

Swarthmore College

Works

Philosophy Faculty Works

Philosophy

2012

Nozick's Defense Of Closure

Peter Baumann

Swarthmore College, pbauman1@swarthmore.edu

Follow this and additional works at: <https://works.swarthmore.edu/fac-philosophy>



Part of the [Philosophy Commons](#)

Let us know how access to these works benefits you

Recommended Citation

Peter Baumann. (2012). "Nozick's Defense Of Closure". *The Sensitivity Principle In Epistemology*. 11-27.
<https://works.swarthmore.edu/fac-philosophy/344>

This work is brought to you for free by Swarthmore College Libraries' Works. It has been accepted for inclusion in Philosophy Faculty Works by an authorized administrator of Works. For more information, please contact myworks@swarthmore.edu.

*Nozick's defense of closure**Peter Baumann*

Robert Nozick's conception of knowledge has triggered a lot of criticism over the last three decades. According to one kind of objection, Nozick's conditions of knowledge are either not necessary or not sufficient or neither necessary nor sufficient for knowledge. It is not clear how serious this kind of objection is: Nozick did not intend his proposal to constitute a final reductive definition of *knowledge* in terms of individually necessary and jointly sufficient conditions; whether Nozick's account can deal with such objections depends on the interpretation of certain aspects of the account as well as on the further development of the account. For all or almost all objections of this kind (see, e.g., Kripke 2011), there have been interesting and perhaps even promising defenses of Nozick's theory (see, e.g., Adams and Clarke 2005). Another kind (see, e.g., Kripke 2011), of objection, however, is more serious. No matter what the exact details of the theory and no matter how one further develops it, it will always be incompatible with a very plausible and important epistemic principle, namely the principle that knowledge is closed under known entailment (the "principle of closure"). Nozick himself openly stated that his account of knowledge is incompatible with closure (see Nozick 1981, 206–11¹) – and stuck with the theory, dropping closure. Most philosophers went the other way and stuck with closure, giving up on any theory like Nozick's which is incompatible with closure. If anything is "the killer objection" to Nozick's account, then it is based on his denial of closure.²

In this chapter I will argue first (section II) that at least in many cases Nozick is not forced to deny common closure principles. Second (section III), and much more importantly, Nozick does not – despite first (and second) appearances and despite his own words – deny closure. On

¹ References by page number only are to Nozick 1981.

² Alternative modal accounts of knowledge, like the safety account (see Sosa 1999; Williamson 2000), face similar problems with closure (see Kvanvig 2004, 209; Murphy 2005; Sosa 1999, 149); it is remarkable that this is generally not seen as a big problem for safety accounts.

the contrary, he is defending a more sophisticated and complex principle of closure.³ This principle does remarkably well, though it is not without problems. It is surprising how rarely Nozick's principle of closure has been discussed. He should be seen not so much as a denier of closure than as someone who's proposing an alternative, more complex principle of closure. Let us start with some remarks on Nozick's account of knowledge and on the principle of closure.

I KNOWLEDGE AND CLOSURE

Nozick holds that a true belief constitutes knowledge just in case it stands in a certain modal relation to the fact that makes it true.⁴ Knowledge, according to Nozick, is true belief which "tracks" the truth across a range of possible circumstances (or worlds); in other words, knowledge is true belief which is "sensitive" to the truth across a range of possible worlds.⁵ Here is a first list of conditions which spells out this idea (with " $\Box \rightarrow$ " for the subjunctive conditional):

Necessarily, for all subjects S and for all propositions p : S knows that p iff

- (1) p
- (2) S believes that p
- (3) Not- $p \Box \rightarrow S$ does not believe that p
- (4) $p \Box \rightarrow S$ believes that p

(See Nozick 1981, 172–78; for an earlier account along very similar lines see Dretske 1971. See also Goldman 1976 and Carrier 1971.)

Nozick adds that the subjunctive conditions need to be taken to hold with respect not to all possible worlds but only to close ones (see 173–74); he says very little, if anything, about what constitutes closeness or remoteness of possible worlds. Apart from that, Nozick argues, using his well-known grandmother case,⁶ that we have to include a reference to

³ Can it still count as a principle of closure? There is not one single agreed-upon formulation of a unique closure principle but rather different though similar formulations of the same basic idea. Nozick's principle is similar enough to these formulations to deserve the title "closure principle." Furthermore, if what makes a principle a principle of closure is that it can explain how we can acquire new knowledge by inference from old knowledge, then Nozick's principle also has to count as a closure principle.

⁴ Talk about "facts" or "truth-makers" should not be interpreted here as having any substantial metaphysical implications.

⁵ The terms "tracking" and "sensitivity" are not always used the same way in the literature; here I am using them as synonyms referring to the conjunction of Nozick's third and fourth condition.

