

Tracking without Concessions?*

DANILO ŠUSTER

Department of Philosophy, Faculty of Arts, University of Maribor, Koroška c. 160, 2000 Maribor, Slovenia
danilo.suster@um.si

ORIGINAL SCIENTIFIC ARTICLE / RECEIVED: 15–04–13 ACCEPTED: 24–10–13

ABSTRACT: In the first, shorter part of the paper I point out some problems and potential misunderstandings connected with B. Berčič's treatment of Nozick's sensitivity condition for knowledge. In the second part of the paper I offer the condition of modal stability or limited sensitivity as a revision of Nozickian conditions for non-accidental connection between our belief and the truth of our belief. "When it is seriously possible for you to falsely believe that p ," that is a good reason for denying that you know that p . Sensitivity within limits requires that we consider *more* possible worlds (all within a sphere of serious or relevant possibilities) than classical sensitivity, but *not* the worlds outside this sphere. The idea of modal stability combines robustness (benefits of safety) with responsiveness to facts (benefits of sensitivity).

KEY WORDS: Modal epistemology, modal stability, Nozick, safety, sensitivity.

1.

To (mis)use a famous quote by Kant, two things fill my mind with admiration when leafing through two volumes (nine hundred pages!) of Boran Berčič's introduction to philosophy. The first is its oceanic breadth—the topics covered include the meaning of life and social contract, philosophy of mathematics and epistemology, fatalism and philosophy of mind, to mention just some of them. The author has a precious talent for presenting a wide variety of complex issues in a readable and accessible form. I have used this material in my pedagogical work and the students really like it. But, as every writer of textbooks knows, one has to steer between Scylla of accessibility and Charybdis of oversimplification. This is not always easy and in the first, critical part of my discussion I will point out some of

* Presented at the "Filozofija—Jam Session," book symposium on Berčič (2012a) and (2012b), Rijeka, 14–15 February 2013.

the problems connected with the treatment of Nozick's tracking account of knowledge and clarify some misunderstandings (Berčić 2012b: 16–18).

The second thing in the book which will impress those of us who are more cautious by nature is the number of exclamation marks scattered on the pages of book—sometimes stylistic, but very often expressing the decisiveness of the author's attitude. To take a random, but typical example: "Philosophy is seeking truth! This is correct! *Philosophy really is seeking truth!* We are not going to accept a philosophical theory in case it is not true! *In order to accept a philosophical theory, it has to be true!*" (2012b: 360). Although critical to Berčić's discussion on Nozick, I agree with his main exclamation marks in the area of knowledge, e.g. "What about the claim that the external world exists? Can we verify this claim? ... the answer is clear: we can! (2012b: 135). ... The possibility of a mistake is compatible with knowledge ... That's is the reason why Descartes was wrong!" To reconcile this spirit of agreement with some of the really problematic features of Nozick's account of knowledge, justly emphasized so by Berčić, I will propose a revised version of tracking which allows for our knowledge of the external world without concessions to skepticism, or so I hope. I shall defend the condition of modal stability or limited sensitivity as a revision of Nozickian conditions for non-accidental connection between our belief and the truth of our belief.

2.

Gettier famously objected to the standard definition of knowledge: *S* knows that *p* iff

- i. *p* is true;
- ii. *S* believes that *p*;
- iii. *S* is justified in believing that *p*.

Gettier showed that you can have a justified true belief and still lack knowledge of what you believe because your true belief was ultimately gained by accident. Berčić uses the example of a blocked thermometer. You enter a room and see that the thermometer on the wall shows 22°C. On this basis you come to believe that the temperature in the room is 22°C and the temperature really is 22°C. Although you have no reason to think that the thermometer is faulty, it is actually blocked—had the temperature been different, the thermometer would still have shown 22°C. So you do not know that the room temperature is 22°C.

Enter Nozick's "tracking" account as a remedy. The agent's belief in the case described is just too luckily true to count as knowledge because her belief does not track knowledge. Tracking is defined in terms of coun-

terfactuals: had what the agent believed been false—if the temperature had been a degree lower or higher, for example, but everything else had stayed the same—then she would have carried on believing what she does, even though it is no longer true. Audi (2011: 260) puts it nicely: “As we track a person in the snow, causally guided by the path, our belief system can be sensitive to the changing evidences that indicate the truth.” Not only can be, but it has to be, according to Nozick, in order to count as knowledge. According to Nozick *S* knows that *p* iff

- i. *p* is true;
- ii. *S* believes that *p*;
- iii. if *p* weren't true, then *S* wouldn't believe that *p*;
- iv. if *p* were true, then *S* would believe that *p*.¹

Knowledge, according to Nozick, requires a modal (counterfactual) connection between our beliefs and facts of the matter. The Gettier counter-examples do not count as cases of knowledge, because they violate condition (iii) also called the *sensitivity* principle.

