

Knowledge and Tracking Revisited

Fernando Broncano-Berrocal

(forthcoming in *Analysis*)

Abstract: An explanatorily powerful approach to the modal dimension of knowledge is Robert Nozick's idea that knowledge stands in a tracking relation to the world. However, pinning down a specific modal condition has proved elusive (e.g., there is an increasing awareness that the safety and sensitivity conditions are not necessary for knowledge). In this paper, I offer a diagnosis and a positive proposal. The root of the problem, I argue, is the unquestioned assumption that tracking is a matter of directly preserving conformity between what is believed and what is the case in certain possible worlds. My proposal is that what we track is whether the conditions for belief formation are appropriate in such worlds. Accordingly, we indirectly track the truth by ensuring that we only use our methods of belief formation in conditions that make it likely that conformity is preserved between what is believed and what is the case.

1. The modal dimension of knowledge: a long-standing problem in epistemology

The idea that knowledge has a *modal dimension* is widespread in epistemology. After all, it is widely acknowledged that one can fail to know a proposition solely because of what might have happened. If you form the true belief that something is a barn by looking at it, but you could easily have encountered fake barns, your belief does not amount to knowledge. Surprisingly, while many epistemologists agree that knowledge features some modal condition, pinning it down has proved very elusive. In this paper, I propose an alternative way to tackle the problem.

A fruitful approach to the modal dimension of knowledge is Robert Nozick's general idea that knowing p lies between believing p truly in the actual situation and

believing p in a way that one's doxastic state *covaries* with p in all possible situations, i.e., in all logically possible worlds (Nozick 1981). In other words, knowledge is a factive state of belief that conforms to what is the case in some but not in every possible world. When the range of covariation of a belief is the one that befits knowledge, that belief is said to *track the truth*.

Pinning down an adequate modal condition on knowledge has proved elusive precisely because pinning down the relevant range of covariation has proved elusive. In particular, it is not obvious how to delimit the relevant set of possible worlds in which agents in possession of knowledge of p are supposed to (necessarily) continue getting things right if they happen to believe p in those worlds, or to (necessarily) avoid believing p if p happens to be false in those worlds.

As it is well-known, the two main candidates in the literature for delimiting the relevant set of possibilities of truth-value covariation are the sensitivity and the safety principles:¹

Sensitivity: If S knows that p via a method of belief formation M , then in the closest possible worlds where p is false, S does not believe that p via M .

Safety: If S knows that p via a method of belief formation M , then in nearly all (if not all) close possible worlds in which S continues to believe that p via M , p continues to be true.²

¹ The idea of knowledge as tracking the truth is typically associated with sensitivity (because it is Nozick's preferred principle), but the way Nozick introduces the truth-tracking metaphor is certainly compatible with safety as well.

² These are simple formulations of (method-relative) safety and sensitivity, which suffice for the purposes of the paper. For a review and critical discussion of several ways to formulate sensitivity, see

Informally, safety says that if one knows a proposition via some method of belief formation, then one's true belief formed in this way could not easily have been false, whereas sensitivity says that a method of belief formation yields knowledge only insofar as one would not use it if the believed proposition were to be false.

If epistemologists have included safety and sensitivity in their accounts of knowledge, it is chiefly because they capture the aforementioned variability of knowledge at least with some modal facts, namely they explain why knowledge is lost when *Gettier-style error possibilities* are salient. Once again, suppose that you believe that the object in front of you is a barn by looking at it, but you could easily have encountered visually indistinguishable fake barns. Why doesn't your true belief amount to knowledge? Safety and sensitivity theorists respectively argue that it is because the modal profile of your belief fails to comply with the safety and sensitivity requirements. After all, in the closest possible worlds where the object in front of you is not a barn, you would still believe it is (because you would be looking at an indistinguishable fake barn), so your belief is insensitive. In addition, you would erroneously take fake barns to be barns in most nearby possible worlds, so your belief is unsafe, or so the story goes.

