Kelp, C. (2008) Classical invariantism and the puzzle of fallibilism. Southern Journal of Philosophy, 46(2), pp. 221-244.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

http://eprints.gla.ac.uk/140928/

Deposited on: 15 May 2017

Classical Invariantism and the Puzzle of Fallibilism

Christoph Kelp
Institute of Philosophy, KU Leuven
christoph.kelp@hiw.kuleuven.be

Abstract. This paper revisits a puzzle that arises for theories of knowledge according to which one can know on the basis of merely inductive grounds. No matter how strong such theories require inductive grounds to be if a belief based on them is to qualify as knowledge, there are certain beliefs (viz. about the outcome of fair lotteries) that are based on even stronger inductive grounds, while, intuitively, they do not qualify as knowledge. This paper discusses what is often regarded as the most promising classical invariantist solution to the puzzle, viz. that beliefs about the outcomes of fair lotteries do not qualify as knowledge because they are too lucky to do so (or, relatedly, because they do not satisfy a safety-condition on knowledge), while other beliefs based on potentially weaker inductive grounds are not too lucky (or, relatedly, because they are safe). A case is presented that shows that this solution to the puzzle is actually not viable. It is argued that there is no obvious alternative solution in sight and that therefore the puzzle still awaits a classical invariantist solution.

1. Introduction: The Fallibilist Puzzle

Let fallibilism be the view that one can know a proposition, *P*, on merely inductive grounds.[1]

By 'inductive grounds' I mean grounds that do not entail the truth of the proposition for which they are grounds. Fallibilism is an appealing view. In fact, it seems to be part of our ordinary conception of knowledge. Consider, by way of illustration, the following case:

The Gift.   A friend of mine whom I have invited for dinner gives me a bottle. I look at the label, read the word 'Vino', see the liquid in the bottle, see that it has the general appearance of a bottle of wine, and come to believe, truly, that there is wine in the bottle.

Intuitively, my true belief that there is wine in the bottle counts as knowledge. At the same time, the grounds I have for my belief are merely inductive. It is possible that I look at a bottle that has the general appearance of a bottle of wine, see the liquid inside, read the word 'Vino' on the label, while at the same time, there isn't any wine in the bottle: By some accident, the

bottle has been filled with some other liquid such as water. Intuitively, I know even though the grounds for my belief are merely inductive. Moreover, there is a vast number of cases in which, intuitively, one knows a proposition although one's underlying belief is based on merely inductive grounds. In this way, fallibilism is part of our ordinary conception of knowledge.

As Stewart Cohen [Cohen 1988] has pointed out, however, fallibilists face a puzzle. The reason for this is that there are cases in which subjects have true beliefs based on as good inductive grounds as (or even better inductive grounds than) the ones I have in *The Gift*, whilst, intuitively, their beliefs do not qualify as knowledge. The cases at issue concern beliefs about the outcomes of fair lotteries. Here then is an exemplary case:

The Lottery.   A friend of mine has bought a ticket in a fair lottery with *n* tickets (where the value of *n* is a suitably large natural number). On the basis of the fact that the probability of winning is 1/*n* I come to believe, truly, that his ticket won't win the lottery.

Intuitively, in this case, my true belief that my friend's ticket won't win the lottery does not qualify as knowledge. I can't know, in advance of the drawing, that a given ticket in a fair lottery won't win. At the same time, the inductive grounds I have for my belief may be as good as (or even better than) the inductive grounds I have for believing that there is wine in the bottle in *The Gift*. After all, the number of tickets in the lottery may be indefinitely large. That means that, for some value of *n*, the probability of losing the lottery is as high as (and for some value of *n* it is even higher than) the probability of there being wine in the bottle given that the bottle is labelled 'Vino', has general appearance of a bottle of wine etc. So, the inductive grounds I have for believing that my friend's ticket won't win the lottery may as good as (or even better than) the inductive grounds on which I believe that there is wine in the bottle. At the same time, intuitively, my belief in *The Lottery* that my friend's ticket won't win the lottery does not qualify as knowledge, while, intuitively, my belief in *The Gift* that there is wine in the bottle does qualify as knowledge. In this way, fallibilists face a puzzle.

