

KNOWLEDGE ESSENTIALLY BASED UPON FALSE BELIEF

Avram HILLER

ABSTRACT: Richard Feldman and William Lycan have defended a view according to which a necessary condition for a doxastic agent to have knowledge is that the agent's belief is not essentially based on any false assumptions. I call this the *no-essential-false-assumption* account, or *NEFA*. Peter Klein considers examples of what he calls "useful false beliefs" and alters his own account of knowledge in a way which can be seen as a refinement of NEFA. This paper shows that NEFA, even given Klein's refinement, is subject to counterexample: a doxastic agent may possess knowledge despite having an essential false assumption. Advocates of NEFA could simply reject the intuition that the example is a case of knowledge. However, if the example is interpreted as not being a case of knowledge, then it can be used as a potential counterexample against both safety and sensitivity views of knowledge. I also provide a further case which, I claim, is problematic for all of the accounts just mentioned. I then propose, briefly, an alternative account of knowledge which handles all these cases appropriately.

KEYWORDS: knowledge, Edmund Gettier, safety, Peter Klein,
William Lycan, Richard Feldman

I. Introduction: NEFA

Richard Feldman¹ and William Lycan² have defended a view according to which a necessary condition for a doxastic agent to have knowledge is that the agent's belief is not essentially based on any false assumptions.³ I shall call this the *no-essential-false-assumption* account, or *NEFA*. Peter Klein⁴ considers examples of what he calls "useful false beliefs" and alters his own account of knowledge in a way which can be seen as a refinement of NEFA. This paper shows that NEFA, even given Klein's refinement, is subject to counterexample: a doxastic agent may possess knowledge despite having an essential false assumption. Advocates of NEFA could simply reject the intuition that the example is a case of knowledge. However, if the example is interpreted, contrary to my own supposition, as not

¹ Richard Feldman, *Epistemology* (Upper Saddle River: Pearson, 2002), 36-37.

² William Lycan, "On the Gettier Problem Problem," in *Epistemology Futures*, ed. Stephen Hetherington (New York: Oxford University Press, 2006), 156-157 and 166.

³ Also see Gilbert Harman, *Thought* (Princeton: Princeton University Press, 1973), 47.

⁴ Peter Klein, "Useful False Beliefs," in *Epistemology: New Essays*, ed. Quentin Smith (New York: Oxford University Press, 2008), 25-61.

being a case of knowledge, then it can be used as a potential counterexample against both safety and sensitivity views of knowledge. I also provide a further case which, I claim, is problematic for all of the accounts just mentioned. Since these types of views are among the most popular analyses of knowledge, this result is significant. I then propose, briefly, an alternative account of knowledge which handles all these cases appropriately.

Although neither Feldman nor Lycan spells out NEFA in great detail, we can stipulate that NEFA is the view that *a doxastic agent possesses knowledge only if there are no assumptions which are essential to the agent's belief which are false*. It is easy to see how NEFA is motivated. Standard Gettier-style cases involve situations where a doxastic agent's belief essentially rests on a false assumption. For instance, in the example from Chisholm⁵ where a doxastic agent's belief that there is a sheep in a field is justified on the basis of the sight of a realistic sheep statue, the agent does not know that there is a sheep in the field even if there is a real sheep hiding behind some trees, because the agent's belief is essentially based on the assumption that the statue is a sheep. This fact does seem to explain why the agent lacks knowledge in that case (and in many other Gettier-type cases). If, on the other hand, the sheep statue co-exists with a flock of genuine and clearly visible sheep, then the fact that the doxastic agent bases her belief that there are sheep in the field in part on the fake sheep does not entail that she lacks knowledge, because the agent's belief that that particular thing is a sheep is not an *essential* feature of her justification. Both Feldman and Lycan, in reviewing the post-Gettier literature, consider no-false-lemma and defeasibility accounts of knowledge and show why those accounts fail. The NEFA account is similar, but arguably avoids the problems that they both claim that no-false lemma and defeasibility accounts face; in particular, those accounts count certain cases as non-knowledge when there is an assumption used by the doxastic agent which is false but which is not *essential* to the agent's justification.