Yet, the story continues, while safety and sensitivity help understand how knowledge, and more precisely, lack thereof covaries with long-studied error

Blome-Tillmann (2017). For several ways to formulate safety, see Rabinowitz (2011). In addition, note that safety and sensitivity are not the only modal conditions for knowledge that have been proposed in the literature. For example, Nozick himself included the so-called receptivity or adherence principle in his analysis of knowledge (if S knows that p, if p were true, S would believe that p) after having experimented with several other principles that he finally discarded because they fell prey to a number of problems; see especially Nozick (1981: 682, fn. 12) for the latter. For critical discussion of several other modal principles, see Egré (2008).

possibilities (viz., with Gettier-style error possibilities), not all instances of ordinary knowledge are instances of safe or sensitive belief.

Firstly, it has been argued that *knowledge based on inductive inference* is not sensitive. Suppose that you suffer from a severe allergy to peanuts and know that each time you have accidentally ingested a peanut you have suffered a life-threatening anaphylactic shock. Suppose that you happen to accidentally swallow a peanut. By reliable inductive reasoning you come to know that you will suffer an anaphylactic shock soon. However, in the closest possible worlds where this is not true (e.g., worlds where the peanut is covered with an unnoticeable protective plastic layer that prevents allergens from being absorbed), you would still believe it is true. Therefore, you have knowledge but your belief is not sensitive.

Secondly, it has also been argued that knowledge may arise in situations where there is a *high risk of epistemic failure* involved. For example, consulting an atomic clock will surely give you knowledge of the time—atomic clocks are the most accurate devices for measuring the time. But if the precise yet fragile internal mechanism of this clock were disrupted by an isotope, if it decayed, you would not know the time. Suppose that the isotope is nearby and could decay at any moment but you luckily happen to look at the clock when it has not decayed yet. By the looks of it, you can know the time (it is the most accurate clock on Earth and is working), but your belief formed on this basis is unsafe. After all, given how easy could it have been that the

isotope had decayed and disrupted the clock before you had the chance to look at it, you would get things wrong in most close possible worlds.³

So if neither sensitivity nor safety state necessary conditions for knowledge, it is hard to see how they can account for the modal profile of known beliefs. And here is the quandary: we have the strong intuition that knowledge has a modal dimension, but the best two candidates for explaining it (safety and sensitivity) make the wrong predictions about the cases. Should we give up this pervasive intuition or find alternative theoretical ways to shed light on the issue? Most epistemologists will surely agree that the latter is preferable. In what follows, I will offer a new twist on Nozick's old idea that knowing is a matter of tracking the truth. Hopefully, this will help shed light on the issue.

2. Directly tracking the truth versus indirectly tracking the truth by tracking appropriate conditions

One often overlooked part of Nozick's theory of knowledge is his motivation for the idea that knowledge stands in a tracking relation to the world. In particular, Nozick is interested in explaining how merely true beliefs could have given rise to *knowledge in a changing world*. To address the issue, he proposes the following thought experiment:

Imagine yourself in the position of God wanting to create organisms who would have (merely) true beliefs in a diverse and changing world. The ways to

³ Several counterexamples against the necessity of safety for knowledge have been proposed in the literature. This case by Bogardus (2014) is the most compelling one. Elsewhere (Broncano-Berrocal 2014), I have offered a formulation of safety that aims to steer clear of it, albeit Bogardus and Marxen (2014) have made a number of criticisms in reply.

accomplish this are to (a) constantly intervene: start them with true beliefs and intervene to change their beliefs each time the world changes; (b) determine what the whole future will be, and create each being with a preprogrammed sequence of beliefs to fit his changing situation in a preestablished harmony; (c) create beings able to detect changes in facts who will change their beliefs accordingly.

(Nozick 1981: 283-4)

Nozick finds (c)—what he calls “the evolutionary process”—the most plausible hypothesis (it is certainly the one that best reflects our world) and goes on to speculating how the adaptive advantages of being endowed with such belief capabilities could have made knowledge possible. The following quote summarizes the essentials of Nozick’s evolutionary explanation:

The evolutionary process can give organisms true beliefs (in a changing world) only by giving them the capability to have true beliefs: so, it will give them more than (merely) true beliefs. In giving them a capability for true beliefs, it makes their beliefs (sometimes) vary somehow with the truth of what is believed; it makes their beliefs somehow sensitive to the facts. However, the evolutionary process will not bestow upon them a capability for true beliefs so powerful that in no logically possible situation would their beliefs be mistaken. Even if such a capacity could arise by random mutation (...), there would not be strong selection for it; there would be no selection against those other organisms whose lesser capacities function just as well in the actual range of situations. A being with