How can the puzzle be solved? One option is to abandon standard epistemology and go for some non-standard epistemological theory. By 'non-standard epistemological theory' I mean theories according to which a sentence of the form "S knows that p" can have different truth-values even though subject, S, proposition, P, and the subject's epistemic position towards the proposition are held fixed. By way of example, contextualism[2], subject-sensitive invariantism[3], contrastivism[4], and relativism[5] are all non-standard epistemological theories so understood. They differ in how exactly the non-standard thesis is implemented. The general strategy by means of which non-standard epistemological theories will attempt to solve the fallibilist puzzle isn't hard to figure out. They will argue that whatever factors over and above subject, proposition, and the subject's epistemic position towards the proposition play a part in the determination of the truth-value of sentences of the form "S knows that p" will conspire in such a way as to predict that in *The Gift* the sentence "Chris knows that there is wine in the bottle" is assigned the value *true*, while, at the same time, in *The Lottery* the sentence "Chris knows that his friend's ticket won't win the lottery" is assigned the value *false*. Be that as it may, what I will be concerned with in the present paper is what on the face of it is the most natural as well as the most promising attempt to tackle the puzzle fallibilism faces within a standard epistemological (or *classical invariantist*) framework, a framework, that is, according to which the truth-values of sentences of the form "S knows that p" do not vary once subject, proposition, and the subject's epistemic position towards the proposition are settled. I will argue that this solution to the puzzle of fallibilism fails and that therefore it is still not clear that classical invariantism is a viable theory of knowledge. In order to do so I will first outline the credentials of the abovementioned solution to the puzzle of fallibilism. I will then adduce a non-standard lottery case which not only shows that this solution to the puzzle of fallibilism does not work, but, which is even worse, also threatens to lead the classical invariantist to contradiction. Finally, I will consider the ways in which the classical invariantist may venture to escape the threatening contradiction. I will argue that none of

these ways is very promising since on each such way classical invariantism either incurs unacceptable intuitive costs or is left without a solution to the puzzle of fallibilism.

2. The Anti-Luck Platitude, Safety, and a Classical Invariantist Solution to the Fallibilist Puzzle

In order to solve the puzzle of fallibilism the classical invariantist has to provide some explanation as to how the two cases, that is, *The Gift* and *The Lottery*, differ. A natural way of doing so and a way that a number of invariantists have found most promising is the following. It is widely regarded as a platitude about knowledge that knowledge excludes luck. Or, to be more precise, knowledge is true belief such that it is not at all a matter of luck that the belief is true.[6] The classical invariantist may now argue that the difference between my beliefs in *The Gift* and *The Lottery* is that in the latter but not in the former my belief is too luckily true to qualify as knowledge. Or, in other words, that in *The Gift* but not in *The Lottery* it is not at all a matter of luck that my belief is true. After all, the classical invariantist may point out, there is always an element of luck in what the outcome of a fair lottery draw turns out to be, that is, which ticket will win and which ones will lose. Accordingly, when I form a belief, in advance of the drawing, that my friend's ticket will lose, whether my belief turns out to be true will be subject to this element of luck, too. So, in *The Lottery* my true belief that my friend's ticket won't win is true in a way that is subject to this element of luck. As opposed to that, when I look at a bottle that is labelled 'Vino', has the general appearance of a bottle of wine etc. and on that basis form a true belief that there is wine in the bottle, it is not at all a matter of luck that my belief is true. The events in the world that have led up to the situation in which I form my belief—the reliable process of wine production, the fraud controls, the fact that my friend bought the bottle in an excellent wine store etc.—guarantee this much. The classical

invariantist may thus seem to be able to resolve the puzzle of fallibilism by appealing to the platitude that knowledge excludes luck.

Alternatively, the classical invariantist may venture to solve the fallibilist puzzle by appeal to the so-called 'safety-condition' on knowledge. The core idea here is that knowledge requires safe belief. The notion of safe belief is then developed in the following way:

(SB)   One's belief that p formed via a method, m, is safe only if not easily would one have acquired a false belief that p via m.[7,8]

Or, in terms of a standard possible worlds semantics of the relevant modal notions, one's belief that p formed via method, m, is safe only if at no nearby possible worlds one acquires a false belief that p via m. Distance of possible worlds here is measured in terms of similarity. The more similar a world is to the actual world, the closer (more nearby) it is.

If knowledge requires safe belief, we can then derive the following condition on knowledge:

(K)   One knows that p only if one (truly[9]) believes that p via m and at no nearby possible worlds one acquires a false belief that p via m.[10]

There is reason to believe that the safety-condition on knowledge captured in (K) also solves the fallibilist puzzle. After all, it would seem that, as regards *The Lottery*, my belief that my friend's ticket won't win the lottery might easily turn out to be false. Although the chances of losing may be indefinitely high, it might easily be the case that his ticket wins the lottery. Or, in terms of a standard possible worlds semantics of the relevant modal notion, my friend's ticket wins at some nearby possible world. After all, the possible world at which my friend's ticket wins is very similar to the actual world. All that needs to differ is that some numbered balls fall into different places. Since at the nearby world at which my friend's ticket wins the lottery I will continue to believe that his ticket will lose on the basis of the probabilistic evidence against winning, there is a nearby world at which I form a false belief even though I

5

form my belief in the same way as in the actual world. Accordingly, (K) will predict that my belief does not qualify as knowledge.