II. Klein's refinement of NEFA

Klein discusses several examples of "useful falsehoods." Although Klein does not quite put it this way, these are cases where the doxastic agent's belief is not essentially based on the falsehood even if the agent's reasoning does in fact pass through the falsehood. Here is one of Klein's examples, the *Appointment Case*:

⁵ Roderick Chisholm, *Theory of Knowledge*, 3rd Edition (Englewood Cliffs: Prentice Hall, 1989), 93.

Knowledge Essentially Based Upon False Belief

On the basis of my apparent memory, I believe that my secretary told me on Friday that I have an appointment on Monday with a student. From that belief, I infer that I have an appointment on Monday. Suppose, further, that I do have an appointment on Monday, and that my secretary told me so. But she told me that on Thursday, not on Friday. I know that I have such an appointment even though I inferred my belief from the false proposition that my secretary told me on Friday that I have an appointment on Monday.⁶

According to Klein, this is a case where a false belief is essential in the causal production of a belief which counts as knowledge: the belief that the appointment is on Monday. Klein then claims that the reason why the agent has knowledge despite the causal role that the falsehood plays is that there is *another* proposition – the proposition that the secretary said that the appointment is on Monday – which meets three conditions: (1) it is also justified by the apparent memory that the secretary said that the appointment is on Monday, (2) it is true, and (3) it justifies the belief that the appointment is on Monday. Because there is available to the doxastic agent a second proposition that meets these conditions – even if the agent doesn't explicitly believe the proposition – the fact that the doxastic agent's reasoning passes through the false belief does not undermine the agent's possession of knowledge.

Although Klein does not see himself as advocating NEFA, his discussion can easily be seen as a clarification of NEFA.⁷ An advocate of NEFA can hold, in light of Klein's examples, that there is a difference between an assumption being essential in the causal production of a belief and the assumption being essential to the justification of the belief – assumptions may satisfy the former but not the latter. The belief that the secretary said on Friday that the appointment is on Monday turns out not to be epistemically essential to the belief; what is essential is just that the secretary said that the appointment is on Monday, and that proposition meets the three conditions above. In general terms, *when a doxastic agent's evidence for a false belief is such that it propositionally justifies a true proposition which on its own propositionally justifies the agent's belief, the false belief itself is not essential to the agent's belief.* Understood this way, Klein is making a clarification of NEFA.

Perhaps inessentiality could also be measured by considering what the agent would believe if the agent were informed that the secretary did not say on

⁶ Klein, "Useful False Beliefs," 36.

⁷ I focus on NEFA in this paper because it is simpler than Klein's own account, though the example I give in Section III of this paper which I claim is a problem for NEFA is also a problem for Klein's account. Thus complexities of Klein's view aside from ones discussed here are not relevant for my purposes.

Avram Hiller

Friday that the appointment is on Monday. If the agent were then to suppose or discover that it is not the case that the secretary made the statement on Friday, then the agent would still believe that the appointment was on Monday because the agent would believe not that the secretary never made the claim at all, but that the secretary made the claim on a different day. Because the belief that the appointment is on Monday would withstand a supposition that the false belief *is* false, the false belief is not really essential to the belief that the appointment is on Monday. If, on the other hand, the agent knows that for some reason, for instance, what the secretary says on days other than Friday is typically unreliable, then the agent would not know that the appointment is on Monday, even if it is still true. If the agent were to suppose or learn that it is false that the secretary said on Friday that the appointment is on Monday, the agent would abandon his belief that the appointment is on Monday. This is another way of demonstrating whether a belief is essential or not.