(some of) God's traits could arise (and be maintained) by an evolutionary process only in a very tough environment". (Nozick 1981: 285)⁴

Nozick's key idea is then that it would be a too demanding condition on knowledge to require covariation with the truth in *every* or *any* possible situation.⁵ Instead, knowledge covaries with the truth in the range of variation of situations of the environment we are adapted to. Nozick's reasoning can be reconstructed as follows: (1) insofar as we are adapted to the environment we inhabit, E, we achieve a high ratio of true to false beliefs by tracking the truth in the range of variation of situations of E; (2) thanks to such reliable belief capabilities, we manage to acquire knowledge in E; (3) covariation with the truth in situations beyond that range would not increase our ratio of true to false beliefs in E. Therefore, covariation with the truth in situations beyond that range is not necessary for acquiring knowledge (in E).

Nozick's idea that our evolutionarily-shaped belief capabilities track the truth across a range of possible situations is certainly not misguided, but after more than 30 years of philosophical explanations it is still unclear how to delimit the specific set of possibilities in which this is supposed to happen.

Perhaps the problem lies in the assumption that *known beliefs directly covary with the truth* of what is believed in certain possible worlds. An alternative approach to the modal dimension of knowledge and specifically to the idea that knowledge stands in a tracking relation to the world has it that our belief capabilities track the truth, not by

⁴ Nozick does not aim for empirical adequacy with this purported evolutionary explanation. Instead, his point is that "knowledge, that is, something of the order of magnitude of tracking is what (on theoretical grounds we should expect) evolution would produce." Nozick (1981: 286).

⁵ This is what the sceptic unrealistically demands, according to Nozick.

directly preserving conformity between what is believed and what is the case in a range of varying scenarios, but by tracking whether the conditions for belief formation are appropriate in those scenarios so that *truth can be indirectly tracked* by ensuring that the relevant belief-forming methods are only used in conditions that make it likely that conformity is preserved between what is believed and what is the case.

Nozick's evolutionary motivation for his own view can be easily revised to back up this alternative view. The key idea would be the following: in the same way as in Nozick's original story the relevant evolutionary process does not bestow upon organisms a capability for true beliefs so powerful that in no logically possible situation would their beliefs be mistaken, in the revamped story the bestowed capability is not so powerful that in no logically possible situation would their conditions for belief formation be inappropriate. After all, organisms with such powerful godlike capacities for tracking appropriate conditions across all possible worlds would exhibit no better appropriateness-tracking skills relative to a specific environment than organisms only adapted to track appropriateness in that environment.

Moreover, a further restriction applies: in order to acquire knowledge it is not required that organisms track *all* the conditions for belief formation in a specific environment, but only *some* of them. To see this, note that there is an ambiguity in the very notion of appropriate conditions for belief formation: it may either mean conditions that are appropriate for forming beliefs (independently of the truth-value of the believed proposition) or conditions that are appropriate for forming beliefs that, so to speak, end up hitting the truth. Accordingly, we can make a distinction between two

types of conditions for belief formation: conditions that merely *enable* one to form beliefs (*enabling conditions*) and conditions that *determine* whether one's (formed) beliefs end up being correct or incorrect (*determining conditions*).⁶

Crucially, knowledge does not require tracking the appropriateness of the former. By way of illustration, the presence of oxygen *enables* one's cognitive processes to produce beliefs (because oxygen keeps organisms alive), but this does not mean that one needs to track this fact in order to acquire knowledge via such cognitive processes. By contrast, factors such as the size, shape, colour or distance of an object *determine* whether one's visual beliefs about that object will be correct or incorrect. When such factors are appropriate for one's reliable visual capacities, one will likely end up forming beliefs that are true. Unsurprisingly, the appropriateness or inappropriateness of these factors (i.e., of the determining conditions) is what one needs to track in order to acquire knowledge.

A final note of clarification is in order. The notion of appropriateness of conditions for belief formation is inherently tied to the notion of *reliability*, insofar as the latter is typically defined as a tendency or disposition to produce a sufficiently high ratio of true to false beliefs *in certain appropriate or normal conditions*. Appropriate conditions for reliability usually (but not exclusively) correspond to conditions shaped by evolution (especially if the method is an innate cognitive process or ability) and to learning conditions, i.e., the conditions under which agents learn to use the method in the first place—e.g., amateur art lovers learn how to tell a painting by Picasso from a

⁶ This distinction resembles Mackie's distinction between causes and background conditions (Mackie 1974). See Broncano-Berrocal (2017; *forthcoming*) for further discussion.

painting by Kandinsky but they do not typically learn how to tell whether a painting is a forgery; the presence of forgeries thus makes their circumstances inappropriate for their ability to recognise painters.