Nozick's theory was faced with opposition from the very beginning and Berčić is right in the general spirit of his criticism. But not in details, or so I shall argue. According to Nozick ordinary subjects know most of the things we ordinarily take them to know. In the case of a reliable thermometer in a normal situation my beliefs track the temperature, if the temperature were different, the thermometer would indicate the change and my beliefs would change accordingly. Now imagine a skeptical scenario in which I am the victim of a Cartesian demon, deceiving me about the temperature in this room. Let 'Bad' be a skeptical scenario of this type (I am brain-in-a-vat being fed all the experiences, etc.). If I were in Bad, I would falsely believe myself not to be in such a scenario. Given Nozick's condition (iii) I cannot know that I am not a victim in such a scenario. Tracking allows for concessions to skepticism.

Many have felt that Bad is somehow irrelevant, too far-fetched to be taken seriously. According to Berčić the main problem seems to be that the third condition is too strong. I agree that this is so for *some* beliefs. My true belief that I am not in Bad is insensitive, so I can never know that I am not in Bad, which seems *implausible*. But Berčić has a *different* complaint (2012b: 17): “It seems that even ordinary cases of knowledge do not satisfy this condition.” How could that be the case?

Whether a true belief counts a case of knowledge depends on whether this belief would still be true in a different possible situation. We said

¹ Slightly more complicated since the theory must take explicit account of the methods of arriving at belief, but we can omit this complications.

earlier that in normal situation my beliefs track the temperature. If the temperature were different, the thermometer (not being tampered with) would indicate the difference and my beliefs would change accordingly. But we have to admit the possibility of bizarre possible situations, to quote Berčić in full (2012b: 17):

There are possible states of affairs such that A would still believe *p*, even if *p* were false and this directly contradicts the third condition. The problem is that *for each of our beliefs there are possible states of affairs such that we would hold the belief in question even if the belief had been false and this means that none of our beliefs can represent knowledge.*

But of course, to quote another *introductory* text, “What is interesting about the sensitivity principle is that while most of our everyday beliefs are sensitive to the truth, our anti-sceptical beliefs, such as our belief that we are not brains in vats, are not sensitive” (Pritchard 2010: 141). How could the existence of bizarre and remote possible states of affairs such that A would still believe *p*, even if *p* were false, directly *contradict* the third condition? The third condition is a *counterfactual conditional* of the form “If *p* were the case, then *q* would be the case,” and if the existence of a bizarre possibility where *p* is true but *q* is not directly contradicts the initial conditional, then (almost) no counterfactual can ever be true!

Consider a well made match in the right conditions (the match is dry enough, oxygen enough is present, etc.). We would like to say that “If that match had been scratched, it would have lighted” is true (or at least highly assertible, in case you have scruples about truth conditions for counterfactuals). But not *all* well made matches light, one in a million (in a billion, in a trillion, etc., if you want), say, does not (for whatever reason). Does the remote possibility of having the malfunctioning match in one’s hand refute the initial conditional? Nearly every counterfactual conditional has to face more or less bizarre refuting possibilities (you scratch the match just when the evil demon soaks the match-box, etc.), but the existence of such remote possibilities does not make the initial conditional false (unassertible).

This fact is rendered by the standard possible world semantics for counterfactual conditionals, accepted also by Nozick (1981: 680–81, fn. 8). According to the Stalnaker-Lewis orthodoxy we test whether the conditional “If it were the case that A, then it would be the case that B” is true in a possible world *i* by considering the closest possible worlds to *i* where A is true. The conditional is then true in *i* just in case B is true at all of the closest possible worlds to *i* where A is true. We ignore far-away possible worlds, where A is also true! To take the Nozickian conditional, when S knows that *p*, then, had that proposition been false, S would not

have believed it. We ignore bizarre (brain-in-a-vat, etc.) far-away worlds, where p is false. The mere possibility of the thermometer being blocked does not endanger our ordinary knowledge of the room temperature based on a normal thermometer.