III. A counterexample to the refined NEFA

However, there are examples similar to ones considered by Klein which cannot be handled by Klein's refinement. Consider the *Spy Case*:

Natasha is a spy in the field. Messages to her from Headquarters often are detected by enemy intelligence, and Headquarters is aware of that fact. Today, Headquarters needs to communicate to Natasha that her contact will be at the train station at 4:00 pm, but Headquarters cannot directly tell her so. However, Headquarters knows that Natasha happens to have a justified false belief that the train from Milan is arriving at 4:00 pm. (It really arrives at 8:00 pm; also, there are no signs posted at the station indicating at what time it will arrive, and thus it is very unlikely that Natasha will find out the truth about the train's arrival time.) Headquarters knows that the enemy does not know that she has this false belief. So Headquarters sends a communiqué to Natasha stating that her contact is on the train from Milan. Natasha goes to the station at 4:00 pm and meets her contact there.

My claim is that Natasha did know that P, the proposition that *the contact will be at the station at 4:00 pm*, even though her belief is essentially based on two false assumptions, the assumption that the contact is on the train and the assumption that the train arrives at 4:00 pm.

Before further analysis of this case, I should note that both sensitivity and safety accounts of knowledge handle this case quite nicely. According to a sensitivity view, it is a necessary condition on knowledge that in the nearest worlds in which the proposition is false, the agent does not believe the proposition

(using the same method).⁸ According to a safety view, it is a necessary condition on knowledge that in the nearest worlds in which the believer has the belief, the belief is true.⁹ We can assume that Natasha does properly track the truth of when the contact would be at the station: if the contact were arriving at 5:00 pm, she would not have formed the belief that P since Headquarters would not have given her the same information. In all the nearest worlds in which she forms the belief that P, the belief is true (assuming that the contact is highly dependable in arriving on time). There are thus no nearby worlds in which Natasha believes that P but P is false. Thus Natasha's belief that P is both sensitive and safe.

Now, is this really a case where Natasha's belief that P depends *essentially* on a false assumption in light of Klein's refinements of NEFA? One might thus wonder whether there is *another* true proposition which is justified by Natasha's evidence and which itself justifies the proposition that P. In particular, Headquarters has attempted to convey to Natasha that the contact will be at the train station at 4:00 pm, and perhaps that proposition on its own is justified by Natasha's evidence and itself justifies P, and thus renders the falsehoods in Natasha's justification inessential.

To support that interpretation, consider the following details added to the case. (Call the case with these added details the *Backup Justification Case*, and I shall refer to the doxastic agent in this case as Natasha_B)

- (A) Natasha_B knows that Headquarters knows that she believes that the Milan train is coming at 4:00 pm.
- (B) Natasha_B knows that Headquarters knows that it is unlikely that that she would get any evidence to the contrary.
- (C) Natasha_B has had experiences with Headquarters in the past where they have given her false information which together with other beliefs they knew she had led her to infer the truth of some important proposition.

(A), (B), and (C), together with the fact that Headquarters has told Natasha_B that the contact is on the Milan train, would be an epistemically adequate basis for Natasha_B to believe that P. We can further add a modal stipulation: if somehow Natasha_B were to stumble upon the fact that the train does not arrive until 8:00 pm, in light of the considerations just given, Natasha_B would still believe that P. Her belief that P is thus modally robust. This indicates that the content of the two false beliefs is not essential to Natasha_B's justification even though her having the false beliefs plays the causal role in her *actually* coming to believe that P. Her

⁸ See Robert Nozick, *Philosophical Explanations* (Cambridge: Harvard University Press, 1981), 167-196.