⁶ "A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly, we must restate our conditions to take explicitly account of the ways and methods of arriving at belief" (179).

methods of belief acquisition used by the subject (see 179–85). How to individuate methods is a difficult and well-known problem (see below). There is also a bit of a debate about how exactly to include a reference to methods in condition (3), the *variation* condition, and in condition (4), the *adherence* condition (see Luper-Foy 1984, 28–29; Williamson 2000, 153–56; Becker 2009, 20; Alfano 2009, 274–75): should reference to the method used appear in both the antecedent and the consequent of the conditional or only in one of those? Not too much depends on this question here and we can include it in both, following Nozick (making adjustments and modifications when necessary and as we go along). This gives us the following explanation of the concept of knowledge, according to Nozick⁷:

- (Nozick-Knowledge) Necessarily, for all subjects S , all methods M , and for all propositions p : S knows via M that p iff
- (1) p
 - (2) S believes via M that p
 - (3) (Not- p and S uses M to settle whether p) $\square \rightarrow S$ does not believe on the basis of using M that p
 - (4) (p and S uses M to settle whether p) $\square \rightarrow S$ believes on the basis of using M that p (see 172–79).

So much for the concept of knowledge as Nozick explains it. What about closure?

Here is a first, very rough, version of a closure principle:

Necessarily, for all subjects S and for all propositions p and q : if S knows that p , and if S knows that p entails q , then S knows that q .

However, this won't do. S might simply not put "two and two together": S might know the first proposition and also know that it entails a second proposition but simply not make the relevant inference and thus not come even to believe the second proposition. Apart from that, we tend to find closure principles plausible because they explain how we can have inferential knowledge. This gives us the following version of a closure principle:

Necessarily, for all subjects S and for all propositions p and q : if S knows that p , and if S competently infers q from p , then S knows that q .

But what if S knows the first proposition and makes the relevant inference but just cannot bring himself to believe the second proposition, for

⁷ We can probably leave the problem aside here how (3) applies to necessary propositions (see 186–87). We can also leave aside complications arising from the use of a plurality of methods (see 180–85). I will call whatever meets (Nozick-Knowledge) "Nozick-knowledge."

instance, because it strikes him as wildly implausible? (See Harman 1986, 11–12.) Let us strengthen the antecedent further:

(Closure) Necessarily, for all subjects *S* and for all propositions *p* and *q*: if *S* knows that *p*, and if *S* competently infers *q* from *p*, and if *S* thereby comes to believe that *q*, then *S* knows that *q* (see, e.g., Williamson 2000, 117).

This is probably not the last word. The subject might lose knowledge of the first proposition before she comes to believe the second proposition (see, e.g., Hawthorne 2004, 29, 33). Should we therefore strengthen the antecedent further and add the condition that the subject retains her knowledge of the first proposition until she comes to believe the second proposition? Or what if the subject learns of some defeater for the second proposition while making the inference? (See Kvanvig 2006, 261–62.) Should we add still another condition that this is not so (or at least: that this is not so in the case of undefeated defeaters)? David and Warfield (2008, sections II–III) argue for further modifications, but we do not need to go into the potentially endless business of adding further conditions to the antecedent. (Closure) is sufficient here as a formulation of a commonly accepted principle of closure. What now is the relation between (Nozick-Knowledge) and (Closure)?

II NOZICKEAN VIOLATIONS OF (CLOSURE):

I am sitting on a chair right now and I believe that I am sitting on a chair right now. Were I lying on a sofa, walking on my hands, or engaged in one of the many realistic alternative ways of not sitting on a chair now, then I wouldn't believe (via the usual methods) that I'm sitting on a chair right now. Were I sitting on a chair now but something else was different (but not too wildly different), then I would still believe that I am sitting on a chair right now (again, via my usual methods for finding such things out). In other words, my belief that I am sitting on a chair right now meets Nozick's conditions for knowledge. According to Nozick, I know that I am sitting on a chair right now. Furthermore, I might competently infer from this and thus come to believe that I am not currently climbing the wall instead of sitting on a chair. This belief, too, constitutes knowledge: were I currently climbing the wall, I wouldn't believe I'm not doing that, and were I abstaining from wall-climbing in different ways I would keep believing that I'm abstinent. In other words, I know that I am not currently climbing the wall. This accords with what (Closure) demands; this is a case where (Nozick-Knowledge) and (Closure) are unproblematically compatible with each other.

Things get much trickier in other cases. I believe that I am currently sitting on a chair and I competently infer from this and come to believe that I am not merely hallucinating that I am currently sitting on a chair. However, were I merely hallucinating I would typically still believe I'm not. I don't know according to (Nozick-Knowledge) but I do know according to (Closure) – given plausible assumptions – that I'm not merely hallucinating that I am currently sitting on a chair. Nozick-Knowledge is not closed. And that, according to many, is a huge problem for Nozick's account of knowledge: shouldn't closure be preserved?