Berčić actually acknowledges this fact in a footnote (2012b: 18, fn 31):

Nozick's strategy might be successful if we were able to select a relevant set of possible worlds (such that A believes that p if p is true and A does not believe that p , if p is false) not too far from the actual world (worlds in which we are not deceived by an evil demon, we are not hypnotized, etc.). In other words the definition might be appropriate in the sense that the definition is satisfied in the actual world if and only if it is satisfied in *every possible world which is similar enough*, but not if it is satisfied in *every possible world*. Of course, the problem is how to select the worlds which are similar enough out of all of the possible worlds.

This is somehow confusing—it is not that the definition, i.e. Nozickian conditions (i)—(iv) have to be satisfied in similar worlds, rather, when evaluating conditions (iii) and (iv) we have to take into account the closest possible worlds where p is false (true). But, more importantly, why is this explanation in a *footnote*? The whole point and the historical novelty of Nozick's approach was exactly to explain, against the skeptic, how *ordinary* knowledge is possible even when faced with the possibility of radical skeptical scenarios!

Of course, we should not forget that we are dealing with a *textbook*, not a scientific article in a prestigious journal. I assume that those passages reveal a conflict between pedagogical considerations and professional standards. The main motivation for tracking is understandable enough, insensitivity of one's belief: "Even if p had been false, you would have believed it anyway," looks like a good reason for denying that you know (cf. Williamson 2000: 147). But it is not easy to explain Nozick's account in any depth without mentioning relevant worlds, counterfactual conditionals, possible-worlds semantics, etc.

3.

According to Nozick I know the temperature of the room because in the closest possible worlds in which the temperature is different the thermometer in question registers this fact and I accommodate my belief accordingly. However, the tracking theory implies that I do not know that I am not a brain in a vat being fed misleading experiences to make it appear as if the temperature has changed. If I were brain in a vat in the contemplated scenario, I would wrongly think that I am not, in violation of sensitivity. Although I know the temperature, I do not know the entailed proposition

that I am not a brain in a vat, deceived about the temperature in the external world. This conjunction of the knowledge-attributing claim about my ordinary knowledge and the knowledge-denying one about the skeptical scenario is, in DeRose's term (1995), *abominable*.

Nozick accepted the conjunction as offering an explanation for the attractiveness of skepticism without capitulating to it entirely. Many would say that this is a serious concession to skepticism and I agree. Can we do without it while still paying due respect to the main motivation for tracking?

I followed Williamson in introducing the motivation for tracking in *negative* terms (you do not know that p if your belief is insensitive). But what explains one's failure to know in a typical skeptical scenario? According to *positive* proposals sensitivity is closely associated with *discrimination* and *reliability*: "A person knows that p ... only if the actual state of affairs in which p is true is distinguishable or discriminable by him from a relevant possible state of affairs in which p is false ..." (Goldman 1976: 774). The capacity to discriminate is the mark of reliability. Now, suppose "one can know that P only if one can tell whether P " (McGinn 1984: 543). The way typical skeptical scenarios are introduced, we are unable to tell them apart from normal environments. Should we then conclude that we do not know that we are in Bad situation? How about Russell's (1921: 159–60) classic:

There is no logical impossibility in the hypothesis that the world sprang into being five minutes ago, exactly as it then was, with a population that "remembered" a wholly unreal past. There is no logically necessary connection between events at different times; therefore nothing that is happening now or will happen in the future can disprove the hypothesis that the world began five minutes ago.

Do we know that the world did not sprang into being five minutes ago? I think we do. Dretske, one of the earliest defenders of tracking, disagrees, as he must: "If there is a past, fossils, memory, and history books tell us what happened in it—the historical details, as it were—but they cannot tell us what is implied by the existence of these details: that there actually was a past" (Dretske 2008: 21). But why should ampliative inferences like induction and abduction (inference to the best explanation) not provide a sufficient basis for preferring certain beliefs over alternative hypotheses? True, we cannot tell apart the world that sprang into being five minutes ago from our "normal" world. But then, something must be wrong with the requirement of *absolute* discrimination. Should we drop the requirement of discrimination entirely? Or just weaken it? I am inclined to accept the second option and *restrict* tracking to the space of relevant possibilities.