⁹ See for instance Duncan Pritchard, *Epistemic Luck* (New York: Oxford University Press, 2005).

evidence justifies the proposition that Headquarters has attempted to convey to her that the contact will be at the station at 4:00 pm, and this proposition, given the background information $Natasha_B$ possesses, propositionally justifies the proposition that P for Natasha even if Natasha does not consciously entertain that line of reasoning.¹⁰

But what if Natasha lacks those other pieces of information? Let's assume that (A), (B), and (C) do not in fact obtain, and thus there is no indirect backup justification that P for Natasha. Thus in the Spy Case I wish to consider, Natasha does not know that Headquarters knows that she believes that the train is arriving at the station at 4:00 pm and she also has not had numerous past experiences where they have conveyed messages to her in the past using false information. However, my own sense is (quite strongly) that this main Spy Case is still a case of Natasha knowing that P. Natasha has a true belief that P and is justified in forming her belief that P. Even if Natasha and $Natasha_B$ differ dispositionally, their actual belief processes are the same; they both consciously form the belief that the contact will be at the station at 4:00 pm on the same basis – that the contact is on the 4:00 pm train. Importantly, as above, Natasha, like $Natasha_B$, tracks the truth that P. In all the nearest worlds in which Natasha believes that P, P is true, and in all the nearest worlds in which P is true and Natasha forms a belief about P using the same method, Natasha believes that P. Typical instances of safety and sensitivity are ones in which the agent herself is primarily epistemically responsible for tracking the truth. This case differs in that it is the epistemically-friendly Headquarters which manages the epistemic environment so that Natasha tracks the truth. But this should not be seen as a reason not to attribute knowledge that P to Natasha – she has knowledge even if it is not her own but someone else's epistemic efforts which ensure the safety/sensitivity of her belief. Although sensitivity and safety accounts have been subject to counterexamples, the intuitions which they employ are still quite strong, and I do not see a reason why this case should be considered to be a counterexample case. My own intuitions thus coincide with safety and sensitivity intuitions in the main Spy Case.¹¹

I take it that this is a case of knowledge even if, in the far-off possible world in which Natasha somehow stumbles upon the information that the train is

¹⁰ The analysis in this paragraph is analogous to Klein's analysis of his Santa Claus case ("Useful False Beliefs," 57), which I discuss below in Section VII.

¹¹ I am no friend of safety accounts of knowledge; see below and also Avram Hiller and Ram Neta, "Safety and Epistemic Luck," *Synthese* 158 (2007): 303-313.

arriving at 8:00 pm,¹² she would come to doubt that the contact would be arriving at the station at 4:00 pm. She might speculate that Headquarters *may* have known that she had a false belief that the train arrives at 4:00 pm and that that is why they told her that the contact is on the train, but that would not be enough to adequately justify her continuing to maintain a high level of credence that P. Natasha's belief really is essentially based upon the two false assumptions that the contact is on the train and that the train arrives at 4:00 pm. Still, Natasha knows that P given the way things have worked in the actual world.

I grant that intuitions may differ; perhaps advocates of NEFA would insist that Natasha does not know that P. I cannot prove that this is incorrect. At the very least, I expect that a neutral reader will still feel some amount of pull in the direction that Natasha knows that P. However, if this case is indeed taken as a definite case of non-knowledge, then it immediately becomes a serious problem for safety and sensitivity views. As above, there are no nearby worlds in which Natasha has a false belief that P. There are, however, some far-off worlds in which Natasha discovers that the train from Milan arrives at 8:00 pm, not at 4:00 pm, but since these are, as stipulated, far-off worlds, they do not undermine Natasha's knowledge that P. Furthermore, they are worlds in which Natasha forms a belief about P using a different method than the one she uses in the actual world. Natasha's belief that P is thus safe and sensitive. In all the nearest worlds in which Natasha forms a belief about P using the same method that she actually uses, P is true, and in all the nearest worlds in which P is true and Natasha uses the same method of forming a belief about P that she actually uses, Natasha truly believes that P.

This would then be a case of safe non-knowledge. Now, advocates of safety merely suggest that safety is a necessary condition on knowledge, and do not suggest that safety is sufficient. The example is thus not a direct counterexample to safety accounts. However, it is unclear how *else* an advocate of safety could show that this is a case of non-knowledge. Moreover, in other examples given by advocates of safety accounts where a doxastic agent has a justified true belief but not knowledge, it is the violation of the safety condition which typically does the work in showing why the belief is not knowledge. Thus this example, if it is interpreted as being a case of non-knowledge, would be seriously problematic for advocates of a safety condition.