Nozick himself states that his tracking account violates closure (see 206–11). He accepts closure under known equivalence (see 229, 690n. 60), under known existential generalization (see 234; but cf. Fumerton 1987, 172, and Garrett 1999 for counterexamples), under known disjunction introduction (230) and conjunction introduction (236), but he denies closure under known universal instantiation or conjunction elimination (227–29).

According to Nozick, I can, for instance, know that I am currently seated and not a brain in the vat on a shelf which falsely believes itself to be currently located on a chair; however, I cannot know that I am not a brain in the vat on a shelf thinking of itself as being located on a chair. Hence, I cannot come to know the latter by competent inference from the former. (Closure) is violated. Similar things hold for other kinds of inferences which have played an important role in recent discussions on skepticism. I truly believe that I have hands, and I can competently infer from that and thus come to believe that I am not being deceived by some Cartesian demon into falsely believing that I have hands. Nozick would grant me knowledge of the former but not knowledge of the latter; hence, closure fails in this case.

It is not so clear whether we should lament the failure of closure in such cases. The alternatives don't look much better: either suspiciously (but cf. Klein 1981) easy denials of skepticism or rampant skepticism denying that we even know that we have hands.⁸ Much more threatening to (Nozick-Knowledge) are cases where no denials of skeptical scenarios are

⁸ But can't I know that I'm not merely dreaming that I'm currently seated (given that I am currently seated) simply because the method I'm using in non-sceptical circumstances in order to figure out whether I'm seated is not applicable at all under sceptical circumstances (see Lipson 1987, 330–31, and Williams 1996, 336–46)? Brains in the vat don't use their senses because they don't have any senses. It is not at all clear how this relates to Nozick's account, given that Nozick assumes that the method can be used no matter whether *p* or *not-p* is the case. Would it still be the same account if we gave up the idea that the variation condition can always be applied (see also 214–17)? These questions also relate to the issue where in the counterfactual conditionals the method should be mentioned (see above).

involved. Take, for instance, Kripke's example (see Kripke 2011, 186–87 and, in addition, 178–79 and 192, 197, as well as Goldman's early dachshund example in Goldman 1976, 779–83 and in Goldman 1983, 84–85, as well as, building on it, Dretske's VW-Saab example in Dretske 1975, 801): *S* is in fake-barn country where only red barns are exempted from being faked; under such circumstances it seems that *S* might know she is facing a red barn, come to believe on the basis of inference that she is facing a barn but not come to know that there is a barn because the latter belief does not meet (Nozick-Knowledge-3).

What can the Nozickean do about such cases? One option is to bite the bullet and simply deny (Closure) or related principles (see, e.g., Becker 2007, ch. 6). But there are less controversial and quite interesting alternatives. One can deny that there is a failure of closure in the red barn case. As soon as one takes the relativity to methods seriously, so this idea goes, one can see that the subject knows that there is a barn. The subject came to believe that there is a barn via the "red barn method" of inferring it from the proposition that there is a red barn. Had there not been a barn but a non-red fake barn, then the subject would not have come to believe that there is a barn on the basis of the red barn method; the subject would have used a different method.⁹ Since there are also no problems with the adherence condition, we can conclude that (Nozick-Knowledge), applied in the right way, does indeed give the right result in the red barn case (see Adams and Clarke 2005, 214–16; see also Roush 2005, 102 as well as 157–59 on Goldman's dachshund case; Goldman himself favours a similar strategy for this kind of problem – see Goldman 1976, 779–80; for the application of this kind of strategy to Moorean anti-sceptical arguments see Black 2002).¹⁰

⁹ The question whether the relativization to methods should appear in the consequent or in the antecedent (or in both) of the subjunctive conditional comes up again here. If it appears in the antecedent and we are checking for the variation condition, then we are dealing with circumstances (the subject facing a non-red barn) in which the subject can only use the relevant method (the red barn method) if it is basically mistaken (sees the wrong color) or confused (chooses a clearly inappropriate method). If we can take (Nozick-Knowledge) in the sense that the use of the method does not allow for such basic mistakes or confusions, then we can conclude that at least in some cases, like the one here, the method should only be mentioned in the consequent. All the cases above also suggest that whether it should appear in the consequent or in the antecedent or in both might vary from case to case.

¹⁰ There are more examples like that. Williamson, for instance, presents the case of someone who knows that Jack is 6 ft 9 in tall and therefore also knows that Jack is taller than 6 ft though the latter belief is not sensitive (see Williamson 2000, 159–60, for this kind of example). Becker 2009, 27–30, uses the above kind of strategy in reply (see also Roush 2005, 71–72). Hawthorne 2004, 45, presents the case of someone who knows that they have eaten less than a pound of salmon and thus also know that they have eaten less than 14 pounds of salmon (even if eating 14 pounds would have produced the illusion that they have eaten less than one pound). Hilpinen

One major problem with this kind of strategy is that it has proved notoriously difficult to identify the method used by the subject in a systematic, principled, and non-arbitrary way; the so-called generality problem (see Feldman 1985; Bach 1985) probably has no solution. Defenders of Nozick might feel free to characterize the method in such a way that the subject comes out knowing, while critics of Nozick might do it in such a way that the subject comes out not knowing (see on this also Williams 2002, 149–51). Lacking an argument for one characterization over all alternatives, this strategy to save Nozick from the red barn objection seems as plausible or implausible as the objection itself. We have reached an impasse here.