Instead of taking Nozick by letter (you fail to know that p if in the closest possible non- p worlds you falsely believe that p) I suggest to remain faithful to the spirit of tracking. “When it is seriously possible for you to falsely believe that p ” that looks like a good reason for denying that you know (cf. also Goldman 1976). As a refutation of knowledge this is in one sense weaker than “Even if p had been false, you would have believed it anyway.” But as Dretske has shown long ago (Dretske 1971: 9–10), “You *might* have believed that p , even if p had been false” is enough to deny your knowledge that p .² In a different sense this refutation poses a stronger requirement: you have to consider *all* of the relevant possibilities, not just the *closest* where p is false. If the connection between your true belief and relevant matters of fact breaks down in one of the seriously possible worlds, then it is not *modally stable*. And knowledge must be underwritten by a stable modal connection between belief and what makes the belief true.

Let me introduce a broader context for this notion. The old Plato’s insight that knowledge is at least non-accidentally true belief is commonly understood as the idea that knowledge requires a modally stable connection between the truth of our belief and the basis for our belief. The connection ensures that the belief is not true merely because of luck. This view is one of the cornerstones of contemporary reliabilism and externalism—you can know something noninferentially, without reasoning from prior knowledge, so long as it is no accident or coincidence that you are right (cf. Sosa 1997: 419). One way to express the idea that S’s belief that p is only accidentally true is to say that S’s belief is not responsive to the truth, it is not sensitive, it does not track the truth of p . A slightly different way is to say that S might easily have been wrong, her belief is not *safe*. Sensitivity is associated with the responsiveness requirement, what an agent believes should be dependent on the truth of what she believes. *Safely* true beliefs, on the other hand, are “error resistant” or “robust”—when one knows one is in no danger (in the relevant sense) of error. Modally inspired approaches to knowledge will ideally converge—if S’s belief that p is the output of a reliable cognitive process then S’s belief is responsive to facts and S’s belief, based on such a process, will not easily be false. And we would normally want to say that when knowledge fails one would easily have a belief that p without being right *because* in that case even if not p , one might still believe that p .

² More accurately: “You believe truly that p on the basis of reasons R , but you might still have the same reasons R even if p were false.” The idea of tracking is better spelled out not in terms of a belief tracking the fact, but one’s reasons or grounds for belief tracking the relevant fact. I follow Berčić in simplifying the discussion and concentrating on Nozickian tracking.

A belief counts as knowledge only if there is a stable modal connection between the belief and the facts of the matter. The connection is not logical—it is possible to believe falsely, but the connection has to be stable, it must not be *seriously* possible to believe falsely. Take, for example, a connection between water and its freezing point: water freezes at zero degrees Celsius. The freezing point depends on pressure, but the connection is stable, the freezing temperature of water would change by less than a degree if you increase the pressure by a factor of 100. Is it possible for water to freeze at the different temperature? Well, in one sense it is, change the environment radically (increase the pressure enormously) and the connection is gone. In another sense, it is not: take this glass of tap water, you can safely put it in the ice cube tray in your fridge, it is not seriously possible that it will not freeze. You can ignore the possibility of a radically different environment. In terms of possible worlds: far-away worlds are not relevant, they do not endanger the stability of the connection. When we test for modal stability of a connection between water and its freezing point we evaluate the conditional “If the temperature were above zero, then the water would not freeze” by considering near-by possible worlds. Now take your belief that you are not in a skeptical scenario. Is it seriously possible for you to believe falsely? Well, in one sense it is, change the environment radically (e.g. you are brain-in-a-vat) and the connection is gone. In another sense, it is not, the mere possibility of being in a radically different environment should not destroy the connection between your belief that you are not in Bad situation and your being in Good situation.

A certain “cosmology” of possible worlds centered around the actual world emerges. We can take our cue from Vogel’s (Vogel 2000: 604) description of “neighborhood reliabilism”: a belief is reliable just in case it turns out to be true whenever it is held in a neighborhood N of worlds not too far away from the actual world. S ’s true belief is sensitive if and only if in the closest possible worlds in which p is false, S does not believe that p . S ’s true belief is safe just in case it turns out to be true whenever it is held in N . For most of the ordinary contingent truths the closest worlds where they are false are “safely” within N . Still, Nozickian sensitivity requires that we always consider at least one non- p world—we have to check for potential *changes* in our beliefs when facts of the matter *change* from p to non- p . Safety poses no such requirement (you can ignore non- p worlds outside N). For “atypical” propositions non- p worlds can extend well beyond the limits of N and skeptical scenarios are as atypical as they can get—worlds where we are brains-in-a-vat are definitely out of N . According to Nozickian sensitivity we do not know we are not in Bad situation (if we were, we would still believe otherwise). Yet this belief is nevertheless

modally stable since changing the fact that we are not in Bad situation would put us beyond the limits N .