Consider now my belief in *The Gift*. Does my belief that there is wine in the bottle satisfy the safety-condition on knowledge captured in (K)? It would seem that the answer to this question is 'yes'. There just isn't a nearby possible world at which I am handed over a bottle labelled 'Vino' which also has the general appearance of a wine bottle at which there isn't any wine in the bottle. The reliable production procedures, the strict fraud controls, the fact that my friend has bought the wine at an excellent wine store etc. guarantee that worlds at which there isn't any wine in the bottle are quite different from the actual world. Accordingly, such worlds are not nearby. That means, however, that there is no nearby world at which I form a belief that there is wine in the bottle via looking at a bottle labelled 'Vino' which also has the general appearance of a bottle of wine etc. and yet acquire a false belief that there is wine in the bottle. Accordingly, (K) will not predict that my belief that there is wine in the bottle does not qualify as knowledge. Since my belief in *The Lottery* that my friend's ticket won't win the lottery does not satisfy (K), while my belief in *The Gift* that there is wine in the bottle does satisfy (K), the safety condition on knowledge can also explain how the two cases differ. In this way, it also seems to be able to provide a solution to the fallibilist puzzle.

We have looked at two solutions to the fallibilist puzzle. Both of them seem promising. According to one solution my belief in *The Lottery* but not in *The Gift* is too luckily true to qualify as knowledge, while according to the other the latter but not the former satisfies the safety condition of knowledge captured in (K). What I will do next is to elucidate the relation between the two solutions to the fallibilist puzzle. More specifically, I will argue that the solution which appeals to the safety-condition on knowledge is a special case of the solution in terms of the anti-luck platitude. The safety-condition, (K), captures in more detail the part of the anti-luck platitude which solves the fallibilist puzzle. To a certain extent my

argument will follow Duncan Pritchard's argument that the safety condition on knowledge captures the sense in which knowledge excludes luck.[11]

In his argument, Pritchard first develops an independently plausible conception of luck. He argues that there is a modal dimension to the notion of luck that is captured in the following conditional:

(L1)   If an event is lucky, then it is an event that occurs in the actual world but does not occur in a wide class of [nearby][12] possible worlds where the relevant initial conditions for that event are the same as in the actual world. [Pritchard 2005: 128]

Pritchard argues that there is a variety of facts that provide independent support for this modal dimension of the notion of luck. First, it appears to be required to match our intuitive verdicts of lucky events across a range of cases. Second, it is continuous with and indeed serves to explain some of the results of recent scientific research on luck. Third, it allows us to explain what is correct about alternative but ultimately unsatisfactory accounts of luck such as accounts that try to analyse (or elucidate) luck in terms of control, accident or chance [cf. Pritchard 2004a, 2005].

Given (L1), we are also in a position to characterise the idea of a non-lucky event and the notion of an event that is not at all lucky. An event is non-lucky only if it occurs in the actual world and there is no wide range of nearby possible worlds with the same relevant initial conditions at which it does not occur. An event is not at all lucky only if it occurs in the actual world and there is no range (however small) of nearby possible worlds with the same relevant initial conditions at which it does not occur.

With the idea of an event that is not at all lucky thus characterised, the notion of a true belief that is not at all luckily true is naturally characterised along the following lines:

(NLTB*)   One's true belief that *p* is not at all luckily true only if one believes that *p* truly in the actual world and there is no range (however small) of nearby possible worlds in which one believes that *p* and the relevant initial conditions for the formation of the belief are the same as in the actual world while one's belief that *p* is not true.

Provided that according to the anti-luck platitude knowledge is true belief that is not at all luckily true, we can now derive the following necessary condition for knowledge.

(NLK) One knows that *p* only if one believes that *p* truly in the actual world and there is no range (however small) of nearby possible worlds in which one believes that *p* and the relevant initial conditions for the formation of the belief are the same as in the actual world while one's belief that *p* is not true. [Cf. Pritchard forthcoming: 7][13]

Provided, as is plausible, that the method of belief-formation is among the relevant initial conditions for the formation of belief, (K) can be derived from (NLK). In this way there is reason to believe that (K) identifies in more detail part of the sense in which knowledge excludes luck. Provided that (K) also solves the fallibilist puzzle there is also reason to believe that (K) identifies in more detail the part of the sense in which knowledge excludes luck that solves the fallibilist puzzle. Accordingly, there is reason to believe that as far as the solution to the fallibilist puzzle is concerned the fates of the anti-luck platitude and of the safety-condition on knowledge captured in (K) are tied together: one won't successfully solve the puzzle unless the other one does.