Another strategy to reply to red barn cases and also to other kinds of cases exploits a further gap in Nozick's theory: the lack of a closeness metric for possible worlds. There have been attempts (see Lewis 1979) to provide one, but it is not clear whether they can be successful (see Barker 1987, 287; Grobler 2001, 293; and Baumann 2009), and Nozick himself does not offer one. So, one might try to propose a closeness metric which would save Nozick from a given counterexample. However, not only is it not clear how to use this in a plausible way to save Nozick from red barn cases and other cases: is a world in which the subject looks at a non-red fake barn really too remote and more remote than a world in which the subject looks at some non-fake barn? One might also suspect that there are at least some cases where the more plausible closeness-rankings allow for an objection against Nozick rather than for his defense. Alternatively, we do face at least the same kind of unsatisfying indeterminacy as in the case of different characterizations of the method used.¹¹

What then about the idea proposed by Roush (2005, 41–51), namely to turn the variation and adherence conditions, conditions (3) and (4), respectively, in (Nozick-Knowledge), into sufficient but not necessary conditions and to add a further sufficient but not necessary condition according to which one can come to know a proposition on the basis of a competent inference? This could still be seen as an account which is broadly Nozickian in spirit and at the same time can deal with red barn cases and other cases (see Roush 2005, 57–74, 93–113). But this wouldn't be Nozick's pure tracking account. Since we're dealing with the prospects

1988, 161–63, presents the case of someone who comes to know, by consulting a thermometer which is reliable for the range between 0°C and 100°C but not beyond, that the temperature is 37° and thus also that the temperature is not below –40°; the latter belief, however, would violate Nozick's variation condition.

¹¹ It does not help the Nozickian to point out that in the red barn case the subject can non-inferentially come to know that there is a barn. What is at issue here is whether the subject can come to know this on the basis of an inference. It seems very plausible to say that this should be possible; hence Nozickian accounts have a problem here.

of such an account in the face of (apparent or real) failures of closure, we should leave these kinds of ideas aside here.

Finally, the Nozickean could argue that at least in some cases closure holds because we do know the conclusion of our competent inference after all. Take the dreaming argument. Assume that I am seated right now and suppose that I know this. It seems that I cannot come to know on the basis of this that I am not merely dreaming that I am seated right now. The latter, many would say, is something I simply cannot know. But why not? Presumably because the variation condition is violated: were I merely dreaming that I am seated right now, then I would still believe that I am not merely dreaming right now that I am currently seated. However, some argue that we don't have beliefs in dreams (see, e.g., Sosa 2007, ch. 1). If that is right, then (Nozick-Knowledge-3) is not violated: if I were merely dreaming that I am seated right now, then I would not believe that I'm not (because I wouldn't believe anything in the first place). There are several problems with this kind of move. First, the Nozickean would have to burden his account with some substantial piece of philosophy of mind. Second, and more importantly, this won't help against other skeptical scenarios: don't brains in vats have beliefs? Or hallucinators? Third, the natural interpretation of the variation as well as the adherence condition is that the subject would have beliefs in the relevant counterfactual circumstances; we were not thinking of accepting such "vacuous" meeting of the conditions (see Barker 1987, 291–92 for a similar point in the case of knowing that one has beliefs).

Related manoeuvres are even less promising and even more problematic. Take the following idea. I am standing in front of Dack, the dachshund. I know that this animal in front of me is a dachshund. I competently infer from this that this animal in front of me is a dog. I even know this latter proposition: despite the fact that I cannot tell dogs from wolves, my belief that this animal in front of me is a dog meets the variation condition. Were I, say, facing a wolf, I might still say "This animal in front of me is a dog" but thereby express a belief different from the belief triggered by Dack. The same demonstrative sentence changes its content (the content of the belief expressed) with a change in the reference of the demonstrative "this" or "this F". So, in a different counterfactual situation like the one above I just would not be able to have the same belief. Hence my belief that this animal in front of me is a dog is sensitive after all. Similar points can be made about indexicals (see for this kind of thought also: Gendler and Hawthorne 2005, 333–34; Manley 2007, 403–6; Kripke 2011, 169–71 (especially notes 20, 23), 192 (note 56), and 213–14).

