that Duchess can know the squirrel ran up the tree.¹⁶

Conclusion

This concludes our brief survey of what we take to be the best of the attempts to kill off the tracking theories of knowledge. We maintain that none of them succeed. True, we do not claim to have provided an exhaustive list of attacks on the theories, and we have not even addressed all of the permutations on the attempts we have discussed. Nonetheless, we do maintain that we have taken a first, strong, sustainable step in the direction of showing that those who believe the tracking theories of knowledge have been shown to be false by the accounts we have addressed, should think again. We maintain that the tracking theories remain viable accounts of

¹⁵Of 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; 2004].

¹⁶Of course, some will claim that this is cheap and dirty animal knowledge, not fancy scientific human knowledge. That's fine with us. Let Sally have only the cheap and dirty variety of knowledge. It is knowledge enough for our purposes.

knowledge. Indeed, we cannot help believe that those who are inclined towards naturalism about knowledge and mind never really gave up the tracking theories. Somewhere in the backs of their minds they likely believed that these theories could be rescued from these sorts of attacks. We believe the rescue is now well under way.¹⁷

University of Delaware
Concordia University

Received: December 2003

References

- Adams, F. 1986. The Function of Epistemic Justification, *Canadian Journal of Philosophy* 16: 465–92.
- Adams, F. 2003. Epistemology, in *Reflections on Philosophy*, ed. Leemon McHenry and Takashi Yagisawa, New York: 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: 228–36.
- Armstrong, D. 1973. *Belief, Truth, and Knowledge*, Cambridge: Cambridge University Press.
- Bonjour, L. 1980. Externalist Theories of Empirical Knowledge, *Midwest Studies* 5: 53–73.
- DeRose, K. 1995. Solving the Skeptical Problem, *Philosophical Review* 104: 1–52.
- DeRose, K. and T. Warfield 1999. *Skepticism*, New York: Oxford University Press.
- Dretske, F. 1970. Epistemic Operators, *Journal of Philosophy* 67: 1007–23.
- 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*, ed. Ernest Sosa and Matthias Steup, Oxford: Blackwell: 13–26.
- Goldman, A. 1976. Discrimination and Perceptual Knowledge, *Journal of Philosophy* 73: 771–91.
- Goldman, A. 1979. What is Justified Belief?, in *Justification and Knowledge*, ed. George Pappas, Dordrecht: Reidel: 1–23.
- Luper, 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: 215–7.
- 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. and M. Swain 1973. Some Conclusive Reasons Against 'Conclusive Reasons', *Australasian Journal of Philosophy* 51: 72–6.
- Sosa, E., 2004. Relevant Alternatives, Contextualism Included, *Philosophical Studies* 119: 35–65.

¹⁷We would like to thank Robert Audi, Ben Chaplain, Phil Hanson, Igal Kvart, Danny Goldstick, Jay Rosenberg, Bill Seager, Ernie Sosa, and Crispin Wright for questions and comments. We would also very much like to thank John Barker and Fred Dretske for conversation and comments on earlier versions of the paper.