3. The Classical Invariantist Solution in Trouble

I will now present a variation of *The Lottery* that raises a problem for the classical invariantist who ventures to solve the fallibilist puzzle by appeal to the anti-luck platitude (or a safety-condition on knowledge for that matter). Here is the case:

The Lottery*. A lottery drawing is carried out in the following way: A computer is programmed to randomly select a natural number within a range *n* of natural numbers. The lottery contains the following tickets: First, for each number within *n* there is a ticket with a numeral denoting that number written on it; and, second, there is a ticket with the sentence "The computer breaks down" written on it. The rules for the lottery are that a ticket with a given numeral wins if and only if the number denoted by the numeral is selected by the computer and the ticket with the sentence "The computer breaks down" written on it wins if and only if the computer breaks down. Suppose that the number of tickets is so large that the probability that the computer-breakdown ticket will win is just as high as the probability that any other individual ticket wins the lottery—even though computers of the relevant type are highly reliable, this particular

8

computer has successfully passed various tests etc. Suppose a friend of mine owns the computer-breakdown ticket. On the basis of the relevant probabilistic evidence I come to believe, truly, that his ticket will lose.

This case, I submit, is a serious cause of concern for defenders of the classical invariantist solution to the fallibilist puzzle outlined above. To see why this is so, recall that, intuitively, one does not know, in advance of the drawing, that a given ticket in a fair lottery won't win. The explanation the defender of the above classical invariantist solution to the fallibilist puzzle offered us was that a belief that a given such ticket won't win, even if true, is too luckily true to qualify as knowledge/does not satisfy the safety-condition captured in (K). The intuition, I take it, carries over to the present case. Among other things that means that I do not know, in advance of the drawing, that the computer-breakdown ticket won't win the lottery.

Provided that I don't know, in advance, that the computer-breakdown ticket won't win, we must expect the defender of the above solution to the fallibilist puzzle to explain my lack of knowledge by the idea that my belief is too lucky to qualify as knowledge/does not satisfy (K). Luck has not been eliminated in the way required for knowledge. Provided, as we have found plausible, that the sense of luckiness at issue is captured in the failure of my belief to satisfy (K) that means that there is a nearby world at which I form a false belief via the same method of belief-formation.

But now notice that my belief-forming situation is otherwise much more like the belief-forming situation in *The Gift*. There I acquire a highly reliable belief that there is wine in the bottle. The reliability is ensured by such facts as the reliable process of wine production, the fraud controls, the fact that my friend bought the bottle in an excellent wine store etc. Moreover, these facts about the belief-forming situation ensure that it is not at all a matter of luck that my belief is true. This point is nicely brought out considering the situation in terms of the corresponding safety-principle. First, the worlds at which there is no wine in the bottle will be rather different from the actual world. That means, however, that they will

not be modally nearby. If they aren't modally nearby, then my belief that there is wine in the bottle satisfies the safety-principle and the resulting safety-condition on knowledge. Since there is no nearby world at which there isn't any wine in the bottle, there also isn't any nearby world at which I form a false belief that there is wine in the bottle on the basis of the inductive grounds I have. My belief that there is wine in the bottle is not at all luckily true and satisfies the safety principle.

In a similar vein, in *The Lottery**, I acquire a highly reliable belief that the computer-breakdown ticket won't win. The reliability of my belief is ensured by the same kind of facts that ensure the reliability in *The Gift*. To be more precise, here the reliability of my belief is ensured by the reliability of computers of this type, the strict production procedures, the concern of the company to lose their excellent reputation, various tests that have been performed on the computer etc. Accordingly, we ought to expect that, just as in *The Gift*, these facts about the belief-forming situation ensure that it is not at all a matter of luck that my belief is true. This expectation is further supported by looking at the predictions of the safety-principle and the safety-condition on knowledge. The worlds at which the computer breaks down are rather different from the actual world. Hence, they are not modally nearby. If they are not modally nearby, then my belief that the computer-breakdown ticket will lose is safe. Since there is no nearby world at which the computer breaks down, there also isn't any nearby world at which the computer-breakdown ticket wins. So, there is no nearby world at which I form a false belief that the computer-breakdown ticket loses when I base a belief to that effect on the inductive grounds I have for my belief. My belief that the computer-breakdown ticket won't win the lottery is not at all luckily true and satisfies the safety condition on knowledge.

For those who are suspicious of the parallelism, notice that the fact that I include a computer-breakdown ticket in some lottery does not bring worlds at which a given computer breaks down modally closer. How close worlds are at which a given computer breaks down

depends on how similar those worlds are. The inclusion of a computer-breakdown ticket in some lottery, however, need have no bearing on how similar worlds in which any computer breaks down are. Rather, how similar such worlds are will depend on how reliable the computers of this type are, in what shape the individual computer is etc. Now it would seem that the case can be described in such a way that none of the closest worlds at which the computer breaks down is nearby. At the very least, however, it can be described in such a way that if in cases like *The Gift* the closest worlds at which the proposition believed is false are not nearby, then they are also not nearby in *The Lottery**. After all, the relevant factors that bear on closeness, such as reliability etc., may approximate perfection much better in *The Lottery** than in cases like *The Gift*. In this way it can be ensured that worlds at which the computer breaks down are much more dissimilar to the actual world than worlds at which there isn't wine in the bottle, say. So, my belief in *The Lottery** will satisfy the safety condition on knowledge if my belief in *The Gift* satisfies it. Provided that the safety condition on knowledge specifies in more detail the part of the anti-luck platitude that is supposed to solve the fallibilist puzzle that means that if my belief in *The Gift* has the relevant non-luckiness features that my belief in *The Lottery* lacks, then so does my belief in *The Lottery**.[14]