This kind of move backfires badly.¹² First, we do have a “vacuous” conformity with the variation condition here. Much more important is a second worry. Compare the demonstrative belief that this animal in front of me is a dog with the non-demonstrative belief that I am currently looking at a dog. The latter belief violates the variation condition while the former doesn't. So, according to Nozick I would not know that I am currently looking at a dog, but I would know that this animal in front of me is a dog.¹³ This sounds like an abominable conjunction: it is very implausible that one could not know the one while knowing the other. (It is also implausible that I can come to know that this animal is a dog by inference from *This animal is a dachshund* while I cannot come to know that there is a dog in front of me by inference from *There is a dachshund in front of me*.) So, the last attempt to save Nozick from closure failure does not help at all and in addition raises general doubts about how Nozick's account would deal with indexical or demonstrative beliefs.¹⁴

Where does all this leave us? Fiddling around with the characterization of the method used or with the closeness metric will at best lead to a situation where neither Nozick nor the critic has an argumentative advantage. Other strategies look rather dim for the original Nozickean account. Should we then simply accept that Nozick denies (Closure) as well as related principles?

III NOZICK-CLOSURE

No, this is not the end of the story. There are good reasons to think that (Closure) is not even close to a satisfying principle of closure. We should therefore wait with our judgment about whether (Nozick-Knowledge) violates closure until we have found a more adequate principle of closure.

¹² It deserves mentioning that it does not help to try to block it by reformulating (Nozick-Knowledge) in terms of Kaplanian belief “character” rather than Kaplanian belief “content” (see Kaplan 1989, 500–7). Character underdetermines content.

¹³ Similar problems arise for my belief that Dax is a dog if we adopt Kripkean ideas about essential properties and argue that dachshunds are essentially dogs: Dax just couldn't be a wolf. Suppose I thus know that Dax is a dog and also know that there is Dax in front of me. I would still not know, so it seems, that there is a dog in front of me. See Hughes 1996 for a series of cases along such lines as well as for the argument that I could, for Kripkean reasons, not be a brain in a vat, so that my belief that I'm not a brain in a vat would vacuously meet the variation condition and I would know that I am not a brain in the vat; see Hughes 1996, 313; see also Goldman 1987, 187.

¹⁴ Consider Russell's old clock example (see Russell 1948, 98). I look at a stopped clock which happens to indicate the current time and I come to truly believe that it is 12.15 now. According to Nozick, I would know this: had I looked at the clock five minutes later, I would not have expressed the same belief by “It is 12.15 now,” given the changed reference of “now” (and, obviously, had the clock indicated a different time, I would not have acquired a 12.15-belief).

I will first explain why (Closure) is not satisfying and propose a modification. Second, I will argue that one can find a more satisfying principle in Nozick's text (contra Nozick himself). Third, I will discuss how this principle deals with relevant problem cases.

Let us start with a case. I wonder what the temperature is and check my reliable thermometer, which indicates that it's 66 degrees Fahrenheit. This is true and I come to know, relying on the thermometer, that it's 66 degrees Fahrenheit. I also come to know that the thermometer indicates that it's 66 degrees Fahrenheit. I competently infer from this and thus come to believe that the thermometer is indicating the temperature correctly in this instance. However, it seems clear that I cannot thus come to know that the thermometer is indicating the temperature correctly in this instance. I can even repeat this exercise and conclude that the thermometer is reliable. However, it seems clear, too, that I cannot thus come to know that the thermometer is reliable. I cannot "bootstrap" myself into knowledge in this way: there is no such "easy knowledge" (see Cohen 2002 or Sosa 2009, chs. 4, 5, 9, 10). But what then? If I competently made the inference to the conclusion, then it would follow that I don't know the premises. But that seems also very implausible. Should we then deny closure?¹⁵

No, but perhaps we should deny (Closure) and replace it by a different principle? What is creating the problem here is that I can only be said to come to know the temperature if it is OK for me to rely on the thermometer. As long as I don't have a positive reason to think that there might be something wrong with the thermometer, I may rely on it and am able to get to know what the temperature is, even if I don't have any antecedent knowledge that the thermometer is reliable; if it was not OK for me to rely on the thermometer, then I could also not be said to know the temperature. (One might protest that I cannot know anything if I don't know that I have come to my belief in a reliable way; this, however, leads directly to skepticism and is not relevant here insofar as we're looking for a plausible non-skeptical way out of our problem.) What I cannot do is gain knowledge of a conclusion based on an inference from premises which I can be said to know only insofar as I can presuppose and take for granted the not-yet-known conclusion. This is the circularity that stands in the way of acquiring knowledge of the conclusion. We can thus reformulate (Closure) in the following way:

(Closure*) Necessarily, for all subjects *S* and for all propositions *p* and *q*: if *S* knows that *p*, and if *S* competently infers *q* from *p*, and if *S* thereby comes to believe that *q*, then *S* knows that *q* – but not if *q* is both

¹⁵ Our case here involves two premises instead of one; however, this deviation does not matter here.

antecedently unknown by *S* but taken for granted and presupposed by *S*'s belief and knowledge that *p*.