In sum, tracking appropriateness is typically a matter of tracking the conditions that are appropriate for the reliable use of one's belief-forming methods, and this obviously neither requires tracking appropriateness in all possible worlds, nor conditions that merely enable the use of such methods.

3. Safety and sensitivity revisited

Nozick's general idea that the modal dimension of knowledge is a matter of directly tracking the truth across a set of possibilities is made specific by the safety and sensitivity principles. We are now in a position to revisit safety and sensitivity in terms of the previous explication of knowing as a matter of indirectly tracking the truth by tracking appropriate determining conditions for belief formation. Let's start with safety.

Safety is the idea that if one knows that *p*, not easily could one's true belief that *p* have been false. We can simply translate this as follows: if one knows that *p* in appropriate determining conditions, not easily could the determining conditions for one's true belief that *p* have been inappropriate. We may call this principle safety from inappropriateness or *i-safety* for short. More formally:

I-Safety: If *S* knows that *p* via a method of belief formation *M* in appropriate determining conditions, then in nearly all (if not all) close possible worlds in

which S continues to believe that p via M, the determining conditions for S's belief that p continue to be appropriate.

Sensitivity could be translated analogously by simply replacing truth with appropriateness of conditions in the formal statement of the principle. It would read as follows: if S knows that p via a method of belief formation M in appropriate determining conditions, then in the closest possible worlds where the determining conditions are inappropriate, S does not believe that p via M.

The problem with this principle is that it makes knowledge sensitive to bizarre conditions for belief formation, because the *closest* possible worlds where the determining conditions are inappropriate may depart too much from actuality. By way of illustration, suppose that you know that you will suffer an anaphylactic shock after having ingested a peanut on the basis of your knowledge that this has happened many times in the past. The closest possible worlds where the determining conditions for your true belief are inappropriate are possible worlds in which the peanut is harmless, e.g., worlds in which it is covered with an unnoticeable plastic layer. This has never happened in the past and will not likely happen in the future, but the revamped sensitivity principle would explain your inductive knowledge—i.e., knowledge of a future event based on knowledge of a repeated sequence of past events of the same type—in terms of the fact that you are able to track such a possible anomaly. Intuitively, however, it is a too demanding condition on knowledge to require sensitivity to such bizarre possibilities.

A better view has it that sensitivity to whether the determining conditions are inappropriate only ranges over nearby possible worlds. Let's call this principle *i-sensitivity*:

I-Sensitivity: If S knows that p via a method of belief formation M in appropriate determining conditions, then in nearly all (if not all) close possible worlds where the determining conditions for believing p via M are inappropriate, S does not believe that p via M.

I-safety and i-sensitivity are the two sides of the same coin: whereas i-safety amounts to assurance of appropriate (determining) conditions for belief formation in close possible worlds, i-sensitivity amounts to avoidance of inappropriate (determining) conditions for belief formation in close possible worlds.

4. The problem set

Without further ado, let's see how well i-safety and i-sensitivity fare with respect to the cases so far considered:

(1) *Gettier-style cases*. Consider the most perplexing Gettier-style case, the fake-barn case. According to i-safety, you ignore that you are in front of a barn: in the actual world the determining conditions for your true belief are appropriate (the light conditions are good, the barn is a real one, it is at a close distance, and so on); in nearby possible worlds, however, they are not, because the presence of fakes is not the sort of scenario that is typically adequate for the reliable exercise of a visual categorization ability.

Analogously, nearby possible worlds are such that you would continue to form the belief that you are in front of a barn when the determining conditions for your belief are inappropriate. So your belief is not *i*-sensitive either. Similar considerations apply to standard Gettier-style cases.