This leaves defenders of the above solution to the fallibilist puzzle in a rather uncomfortable situation. According to their solution, the intuition that my belief in *The Lottery** does not qualify as knowledge has to be explained by the fact that my belief that the computer-breakdown ticket won't win the lottery is too lucky to qualify as knowledge. Luck has not been eliminated from my belief in the way requisite for knowledge/it does not satisfy (K). At the same time, the parallel with *The Gift* compels defenders of this solution to acknowledge that my belief that the computer-breakdown ticket won't win is not at all luckily true/it does satisfy (K). So, *The Lottery** seems to force defenders of the above solution to the puzzle of fallibilism to maintain both that my belief does and does not eliminate luck in the

way requisite for knowledge. Or, in terms of the terminology of the safety-condition on knowledge, *The Lottery** seems to present a case in which one does and does not form a false belief via one and the same method at nearby world.

I take it that this consequence of the above classical invariantist solution to the puzzle of fallibilism provides us with excellent evidence that the solution is not viable. In fact, things are worse. For *The Lottery** threatens to establish that the classical invariantist defender of the anti-luck platitude (and of a safety-condition on knowledge) runs into a flat-out contradiction. After all, the above considerations seem to suggest that, on the classical invariantist interpretation of the anti-luck platitude, in *The Lottery** luck both *is and is not eliminated* from my belief in the way requisite for knowledge. So the pressure is on the classical invariantist to avoid the charge of contradiction. In the next section, I will look at the prospects of different ways in which the classical invariantist may attempt to defuse the charge of contradiction.

4. Escape Routes for Classical Invariantists

It seems that in order to avoid the charge of contradiction, classical invariantists have to choose one of the following four options:

(1) They can surrender the intuition that we don't know, in advance of the drawing, that a given ticket in a fair lottery won't win. Classical invariantists who choose this option will claim that, in *The Lottery*, for instance, my belief that my friend's ticket won't win does qualify as knowledge.

(2) They can surrender the intuition that, in cases like *The Lottery**, we don't know, in advance of the drawing, that tickets like the computer-breakdown ticket won't win (and stick to the intuition that we don't know, in advance of the drawing, that tickets like the numbered tickets won't win).

(3) They can surrender the intuition that in cases like *The Gift* we know the proposition at issue.

(4) They can surrender the intuition that the problem with beliefs in lottery propositions, the reason why they don't qualify as knowledge, is that these beliefs are too lucky to count as knowledge/fail to satisfy (K). Classical invariantists choose this option will claim that we must

not expect the anti-luck platitude/the safety-condition captured in (K) to solve the fallibilist puzzle in the first place.

Let us see what happens on each of the options. First, consider classical invariantists who surrender the intuition that we can't know, in advance of the drawing, that a given ticket in a fair lottery won't win and claim instead that we can know, in advance, that a given such ticket will lose. Defenders of this proposal will, presumably, maintain that one's belief that a given such ticket won't win isn't too lucky to qualify as knowledge after all. They may substantiate their proposal by weakening the anti-luck platitude. They may deny that knowledge requires true belief that is not *at all* luckily true and instead maintain that knowledge merely requires true belief that is non-luckily true. Moreover, they can exploit Pritchard's arguments that the notion of luck involves a modal dimension and develop the weakened anti-luck platitude in the following way. First, recall, that we said that an event is non-lucky only if it occurs in the actual world and there is no wide range of nearby possible worlds with the same relevant initial conditions as in the actual world at which it does not occur. Accordingly, one's true belief that p is non-luckily true only if one truly believes that p in the actual world and there is no wide range of nearby possible worlds at which one believes that p and the relevant initial conditions for the formation of the belief are the same as in the actual world, while one's belief that p is not true. If this conception of non-lucky true belief is slotted into the weakened anti-luck platitude we get:

(NLK*)   One knows that p only if one truly believes that p in the actual world and there is no wide range of nearby possible worlds in which one believes that p and the relevant initial conditions for the formation of the belief are the same as in the actual world, while one's belief that p is not true.[15]