This formulation might not give us the final version of an acceptable closure principle, but it is good enough and close enough for our purposes here (see for ideas in this direction: Dretske 1970, 1014; Wright 1985, 432–38, Wright 2000; Davies 1998; Davies 2000; and more recently Barke 2002, 164–66; the above principle is obviously inspired by ideas from the recent debate on failure of transmission of warrant).¹⁶

One advantage of (Closure*) over (Closure) is that the former but not the latter can explain in a plausible way how we can know the premises in the above case but not the conclusion. It also has an advantage in other important problem cases. Consider the following one (see Harman 1973, 161; Vogel 1990, 15–20; Hawthorne 2004). Suppose you know that you don't own and will never own 50 million dollars. Suppose you know that this entails that you will never win 50 million dollars in a lottery. But it seems that you can never come to know that you will never win 50 million dollars in a lottery (given that you own a ticket in a fair lottery). This can be generalized to many other "lottery propositions" (see Vogel 1990, 16–17). I am talking to a shopkeeper and all the circumstances are normal. It is thus hard not to grant me knowledge that I am talking to a shopkeeper. But do I know what is entailed by that, namely, for instance, that I am not talking to a robot programmed to behave like a shopkeeper which was coincidentally sent to exactly the shop I frequent? It is not difficult at all to find lots of pairs of propositions such that one is an ordinary proposition which we take ourselves to know and the other one is a highly probable but unknown "lottery proposition" which can be inferred from it.

Again, with (Closure) we face the trilemma of either having to give up an ordinary knowledge claim, to claim to know what we probably don't know, or to give up closure. (Closure*), by contrast, offers a way out. I can be said to know that I will never own 50 million dollars only insofar as it is OK for me to rely on the antecedently unknown but presupposed proposition that my financial matters will take a normal course. I cannot without unacceptable circularity make an inference from my premise to this presupposition and thus come to know it.

(Closure*) might even help us with skeptical puzzles when (Closure) doesn't (but cf. also Brueckner 1994, Cohen 1998 or David and Warfield 2008 on the question whether skepticism has much to do with closure

¹⁶ Some might object that (Closure*) is not a closure principle but a transmission principle (like (Closure), too). I don't want to debate word choice here but just point out that any plausible principle of "closure" will be saying something about transmission.

principles). How can I, if knowledge is closed, know that I am currently seated but not know what I can infer from that, namely, for instance, that I am not a brain in a vat merely hallucinating that it is currently seated? (Closure) has a big problem here but (Closure*) doesn't. I can be said to know that I am currently seated only insofar as it is OK for me to rely on the antecedently unknown but presupposed assumption that I am not being radically deceived. Again, I cannot make an inference from my premise to this presupposition and thus come to know it.

Nozick himself also offers a more subtle and adequate principle of closure (whether or not he sees it as a principle of transmission rather than and in contrast to a principle of closure in the more narrow sense). It is astonishing how rarely this has been discussed or even mentioned (see as exceptions: a short paper by Thompson 1986/87 and brief passages in Kripke 2011, 194–203; Mazoué 1986, 211–12; Klein 1987, 272–73; and Roush 2005, 63n. 29). Here it is:

S knows via inference (from p) that q if and only if

- (1) S knows that p
- (2) q is true, and S infers q from p (thereby, we assume, being led to believe that q). (231)

The right-to-left direction of this is remarkably close to (Closure). But Nozick adds two further “inference conditions”:

- (i) If q were false, S wouldn't believe that p (or S wouldn't infer q from p) (231)

plus the condition that if q were true, S would believe that p (and would infer q from p if he were to infer either q or *not- q* from it) (233–4).¹⁷ In other words:

S knows (via inference from p) that q if and only if

- (i) S knows that p ,
- (ii) competently infers q from p ,
- (iii) thereby comes to believe that q ,
- (iv) if q were false, S wouldn't believe that p (or S wouldn't infer q from p), and
- (v) if q were true, S would believe that p (and would infer q from p if he were to infer either q or *not- q* from it).

Since closure (or transmission) principles are usually formulated merely as sufficient conditions for knowledge of some inferred proposition q , we can formulate Nozick's closure (“Nozick-Closure”) in the following way:

¹⁷ Nozick also uses the simplified conditions “*not- q* \Rightarrow *not*-(S believes that p)” and “ q \Rightarrow S believes that p ” (see 234). He adds that he is less confident about the latter than about the former (see 692n.95).