Let us say that a belief is modally stable if and only if it is sensitive in all of the worlds within neighborhood N of worlds not too far away from the actual world. When not- p is within N , safety requires sensitivity: in worlds in which it is not the case that p , one must not believe that p for otherwise it would be easy (remember, worlds in N are nearby!) to be mistaken about p (cf. Yamada 2011: 78). But the opposite is also true—under the assumption that p is truly believed, if p is sensitive (within N) then it is also safe (cf. Comesana 2007: 791). Some would disagree (e.g. Alspector-Kelly 2007: 130, Zalabardo 2012: 115) but I would be prepared to defend an even stronger claim—within N safety and sensitivity yield the same condition of modal stability.³

Nearly all of the putative cases of sensitive but unsafe beliefs or cases of tracking without knowledge in relevant literature postulate that the closest (one and only) non- p world is a world where one does not believe that p . And that is supposed to be consistent with one believing that p in more remote but still nearby not- p worlds which falsify the relevant claims of knowledge that p . Take a typical example—Goldman (1983: 84) offers the following counter-example to sensitivity being sufficient for knowledge:

Sam correctly believes that Judy is before him, but if it were Judy's twin sister instead, he would mistake her for Judy. Then, as long as the twin sister's being there is a serious possibility, Sam doesn't know that Judy is before him. But Nozick's analysis doesn't make reference to "serious possibilities"; it talks about what would be the case if p weren't true. Suppose that what would be the case if Judy weren't before Sam is that nobody would be there, and if nobody were there, Sam wouldn't believe that Judy is there. Then Sam's bid for knowledge survives Nozick's condition, and nothing else in the analysis is able to defeat it.

According to Goldman Sam's belief that Judy is before him is sensitive, but Sam still does not know that Judy is before him. A friend of safety might say that the presence of Judy's twin sister should be taken as a realistic possibility within N which defeats Sam's knowledge. But so should a friend of sensitivity! Goldman takes "If p weren't true ...," in the sensitivity conditional to be about *the* closest non- p world. Knowledge is defeated if there are any serious or relevant counterpossibilities in which the belief would be held, even if these are not the *closest* counterpossibilities to the actual situation. I agree, remember, a serious possibility for believ-

³ When safety and sensibility are expressed in terms subjunctives it is often said that such conditionals do not contrapose which is given as an argument for their difference. I have no space to discuss this argument, let me just say that they do contrapose within N .

ing falsely is a good reason for denying that you know. When we test for modal stability of a certain connection we evaluate the conditional “If it were the case that not- p , then it would not be the case that S believes that p ” by considering *all* possible worlds in N where p is false (and not just the closest possible world to the actual world where p is false).

Another question becomes pertinent immediately—how to determine the boundaries of N ? One can find notions like nearby, in the neighborhood, (very) close, similar, not very dissimilar, relevant, seriously possible, the least remote ..., is this lack of precision not a cause for despair? Probabilistic likelihood is of no help either. According to Pritchard low probability events, like lottery wins, occur in nearby possible worlds, for all that needs to be different in order for one’s ticket to be a winning ticket is that a few numbered balls fall in a slightly different configuration (e.g. Pritchard 2009: 35). So what counts as close is not determined in terms of probabilistic likelihood. Take Vogel’s example of an atypical proposition: “The Earth is not governed by cows.” In order to see whether this true belief is sensitive we have to consider what S would believe in the closest possible world in which the Earth was governed by cows. Such a world would be very different from the actual world, it might well lie beyond the boundary of N , so it remains open that S is wrong about this matter in the closest world where cows do govern, says Vogel. Thus S’s belief is not sensitive. Or is it? Based on *Planet of the Apes* we can say that in the closest ape governed worlds humans do believe they are governed by the apes. We might extend this to, say, dolphins. In the closest dolphin governed worlds humans do believe they are governed by dolphins. And then, why not cows? In the closest possible world in which the Earth was governed by cows S would believe that the Earth was governed by cows, after all. The fit will eventually end somewhere, perhaps with sea stars or worms. Or perhaps not—enough fleshing out of examples can tweak intuitions in any direction.