Classical invariantists who endorse (NLK*) can typically[16] allow that one's belief that a given ticket in a fair lottery won't win qualifies as knowledge. After all, provided that the chances that such a ticket will win are small, that means that there will only be a small fraction of nearby worlds at which a given ticket wins. If there is only a small such fraction, however,

then there is no wide range of such worlds. Accordingly one's belief that a given ticket in a fair lottery won't win satisfies (NLK*). In this way, endorsing (NLK*) may allow classical invariantists to hold that we can know that a given ticket in a fair lottery won't win. At the same time, classical invariantists who are attracted by the present escape route may still appeal to the anti-luck platitude in order to explain other problematic cases they would want their theory to explain—such as, for instance, a range of Gettier cases.[17] Of course, there is some work that defenders of this option will have to do. For instance, they need to explain how it is that we have the intuition that we can't know, in advance, the details of the outcomes of fair lotteries. However, it may seem that this is not an impossible task. So, perhaps there is hope for classical invariantists.[18]

Unfortunately, however, this proposal does not stand up to closer inspection. To see why this is so, consider the following case:

*The Lottery **.*	A lottery with ten tickets and one winner has been held but the winner has not yet been announced. I come to believe on the basis of the probabilistic evidence that ticket 1 did not win the lottery. My belief is true. According to the present suggestion, I now know that ticket 1 didn't win. I then come to believe on the basis of the probabilistic evidence against winning that ticket 2 did not win. My belief is true. I know that ticket 2 did not win. I acquire further true beliefs that tickets 3 to 9 did not win. Each of my beliefs qualifies as knowledge. I now competently deduce and on that basis come to believe that neither ticket 1 nor ticket 2 nor ticket 3 … nor ticket 9 won the lottery. Since knowledge is closed under competent deduction, I thereby come to know that what I believe. I now competently deduce and on that basis come to believe that ticket 10 won the lottery. Since knowledge is closed under competent deduction, this belief of mine also qualifies as knowledge.

Now, first of all, notice just how implausible this result is. For now it is not only the case that I can come to know, on the basis of merely probabilistic grounds, that a given ticket in a fair lottery will lose, I can also come know, even in the absence of further grounds such as hearing the announcement of the winner, which ticket will *win* the lottery. *That*, however, appears to be massively implausible. At second glance, it is not clear that the defender of this option can even bite this massive bullet and maintain that I do know that ticket 10 won the lottery. After

all, it would seem that my belief, based on competent deduction, that ticket 10 won the lottery, although true, is only luckily true. For it might easily have been the case that a different ticket won the lottery in which case my belief, based on competent deduction, would have been false. The point can be brought out nicely by consulting the predictions of the anti-luck condition on knowledge as captured in (NLK*). At most nearby worlds a different ticket wins the lottery. So, at most nearby worlds at which I still believe that ticket 10 won the lottery and the relevant initial conditions for the formation of my belief are held fixed—which include, among other things, that I still form a belief that ticket 10 won on the basis of competent deduction—my belief that ticket 10 won is false. So, given that knowledge requires non-lucky true belief in the way envisaged by a defender of this escape route, by his own lights, my belief that ticket 10 won the lottery does not qualify as knowledge. Great, the defender of this escape route might want to reply. So, there is a way of holding on to the counterintuitive but at least not massively implausible idea that one can know on the basis of merely probabilistic grounds that a given ticket in a fair lottery did not win, whilst being able to resist the massively implausible idea that one can come know, even in the absence of further grounds such as hearing the announcement of the winner, which ticket won the lottery. In fact, however, things are not as rosy as they may appear to the defender of the present escape route. For what these considerations effectively show is that, on the weakened anti-luck platitude, knowledge is not closed under competent deduction. One can know a set of propositions, $P_1 \ldots P_n$, competently deduce another proposition, Q, from $P_1 \ldots P_n$, and on that basis come to believe Q, whilst not knowing Q. In this way, the counterintuitive concessions the defender of this solution to the fallibilist puzzle has to make pile up. The question then arises as to how much classical invariantists of this breed can swallow whilst still laying claim to advancing a plausible theory of knowledge.[19],[20]

Let us turn to the second option, then. Might the classical invariantist be able to avoid the charge of contradiction by maintaining, contrary to intuition, that in *The Lottery** I know

that the computer-breakdown ticket won't win but I don't know of any of the numbered tickets that they won't win? This is the option they will want to choose if they want to hang on to the anti-luck platitude as it was initially conceived and safety-condition on knowledge captured in (K). Now, one might think that this option looks a lot better than the first option in that the extent of counterintuitive results is much lesser: Rather than having to say, contrary to intuition, that for each losing ticket in a fair lottery we can know that it won't win, we now only have to admit that there are some tickets in some (fair) lotteries of which we can know, in advance, that they won't win. Moreover, this option appears to be able to preserve closure under competent deduction: If in order to know one has to have *eliminated* risk of error, there is reason to believe that competent deduction will preserve elimination of risk of error with respect to propositions entailed by the proposition known. Of course, a defender of this option will also have work to do. He will have to explain to us how it is that we have the intuition that I don't know that the computer-breakdown ticket won't win. Again, however, that may not be an insuperable difficulty.