- (Nozick-Closure) Necessarily, for all subjects *S* and for all propositions *p* and *q*: if *S* knows that *p*, competently infers *q* from *p*, thereby comes to believe that *q*, and if it is also true that
- (1-1) if *q* were false, *S* wouldn't believe that *p* (or *S* wouldn't infer *q* from *p*), and
 - (1-2) if *q* were true, *S* would believe that *p* (and would infer *q* from *p* if he were to infer either *q* or *not-q* from it), then *S* knows that *q*.¹⁸

Instead of the presupposition clause in (Closure*), (Nozick-Closure) contains the condition, roughly speaking, that the belief in the premise tracks the truth of the conclusion. This is certainly not an ad hoc or arbitrary move, especially not for a tracking theorist. Nozick himself points out (see 234) that such a principle allows us to make ordinary knowledge claims, such as *I am seated now*, at the same time deny entailed negations of skeptical scenarios such as *I am not a brain in a vat merely hallucinating that it is seated now*, while holding on to closure: were I a brain in a vat merely hallucinating that I am seated I would still believe that I am seated. (Nozick-Closure-1-1) is not met. This is all the more remarkable because Nozick thought that in order to be able to claim to know ordinary propositions and to deny entailed negations of skeptical propositions he would have to give up closure; but his own, more developed principle allows for keeping closure under such circumstances.

We also get the desired results for the above-mentioned lottery problem as well as the problem of easy knowledge and bootstrapping: were the thermometer not indicating the temperature correctly or were it not reliable, I would still believe (within certain limits, of course) that the temperature is as indicated (and not know it); and were I to win 50 million dollars in the future, I would still believe now that I will never be rich (and not know it). Like (Closure*), (Nozick-Closure) has the great advantage of dealing satisfactorily with three serious problems: Harman's lottery problem, the problem of easy knowledge and bootstrapping, and the problem of closure-based skepticism. (Closure) fails in all these respects.

But how optimistic can we be about (Nozick-Closure), all things considered? Is it in the end too weak or too strong? Does its left-hand side offer necessary or sufficient conditions for inferential knowledge?

¹⁸ I am skipping – for the sake of simplicity and because nothing hinges on it here – a further condition and strengthening of the antecedent which Nozick mentions later (239–40): an anti-question-begging condition according to which it is not the case that I would not know *p* would I not know *q*. This additional condition is somewhat similar in some respects to the final condition in (Closure*) but still markedly different.

Consider the following case, which is similar to the dachshund and the red barn cases mentioned above. I know – as most of us would admit – that my friend Hans owns a dog because I know that he owns an animal that barks (and only dogs bark). If Hans were not to own a dog, he would own a wolf which does not bark (or make any other characteristic noises). I cannot tell such silent wolves from silent dogs. Let us assume that under such alternative circumstances I would still believe that Hans owns a dog (a silent one). According to Nozick, we have to conclude in this case – if we don't do tricks with methods (see above) – that I don't know that Hans owns a dog.¹⁹ However, (Nozick-Closure) tells us that I do know that Hans owns a dog: if Hans would not own a dog, I would not believe that his animal barks. Nozick-Knowledge gets in the way of Nozick-Closure. And we should side with Nozick-Closure, considerations of plausibility suggest.

Here is another case (see McKinsey 1991, 15–16). I am thinking about water. Suppose I know that I am thinking about water and suppose I also know the truth of some form of semantic externalism which entails that I can only have thoughts about water if I have been in contact with water (it doesn't matter here whether this kind of semantic externalism is really true; it only matters what would and what wouldn't follow). Can I then infer and come to know that I have been in contact with water, only on the basis of introspection and some philosophical knowledge? There is good reason to deny this or at least to be skeptical about a positive answer. However, according to (Nozick-Closure) I can come to know this. (Had I not been in contact with water, I would not be able to think thoughts about water.) (Nozick-Knowledge) agrees, perhaps vacuously, with that verdict for the same reasons. So, (Nozick-Closure) does not formulate strong enough or sufficient conditions for inferential knowledge. Perhaps one can repair this by adding a further condition? It is not obvious whether one can.

Roush (2005, 63n. 29) presents a case where the conditions on the left-hand side of (Nozick-Closure) seem too strict for necessary conditions for inferential knowledge: someone can know an (inductive) generalization from which she can infer and thus come to know that the ice cubes left in the sun will melt – even if she would have continued to believe in the generalization had the ice cubes not melted in the sun (see also Klein 1987, 272–73 for a more intricate example leading to the same general conclusion). Both (Nozick-Knowledge) and (Nozick-Closure) seem

¹⁹ Nozick does not comment on this type of case. If he had claimed that the subject knows that Hans owns a dog, all the worse for his account.

to give the wrong, negative verdict about knowledge here.²⁰ If we're only interested in sufficient conditions for inferential knowledge (as we are when we're interested in a closure principle), then this might not be such bad news for (Nozick-Closure), but it would still be a problem for (Nozick-Knowledge).