There is no reasonably precise characterization of the boundary between worlds nearby and remote in the relevant literature (cf. Alspecter-Kelly 2011: 129). Sensitivity predicts that the limits of N are undetermined and doubly so. Our capacities of discrimination are limited and it is not clear how reliable we are in hostile environments. There will always be a gray area, where our powers of discrimination may fail us and this will affect our bid for knowledge in some cases but not in the case of skeptical hypotheses.

Let me introduce an analogy. Suppose that you are snorkeling somewhere along the Adriatic coast looking for sea stars. The water is clean, you have to dive only for a meter or so in order to touch the rocky seabed which is clear and easy to inspect. Disappointingly you form a true belief:

“There are no sea stars on this seabed.” You are not aware of the existence of very small sea stars, at only five mm across (species *Patiriella parvivipara*), in Australia. Had one of those been present, you would not have been able to detect it. But the closest world with *Patiriella parvivipara* on the seabed is the one where you snorkel somewhere along the Australian coast. Even if you are unaware of *Patiriella parvivipara*, given your environment, you know there are no sea stars on this seabed. Had there been one it would have been an ordinary red starfish (*Asterias rubens*), and you would have seen it. At least that’s what an externalist ought to say.

Now suppose that you form an additional true belief: “There are no sea stars with telepathic powers on this seabed.” This belief follows logically from your first belief, but let me bypass the problem of closure in this paper. Suppose that if there were a red sea star with telepathic powers on the seabed, it would want to hide itself from your gaze, so it would use its powers to affect your perception. You would not be able to detect its presence, you would still have believed there were no sea stars on the seabed. Does this possibility affect your knowledge? I do not think so. Given our best scientific knowledge there is the ampliative relation of epistemic support for our observational beliefs like: there are no starfish with telepathic powers on this seabed. Existence of such creatures is *implausible* and the same is true of the skeptical hypotheses.

The two ways you can know “There are no sea stars on this seabed,” (despite the existence of *Patiriella parvivipara*) and “There are no sea stars with telepathic powers on this seabed,” represent two strategies of dispelling with skeptical worries. The first one is used by externalists: externalistically considered, there would be great differences between the vat world and the actual world. According to externalism with regard to mental content in order to have beliefs, it is necessary to be related to the environment in the right way. Given radically different environment in skeptical scenarios we would likely have access to few if any of the contents we actually have. So it is true, after all, that if I were a brain in a vat I would not believe I was not one. I would lack the conceptual resources required in order to believe anything about brains or vats (Sainsbury 1997: 918–919). Black, on the other hand, employs the fact that Nozickian sensitivity must take explicit account of the methods of arriving at belief (Black 2002). The only worlds that are relevant to whether or not I know that *p* are those in which my belief is produced by the method that *actually* produces it. Skeptical scenarios are supposed to be realized in possible worlds in which my belief is produced by different methods, so they are irrelevant.

I prefer to restrict sensitivity to worlds within *N* (contra Sainsbury). I agree with Black that brain-in-vat worlds are irrelevant but I prefer

to avoid discussion about methods.⁴ So I opt for the second strategy—(im)plausibility and explanatory deficiencies. We know that propositions like “At least one spaceship exists that can be used for traveling to another galaxy and coming back within a couple of months” are false. So why don’t we know that claims like “Brains-in-vat exists,” “The technology needed for envatting people exists,” etc. are also false? (cf. Steup 2012). Explanatory lacuna are the essence of implausibility and if plausibility is an acceptable criterion of relevance, skepticism loses (Leplin 2009: 141). Bad situations are definitely outside *N*.

Nozick readily accepts the fact that skeptical scenarios are remote, irrelevant and implausible. When evaluating claims about modal stability possible worlds are ordered relative to their degree of closeness to the actual world. Some take this order to be objective, others argue that the order is merely settled within, and thus relative to, a particular context. But modal epistemology in general and Nozick in particular is committed to a certain metrics of possibilities, for how could we otherwise gain our ordinary knowledge? I know the temperature of the room because in the closest possible world in which the temperature is different, the thermometer registers this fact and I accommodate my belief accordingly. But if I am in Bad situation or Bad is a serious possibility then sensitivity fails. That is, the approach assumes that one is not, say (actually) brain in a vat, and that is not an assumption that the skeptic is willing to concede. The idea of modal stability adopts this metrics and adds the idea that sensitivity is limited, it only makes sense up to a (not precisely determined) point.