(2) *Inductive knowledge*. Beliefs that arise from inductive inference, if they amount to knowledge, are also *i*-safe. Consider the peanut case. You know that you will suffer a life-threatening anaphylactic shock after ingesting a peanut because you know that it has happened many times in the past. The actual determining conditions for belief formation are appropriate: as in past cases, you have ingested a standard peanut, you still suffer from a severe allergy to peanuts, you start to feel the first symptoms of anaphylaxis, and so on. Nearby possible worlds in which you believe the same proposition are such that the determining conditions continue to be appropriate, so your belief is *i*-safe.

In addition, beliefs from reliable inductive inference are (if they amount to knowledge) also *i*-sensitive. Although you would certainly believe that you will suffer a life-threatening anaphylactic shock even if this were not the case (e.g., in circumstances where the peanut is covered with an unnoticeable plastic layer), in no close possible world anomalies of this sort occur, so the consequent of *i*-sensitivity is vacuously true. That is, while your belief is not sensitive, it does count as *i*-sensitive.

(3) *Knowledge in situations of high risk of epistemic failure*. In the clock case, you know the time by consulting a functioning atomic clock in the presence of an undecayed isotope despite the fact that the isotope could easily have decayed, which would have disrupted

the clock's internal mechanism. Arguably, the fact that an undecayed isotope *does not* disrupt the clock is a mere *enabling condition* for the normal functioning of the clock. To see this, consider the following analogy. Someone throws a stone at you. Seconds before the stone hits your head you see a cat and form the true belief that there is a cat. That person could easily have thrown the stone a few seconds before so that the hit on your head would easily have made you hallucinate that there is a dog instead.

The point then is that in the same way as the stone hit would prevent you from using your functioning visual apparatus to form a belief about the animal in front of you, the decayed isotope would prevent you from using your method of belief formation to form a belief about the current time. After all, this method not only includes the use of your visual apparatus, but also of a functioning device, which is rendered useless. This indicates that the fact that the isotope has not decayed is a mere enabling condition for your method. We have seen, however, that we need not track enabling conditions, but only conditions that determine whether one's beliefs (once one's belief-forming method is usable) end up being correct or incorrect.

Crucially, if something renders your method of belief formation useless, i.e., if your enabling conditions are not in place, the corresponding determining conditions will not be in place either, because the former are preconditions for the latter. This explains why your true belief in clock is i-safe: the consequent of i-safety is vacuously true because in no close possible worlds in which you believe the same proposition about the time the relevant determining conditions for your belief are in place. After all, it makes no sense to talk about appropriate conditions for forming true beliefs via a

method of belief formation if you have been stripped of it. In addition, your belief is i-sensitive for the same reason: there are no close possible worlds in which you form a belief in inappropriate determining conditions because such conditions are not in place.

5. Conclusion

The modal dimension of knowledge is captured by Nozick's idea that knowledge stands in a tracking relation to the world, but this tracking relation is not direct, as Nozick and many others following him have for so long assumed: our belief capabilities do not directly preserve conformity between what is believed and what is the case in a given set of possibilities. Instead, we indirectly track the truth by tracking whether the conditions for belief formation are appropriate in those possibilities, thus ensuring that our beliefs would be formed only in conditions that make their correctness likely. Borrowing Nozick's words, this is what (on theoretical grounds we should expect) evolution would produce.

References

- Broncano-Berrocal, F. (forthcoming). Purifying impure virtue epistemology. *Philosophical Studies*, DOI: <https://doi.org/10.1007/s11098-017-0873->.
- Broncano-Berrocal, F. (2017). A Robust enough virtue epistemology. *Synthese* 194: 2147-2174.
- Broncano-Berrocal, F. (2014). Is safety in danger? *Philosophia* 42: 1-19.
- Blome-Tillmann, M. (2017). Sensitivity actually. *Philosophy and Phenomenological Research* 94: 606-625.

Bogardus, T. (2014). Knowledge Under Threat. *Philosophy and Phenomenological Research* 88: 289-313.

Bogardus, T. & Marxen, C. (2014). Yes, safety is in danger. *Philosophia* 43: 321-334.

Egré, P. (2008). Reliability, margin for error, and self-knowledge. In V. Hendricks & D. Pritchard (eds.), *New Waves in Epistemology*. Palgrave-Macmillan: 215-250.

Mackie, J. L. (1974). *The cement of the universe*. Oxford University Press.

Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

Rabinowitz, D. (2011). The safety condition for knowledge. *Internet Encyclopedia of Philosophy*.