Unfortunately, by going down this escape route the classical invariantist also incurs significant further intuitive costs. First, there is reason to believe that by the lights of the present escape route, some very plausible relations between knowledge and assertion, viz. that knowledge warrants assertion, will break down.[21] As I have already pointed out, since the number of tickets may be indefinitely large, the probability that the computer-breakdown ticket will win may be as high as the probability that any of the other tickets will win. Now suppose that the number of tickets is so large that the probability that the computer-breakdown ticket wins is just as high as the probability that any other ticket wins—even though the computer is highly reliable etc. Suppose I truly believe on the basis of the probabilistic evidence against winning that the computer-breakdown ticket won't win. According to the present proposal, my belief qualifies as knowledge. Suppose, furthermore, that I am also informed about the further particulars of the lottery. In particular, I know that

the probability that the computer-breakdown ticket will win is just as high as the probability that any other ticket will win. In this case, by the lights of the present classical invariantist proposal, I am in a position to warrantedly assert: "It is just as likely that the computer breakdown ticket will win as it is likely that any other ticket will win, the lottery is fair, but of all tickets the computer-breakdown ticket won't win." However, it would seem that I could have no warrant for such an assertion. For if I have warrant for the first two conjuncts, I won't have warrant to assert the third conjunct. So, the envisaged relation between knowledge and assertion breaks down.

Moreover, on the premises of the present escape route, highly plausible relations between knowledge and rational action are bound to break down. Suppose there is a number of courses of action, $a_1 \ldots a_n$, I can take in a given situation. Suppose that for each course of action, $a_1 \ldots a_n$, there is a proposition, $P_1 \ldots P_n$, such that a given action, a_i, will be successful only if the allocated proposition, P_i, is true. Now suppose that there is some proposition, say, P_1, that I know to be false. At the same time, I do not know of any of the other propositions, $P_2 \ldots P_n$, whether they are true or false. It seems overwhelmingly plausible that in this situation, whatever I do, if I want to act successfully, I ought not opt for course of action a_1. It will be irrational for me to opt for a_1. After all, a_1 will be successful only if P_1 is true and I know that P_1 is false. I can rule out that a_1 will be successful. At the same time, since I don't know whether any of the other allocated propositions are true or false, I can't rule out (or rule in) the success of any of the other courses of action. Since if one is presented with a number of possible courses of action of which one can rule out one and only one as unsuccessful, then it is certainly irrational to opt for the one course of action one can rule out, it will be irrational for me to opt for a_1.[22]

Now suppose a friend of mine wants to give me a ticket in a computer-based lottery with a large number of numbered tickets and a computer-breakdown ticket. Suppose the number of tickets is so large that it is as probable that the computer-breakdown ticket will win

as it is probable that any other ticket will win (although the computer is highly reliable etc.). Suppose I believe, truly, on the basis of the relevant probabilistic evidence that the computer-breakdown ticket won't win. According to the present account, my belief qualifies as knowledge. I also believe of each of the other tickets that they won't win. According to the present account, I don't know of any of the other tickets that they won't win. Suppose now that my friend has offered me a ticket, I also want to win the lottery. Suppose, finally, my friend lets me choose my own ticket. What should I do? What is the rational thing for me to do? Intuitively, it does not matter. Since the chance of winning is equal for all tickets, it would seem that it whichever ticket I pick my choice will be equally rational/irrational. Since, if I want to win the lottery, it will not be irrational for me to make a choice, that means that it won't be irrational for me to choose the computer-breakdown ticket either. At the same time, given the highly plausible relations between knowledge and rational action outlined above, by the lights of a defender of the present escape route, it is *irrational* for me to pick the computer-breakdown ticket. After all, the only thing I know about the outcome of the lottery is that the computer-breakdown ticket will lose. The only course of action I can rule out as unsuccessful is the one of choosing the computer-breakdown ticket.[23] In this way the highly plausible relations between knowledge and rational action break down. Again, we find that the costs of the present escape route pile up. Accordingly, I take it that it is not a very attractive option either.

Let me then turn to the third option. The classical invariantist may maintain that, contrary to what initially seemed plausible, we do not know propositions for which we have only inductive evidence because our beliefs are too luckily true to qualify as knowledge. For instance, in *The Gift*, I don't know that there is wine in the bottle because my belief to that effect, although true, is too luckily true to qualify as knowledge. Or, put in terms of the safety-condition on knowledge captured in (K), defenders of the present escape may maintain that, contrary to first appearances, there is a minority of nearby possible worlds at which there isn't

any wine in the bottle. Since at those worlds I continue to believe that there is wine in the bottle my belief does not satisfy (K). Accordingly, it does not qualify as knowledge. The advantages of this escape route vis-à-vis the first two options is that there is reason to believe that a defender of this escape route can hold on to closure under competent deduction and rescue the relations between knowledge and assertion and knowledge and rational action. Provided that knowledge is true belief that is not *at all* luckily true, there is reason to believe that the defender of the present escape route can hold on to closure—for the same reasons as we had reason to believe the defender of the second option to be able to do so. At the same time, by the lights of this escape route, in the above lottery case I do not know that the lottery-breakdown ticket won't win—and I do not know of any other ticket that it won't win either. Accordingly, it would seem as if defenders of the present proposal can avoid the problems that beset the second proposal. The task that defenders of this option face is to explain how it is that we have come by the intuition that we can have knowledge on the basis of merely inductive grounds.