How about the problem of begging the question against the skeptic? Provided the world is in fact pretty much as we take it to be, skepticism with respect to ordinary knowledge (“Temperature in this room is 22°C”) is false. But the issue what is the actual world like is unresolved—we just *assume* that skeptical scenarios are not seriously possible. The common response on Nozick’s behalf is to point out that the sceptical conclusion is that knowledge is *impossible*. The sensitivity-based antisceptical response is clearly the denial of this claim— it presents a theoretically adequate account of how knowledge is *possible*. Knowledge in question is actually possessed only if certain further conditions obtain, but if the sceptical argument were correct then even this conditional knowledge would not be possible (cf. Pritchard 2005, 50). Rhetorically, pragmatically or (perhaps) even dialectically, in argumentation, we are not allowed to use the metrics assumed (we take for granted claims which the skeptic would not accept). But neither tracking nor modal stability is meant to *convince* or *persuade* a skeptical person.

⁴ For some of the worries about methods cf. Zalabardo (2012: ch.3).

4.

“When it is seriously possible for you to falsely believe that p ,” that is a good reason for denying that you know that p . I tried to unpack this intuition in terms of modal stability, offered as a revision of Nozickian conditions for non-accidental connection between our belief and the truth of the belief. Sensitivity within limits requires that we consider *more* possible worlds (*all* within a sphere of serious possibilities) but *not* the worlds outside this sphere. This modal condition is not entirely new, it was briefly welcomed but never really endorsed by Pritchard (he calls it super-safety, e.g. Pritchard 2002: 297 and 2012: 257) and criticized recently as inferior to safety by Greco (2012: 200), who coined the name restricted sensitivity. But I think that modal stability deserves more attention, it combines robustness (benefits of safety) with responsiveness to facts (benefits of sensitivity). True, in order to know you have to be able to tell apart and discriminate, but up to a point. Freud introduced a controversial notion of a drive to death. Perhaps we could in a similar vein speak about a drive to discriminate which explains the attraction of skepticism. As reflective creatures we want to be able to discriminate between possibilities even in circumstances where there is nothing to discriminate. We should resist this drive! Sometimes we can know without being able to tell apart.

Finally, let me illustrate the position proposed with a story. Some time ago Richmond Thomason introduced Sally, the smart wife, for epistemological considerations (“if Sally, my wife, were unfaithful to me, I would (still) believe that she was faithful for her deception would be so clever”). The case reappears in a modernized epistemological remake.

Sally is invited to dinner to Jenny’s house. The two friends are comfortably sited in a living room, each with a glass of aperitif in her hand. They hear rumbling of plates from the kitchen.

Sally: You are such a happy woman, Jenny! Frank is a wonderful person and such a good cook! I just heard him cutting and stirring—must be that delicious golden fruit salad of his.

Jenny (jokingly): And what if it is not really him in the kitchen?

Sally: Oh, I recognize those sounds when I hear them! I would not hear those sounds unless Frank were really in the kitchen.

Jenny: And if he were not in the kitchen, would you still believe it?

Sally: In these circumstances, of course not.

Jenny (having a second glass): But what if those sounds were pre-recorded just in order to fool you?

Sally: Then I would, of course, be under the false impression that he is in the kitchen. But what a fanciful idea! You two were never known for practical jokes. If that were so, anything you like would be true. No, if he were not in the kitchen, one would not hear those appetizing sounds. I know he is in the kitchen.

After the dinner Sally and Jenny are comfortably sited in the living room again, each with a cup of coffee in her hand. They hear the sounds of rumbling and water running in the kitchen.

Sally: You really are a happy woman, Jenny! Frank is a wonderful person, so diligent and so competent. I am still amazed by his discussion on how to grow *Camellia japonica*! Honestly, I could not follow him. And so faithful, too, so I heard from Susan, his new secretary.

Jenny: What—have there been any rumors??⁵ Come to think of it—if Frank were unfaithful to me, I would still believe in his faithfulness, for his deception would be so clever.