This is also a case of sensitive non-knowledge. Interestingly, Nozick provides a further condition on knowledge that if P were true, and S were to form

¹² I am assuming that her stumbling on that information is a very unlikely and modally distant possibility but not an impossibility.

Avram Hiller

a belief about P using the same method, S would still believe that P. This condition comes close to not being met in the Spy Case – but not close enough. Again, firstly, the worlds in which P is true but Natasha does not believe that P are far off. Second, in those worlds, Natasha forms a belief about P using a different method – she deduces the information about the contact’s arrival time from a different source than in the actual world. Thus if the Spy Case is viewed as a case of non-knowledge, it is a counterexample to Nozick’s account.

IV. A problem case for safety and sensitivity

As I have said, I do view the Spy Case as a case of knowledge, and thus I do not regard it as a problem for safety or sensitivity. However, there is a somewhat similar case which I believe is a genuine case of non-knowledge and which appears to be a more serious problem for safety and sensitivity views. Imagine a case which differs only slightly from the Spy Case. In what I will call the *Cognitive Defect Case*, Natashac has a cognitive defect: whenever she hears about a train, she believes that the train arrives at its destination at 4:00 pm. She has had plenty of evidence that trains arrive at other times, but she always forms an unjustified belief that any future train will arrive at 4:00 pm. Headquarters is aware of this defect, but the enemy is not. In this example, Natashac does not have any prior belief that the train from Milan is arriving at 4:00 pm – she is, perhaps, unaware that there even is a train from Milan – but Headquarters tells her simply that the contact is on the train from Milan. She then infers the unjustified but true belief that the contact will be at the station at 4:00 pm.

The Cognitive Defect Case involves beliefs which are just as safe and sensitive to the truth as in the Spy Case, because Headquarters would not have given Natashac the false information if it didn’t know that she would make an inference to the true belief. There are no nearby possible worlds in which Natasha believes that P and P is false. Nevertheless, this case appears to be a case of non-knowledge – Natashac’s belief in P is unjustified. Thus the Spy Case is a case of knowledge and the Cognitive Defect Case is not, but in the two cases, the two Natashas’ beliefs are equally safe and equally sensitive. Even if the reader does not share these clear intuitions, as long as the reader feels more of a pull for the Spy Case to be a case of knowledge than the Cognitive Defect Case, then safety/sensitivity views are cast into doubt. Advocates of safety and sensitivity could add a justification condition as another necessary condition on knowledge. However, the spirit of safety and sensitivity views is that they avoid the need to add a justification condition on knowledge; and furthermore, such a resulting theory would be theoretically inelegant since many cases of unjustified belief are

also cases of unsafe/unsensitive belief; the two conditions would do overlapping work.

V. Another problematic case for the refined NEFA account

NEFA does get the Cognitive Defect Case right – there is an essential false assumption that the train arrives at 4:00 pm – and it is not a case of knowledge. But if it is indeed a case of non-knowledge, it is not the falsehood of the assumption which makes it non-knowledge; it is the fact that the false assumption was acquired in an unjustified manner. Thus this case should not be helpful for an advocate of the refined NEFA account even though NEFA does deliver the correct answer to it.

This fact can be brought to light further by considering a final case, the *Backup Justification Cognitive Defect Case*. In it, Natasha_D has the same cognitive defect as Natasha_C has. However, Natasha_D also has the kind of backup information that Natasha_B has – (A), (B), and (C) above hold for Natasha_D. Thus on a refined NEFA view, even though Natasha_D's own belief process passes through a false belief, she still has at her disposal a backup justification which meets the condition stated in Section II – her evidence is such that it propositionally justifies a true proposition which itself is an adequate basis for Natasha_D to believe that P. Thus this case would be deemed by the refined NEFA view as a case of knowledge even though it is not.