Thompson (1986/87, 262–64) presents counterexamples to (Nozick-Closure-1-2): a subject knows that he was born in San Diego and infers and comes to know that he was born in the USA; however, the fact that he could have easily been born in San Francisco and come to believe this does not threaten inferential knowledge that he was born in the USA. (Nozick-Closure) seems too strong, again, and not to indicate necessary conditions for inferential knowledge (see also Thompson 1986/87, 264, for an attempt to repair this problem). However, one could reply that closure principles are typically formulated as sufficient (and not necessary) conditions for inferential knowledge. Overall, (Nozick-Knowledge) gives the right verdict in this case.

Nozick accepts closure under known equivalence (see 229) but rejects closure under known conjunction elimination (see 228). According to Nozick, he can know that he is in Emerson Hall, which entails that he is in Emerson Hall and not a brain in a vat far away from Emerson Hall; relying on the equivalence of the latter with the former, Nozick can infer and come to know that he is in Emerson Hall and not envatted far away from Emerson Hall. However, he cannot come to know by conjunction elimination that he is not envatted far away from Emerson Hall. (Nozick-Closure) supports this verdict, as does (Nozick-Knowledge). Hawthorne (2005, 31–32) holds that this is unacceptable, while Adams, Barker, and Figurelli (2011) argue that one should stick with closure under conjunction elimination but give up closure under known equivalence. What should one do here? I think that (Closure*) rather than (Nozick-Closure) makes it plausible that, and explains why, Nozick could be right at least in his verdict about the skeptical scenario: Nozick can be said to know that he is in Emerson Hall, given that it is OK for him to rely on certain (anti-skeptical) presuppositions, but he cannot come to know these presuppositions by inference from the relevant premises.

Finally existential generalization. Nozick thinks he can hold on to closure under known existential generalization (see 234), but Fumerton (1987) and Garrett (1999) present counterexamples. Let us take Fumerton's

²⁰ If one were to argue in favor of the negative verdict, then it is very hard to see how one can hold on to the possibility of inductive knowledge in general.

example (see Fumerton 1987, 172). Suppose that Richard knows (from reading the newspaper) that Jones murdered Smith. Had Jones not murdered Smith, then one of the nine colleagues of Jones would have done it and Richard would have found out (from the newspaper) that Jones did not murder Smith. However, had none of the ten would-be killers managed to kill Smith, then the newspaper would have reported the false rumor that Jones murdered Smith. Under these circumstances, Richard would know that Jones murdered Smith but not that someone murdered Smith. Given Fumerton's version of the example, this would be due to (Nozick-Knowledge) and would be compatible with (Nozick-Closure) (because its condition 1-1 would be violated in the case of the existential belief). However, one can easily change the story (by assuming, say, that the rumor would say that someone but not Jones had murdered Smith) in such a way that only (Nozick-Knowledge) but not (Nozick-Closure) forces us to say that Jones doesn't know that someone murdered Smith. However, if one were to try to defend Nozick here, one should rather stress the intricate nature of the example and argue that it is not so clear whether Richard can come to know by reading the newspaper whether Jones murdered Smith. Could he really, given that this newspaper would spread false rumors in certain cases?²¹

IV CONCLUSION

Where does all this leave us? Nozick does not deny closure,²² despite his own announcements to the contrary. Instead, he proposes a rather sophisticated and useful principle of closure, (Nozick-Closure), which can deal surprisingly well with certain skeptical puzzles (again, against some of Nozick's own statements) as well as with some other problems, like the problem of easy knowledge and bootstrapping or Harman's lottery problem. More standard principles like (Closure) have serious problems here (but not (Closure*)), which is not that different from Nozick's principle.

²¹ Nozick's scepticism about closure under known universal instantiation (see 227–28) remains as a problem. My true belief that all members of the club have paid their dues might track the truth even though my true belief that member Wilson has paid his dues might not track the truth because I would never believe in his tardiness. It still seems very plausible that I can come to know, via inference from the general proposition, that Wilson has paid his dues. Both (Nozick-Knowledge) and (Nozick-Closure) seem to give the wrong verdict in such cases.

²² I am ignoring problems with multi-premise closure here. It seems that everyone has a problem with that and good reason to doubt that there are valid principles of multi-premise closure (see Hawthorne 2004).

But there are also problems indicating that Nozick's view of knowledge is in tension with his closure principle; the problem here is that we should often go with the latter against the former. And then there are problems like the (at least *prima facie*) failure (in Nozick's theory) of closure under known existential generalization.²³

However, two points should be stressed, again, at the end. First, whatever one thinks about the above problems of Nozick's theory, they're – as (Nozick-Closure) helps to explain – not knock-down problems. (There might be other such problems elsewhere in his theory.) Second, Nozick's theory does not entail the denial of closure; on the contrary, he proposes a very interesting alternative closure principle which has been widely neglected, even by Nozick himself.²⁴

²³ I think (Closure*) does very well in all those cases.

²⁴ I am grateful for comments by Fred Adams, Kelly Becker, Audre Brokes, and Byeong D. Lee.