Sally: Oh, why bother with fancies like that! I told you, everybody admires his dedication. “If Frank were unfaithful to you . . .,” is just like “if pigs had wings . . .,”—anything you like would be true. One might then as well say that if Frank were unfaithful to you, you would not believe that he was.

Jenny: But I do bother! Him being so clever and me being so unsuspecting, . . .

Sally: What, would you like him to be *less* clever? Would you like to be *obsessed* by jealousy?

Jenny: No, but I would like to be *certain*.

Sally: Oh, you know he is faithful. The *way* he is, you would not believe that Frank is faithful unless he really were.

Jenny: Oh, I know *that*. But do I really know he *is* the way I think he is? For if he were not, I would still believe he is!

Sally: Forget the whole damn thing! A little credulity here and a little ignorance there makes for success in marriage. Do not spoil that!

⁵ Cf. Hawthorne (2005: 40): “Someone might know that his partner will never leave him and this may entail that she will never leave him for Mr X. But if he runs though the relevant inference, that may induce anxiety about both the inferred proposition and the original (we all know how to fill in the details!)”

References

- Alspector-Kelly, M. 2011. "Why safety doesn't save closure", *Synthese* 183, 127–142.
- Audi, R. 2011. *Epistemology. A Contemporary Introduction to the Theory of Knowledge* (London: Routledge).
- Berčić, B. 2012a. *Filozofija: svezak prvi* (Zagreb: Ibis grafika).
- . 2012b. *Filozofija: svezak drugi* (Zagreb: Ibis grafika).
- Black, T. 2002. "A Moorean Response to Brain-in-a-Vat Scepticism", *Australasian Journal of Philosophy* 80, 148–163.
- Comesana, J. 2007. "Knowledge and Subjunctive Conditionals", *Philosophy Compass* 2/6: 781–791.
- DeRose, K. 1995. "Solving the Skeptical Problem", *The Philosophical Review* 104, 1–52.
- Dretske, F. 1971. "Conclusive Reasons", *Australasian Journal of Philosophy* 49, 1–22.
- . 2008. "The Case against Closure", in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Oxford: Wiley-Blackwell), 13–26.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge", *Journal of Philosophy* 73, 771–791.
- . 1983. "Review of Philosophical Explanations", *The Philosophical Review* 92, 81–88.
- Greco, J. 2012. "Better safe than sensitive," in K. Becker and T. Black (eds.), *Sensitivity Principle in Epistemology* (Cambridge: Cambridge University Press), 193–206.
- Hawthorne, J. 2008. "The Case for Closure", in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Oxford: Wiley-Blackwell), 26–43.
- Leplin, J. 2009. *A Theory of Epistemic Justification* (New York: Springer).
- McGinn, C. 1984. "The Concept of Knowledge", *Midwest Studies in Philosophy* 9, 529–554.
- Nozick, R. 1981. *Philosophical Explanations* (Cambridge: Cambridge University Press).
- Pritchard, D. 2002. "Resurrecting the Moorean Response to the Sceptic", *International Journal of Philosophical Studies* 10, 283–307.
- . 2005. *Epistemic Luck* (Oxford: Oxford University Press).
- . 2009. "Safety-based epistemology: whither now?", *Journal of Philosophical Research* 34, 33–45.
- . 2010. *What is This Thing Called Knowledge?* (London: Routledge).

———. 2012. “Anti-luck Virtue Epistemology”, *Journal of Philosophy* 109, 247–279.

Russell, B. 1921. *The Analysis of Mind* (London: G. Allen & Unwin).

Sainsbury, R. 1997. “Easy Possibilities”, *Philosophy and Phenomenological Research* 57, 907–919.

Sosa, E. 1997. “Reflective knowledge in the best circles”, *Journal of Philosophy* 94, 410–430.

Steup, M. 2012. “Epistemology”, *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2012/entries/epistemology/>>.

Vogel, J. 2000. “Reliabilism Leveled”, *Journal of Philosophy* 97, 602–623.

Williamson, T. 2000. *Knowledge and its Limits* (Oxford: Oxford University Press).

Yamada, M. 2011. “Getting It Right By Accident”, *Philosophy and Phenomenological Research* 83, 72–105.

Zalabardo, J. 2012. *Scepticism and Reliable Belief* (Oxford: Oxford University Press).