Of course if NEFA is explicitly stated¹³ as a view that a doxastic agent has knowledge if and only if the doxastic agent has a *justified* belief which is not essentially based on a false assumption, then the Backup Justification Cognitive Defect Case isn't a counterexample to it since Natasha_D is unjustified and does not have knowledge. But the cases considered above together bring out something odd about such a refined NEFA view – why does the agent need to be justified at all when the agent can have knowledge merely on the basis of the existence of a potential justificatory structure which she does not actually employ? If Natasha_B knows that P because her evidence is such that it justifies a chain of reasoning that she does not employ (the assumptions that Headquarters has attempted to convey to her the message that the contact will be at the station at 4:00 pm and that Headquarters is usually correct when they convey the message) and not because of her actual conscious justification (that the contact is on the train and that the train arrives at 4:00 pm), then why is actual justification needed for an agent to have knowledge?

¹³ As Feldman does in *Epistemology*, 37.

VI. An alternative to both NEFA and safety/sensitivity

Since these cases are fairly complex and intuitions may differ, one might wonder whether they may be instances of “to the winner go the spoils.” In other words, if either NEFA or safety/sensitivity were shown to be mostly successful accounts of knowledge on the basis of independent argumentation, then we should then simply adopt the intuitions that the otherwise correct view tells us we should have in these cases. What I’d like to suggest is that the intuitive responses I report can be systematized in another way; this provides evidence that we should not adopt a ‘spoils’ attitude in the cases considered above.

On an account of knowledge I develop elsewhere,¹⁴ knowledge involves (internalistically) justified belief which is formed in an epistemic environment which is conducive to (internalistically) justified believers forming true beliefs relevantly similar to the belief in question. Although a complete elaboration of that view is well outside the scope of this paper, I should note that the cases considered above fit this view of knowledge quite well. The Spy Case and the Backup Justification Case are both cases of justified true beliefs, and furthermore, they are both cases where, thanks to helpful oversight of Headquarters, Natasha is in a good epistemic environment to form true beliefs about the whereabouts of her contact. These cases trade on a peculiar feature of testimonial evidence. Typically, one acquires knowledge via testimony when a cooperative testifier states truths. These two cases maintain the spirit of knowledge via testimony – the testifier, Headquarters, is being epistemically cooperative in conveying a true target proposition to a doxastic agent, albeit using some unusual means. Headquarters sees to it that the epistemic environment is a good one for Natasha to form beliefs about the arrival time of her contact.

In the Cognitive Defect Cases as well, Headquarters is also overseeing the epistemic environment in a cooperative way. However, because Natasha is unjustified in using the internal processes that she does, she does not count as having knowledge despite being in a good epistemic environment. To possess knowledge, one must not only be in a good epistemic environment, but one must process information in a justified manner. Thus these cases taken together support the idea that there are two necessary conditions – the agent’s having a justified belief and the agent’s being in a proper epistemic environment – on knowledge. The fact that NEFA needs to appeal to a possible alternate route of justification of a belief in order to show that certain cases of “useful falsehoods” are still knowledge appears to undermine the inclusion of a justification condition on

¹⁴ Avram Hiller, “Knowledge as Justified Stable Belief,” manuscript.

knowledge. But on a view which in letter and spirit is quite different from NEFA, the inclusion of a distinct justification condition is not undermined by such cases.

This of course is a very brief suggestion, and I do not introduce the cases in this paper to prove my own account. Rather, I mention this view to show that there is a plausible alternative to both the NEFA and safety/sensitivity views which handles all these cases successfully. At the very least, one should not jump to a “winner gets the spoils” response without considering other potential ways of accounting for all the cases.

VII. Klein’s Santa Claus case

Klein himself considers a case similar to the Spy Case, the Santa Claus Case:

Mom and Dad tell young Virginia that Santa will put some presents under the tree on Christmas Eve. Believing what her parents told her, she infers that there will be presents under the tree on Christmas morning. She knows that.¹⁵

One might view this case as fully analogous to the original Spy Case, and if it is, then my own discussion would not be original. Klein himself analyzes the case in a way analogous not to the original Spy Case but to the Backup Justification Case. Klein claims that the facts that Virginia’s parents told her that Santa would put presents under the tree and that her parents are normally reliable truth-tellers justify for her the proposition that someone will put presents under the tree.¹⁶ For Klein, Virginia’s evidence gives her an adequate basis to believe that there will be presents under the tree even if her own actual belief process is causally based upon a falsehood. If this is the proper analysis of the case, then it is simply a different case than the Spy Case since in the Spy Case there is no backup justification available to the doxastic agent.

I myself am unsure whether to say that Virginia has knowledge. Perhaps Virginia’s belief may be justified on the basis of memories of past years’ Christmas gifts. Even if she were to discover that there is no Santa, she knows that someone has been putting presents under the tree in years past. This would be another backup justification for Virginia. On the other hand, if Virginia is so young that this is the first Christmas where she is capable of understanding what her parents tell her, then arguably she doesn’t know that there are presents under the tree because it is essentially based on her false beliefs and not on her parents’ reliability, which she is perhaps too inexperienced to grasp. Furthermore, if Virginia believes in the highly far-fetched belief that a fat elf in a red suit flies

¹⁵ Klein, “Useful False Beliefs,” 37.

¹⁶ Klein, “Useful False Beliefs,” 57.

Avram Hiller

from the North Pole into every good Christian child's chimney to place presents under the tree on Christmas Eve, she may count as being in the same category as Natashac – not rational enough to count as having knowledge in the arena in question. Thus even if we accept that Virginia's parents have created a good epistemic environment for her to form beliefs about whether there will be presents under the tree, we need not grant that Virginia knows that there will be presents under the tree.

Another complication is that it may also be that Santa *does exist*; in particular, Santa is Virginia's parents, if 'Santa' in Virginia's idiolect can be analyzed as being an abbreviation for the definite description "whoever places presents under the tree." As evidence for this, some kids might say, upon finding out that no one sneaks into chimneys on Christmas Eve, "Mom and Dad are Santa!" Of course, some aspects of the typical description of Santa will not be true of Virginia's parents, such as the fact that Virginia's parents do not live on the North Pole, but perhaps the core part of the description – the only part which determines the referent in Virginia idiolect – is that Santa is whoever it is who places presents under the tree. Thus Virginia's belief would not be false, since even though her parents don't intend the word 'Santa' to be a definite description, it may function that way in Virginia's idiolect. In that case we would have a case of knowledge which is not based in any way on a falsehood.

I am unsure which of these considerations do and do not apply to the Virginia case. I bring up these considerations to explain why our intuitions may be pulled in one way or another by that case. One who accepts that Virginia does have knowledge need not accept my (or a refined NEFA) view of knowledge, and one who does not accept that Virginia has knowledge need not reject my own view. The reason I focus on the Spy Cases is that they clearly distinguish between several possibilities which may be at play in the Santa Case.

VIII. Conclusion

In sum, I have discussed several cases which show that either NEFA or safety (or both) is an inadequate account of knowledge. Since my methodology involves the use of thought experiments about which intuitions may differ, I cannot claim that the examples provide a definitive refutation of both of those views. However, they do undermine at least one of the views, since intuitions about the cases which are friendly to one will be unfriendly to the other. At the very least, advocates of each of these views need to account for why there is *some* intuitive pull in the direction opposite to the intuitive responses given by the views in these cases. My brief suggestion in Section VI is that we should not try to make minor tweaks of one or

the other of these accounts, since there is a plausible alternative account of knowledge which handles the cases appropriately and which thus demands further examination.¹⁷

¹⁷ I'd like to thank Ram Neta and an audience at the Reed College Philosophy Department for discussion of these issues.