The problem with this proposal is that if my belief that there is wine in the bottle and related beliefs are too luckily true to qualify as knowledge/do not satisfy the safety condition on knowledge it is hard to see how one could avoid radical scepticism or something very close to it. We have seen that if my belief that there is wine in the bottle does not satisfy the safety-condition on knowledge captured in (K), then there is a nearby world at which there isn't any wine in the bottle. And if, as we have found plausible, the safety-condition (K) captures the part of the anti-luck platitude that solves the fallibilist puzzle, that means that what makes my belief too luckily true to qualify as knowledge is that there is a nearby world at which my belief that there is wine in the bottle is false, that is, there is a nearby world at which there isn't any wine in the bottle, while I continue to believe that there is. Now if a world that is so dissimilar from actuality as a world in which there isn't any wine in an unopened bottle labelled 'Vino' which also has the general appearance of a bottle of wine, has

been bought in an excellent wine store etc. still counts as nearby, then it is hard to see how the following scenarios, which obtain only at worlds that are quite dissimilar to actuality, could fail obtain at nearby worlds, too:

(1) a scenario in which the object I am looking at isn't a sheep but a look-alike stone,

(2) a scenario in which the objects in the bowl aren't fruit but pieces of wax made to look like fruit, and

(3) a scenario in which my watch has stopped working five minutes ago.

If those scenarios obtain at nearby worlds, however, then my beliefs that there is a sheep over there, that there are some pieces of fruit in the bowl, and that it's half-past two do not satisfy (K)/are too luckily true to qualify as knowledge. After all, in all of these scenarios the method of belief formation that supplies me with a true belief in the actual world will give me a false belief. So, if the worlds at which these scenarios obtain are nearby, my beliefs in the actual world will not satisfy (K)/are too luckily true to qualify as knowledge. Now the worry arises that, for virtually any belief in an ordinary everyday proposition, with some ingenuity one could come up with a scenario that is like a scenario in which there isn't any wine in the bottle in *The Gift* and the scenarios described in (1)-(3) in that it obtains only at worlds that are rather (but not extremely) dissimilar to the actual world, while it is still similar enough to actuality that one uses the same method of belief-formation. If with some ingenuity one can come up with such scenarios, then if we want to save classical invariantism from running into paradox by maintaining that beliefs in cases like *The Gift* are too luckily true to qualify as knowledge/do not satisfy the safety-condition captured in (K), it would seem that we will be forced to accept that virtually none of our beliefs in ordinary empirical propositions exclude luck in the way requisite for knowledge/satisfy (K). That is, the worry arises that with a bit of ingenuity we could show that a defender of this escape route is committed to radical scepticism or something very similar to it. Again, it turns out that on closer inspections the

consequences that ensue for defenders of the present escape route are too unattractive to make the proposal a live option.

What we have seen so far is that the fallibilist puzzle threatens to establish that the most natural and promising classical invariantist solution to it runs into a flat-out contradiction. At the same time, classical invariantists cannot take the fallibilist puzzle to suggest that he will have to bite the bullet and accept some apparently counterintuitive consequences of his view and explain them away by some other means. The further intuitive costs that he would incur are too significant to make any such option viable. That means that the classical invariantist must accept that the most natural and promising solution isn't the right solution after all. And, in order to convince us that he is offering a viable theory of knowledge he owes us an alternative solution to the fallibilist puzzle.

Although I certainly won't be able to consider every possible alternative solution to the fallibilist puzzle, I will finish by looking at some of the routes the classical invariantist may want to pursue in this vein. Recall that what the alternative solution will have to do is explain how my belief in *The Gift* differs from my beliefs in *The Lottery** (and in *The Lottery*) such that the former but not the latter qualifies as knowledge. Since my evidential situation, that is, the facts that lead me to my respective beliefs such as my background knowledge and other facts that are on reflection accessible to me, appears to be of the same kind in *The Gift* and *The Lottery**, it would seem that placing an additional internalist justification condition on knowledge won't solve the problem. On the other hand, since, in addition, the facts that ensure the reliability of my respective beliefs are of the same kind and both beliefs can enjoy an equal degree of reliability, it would seem that the prospects for an additional externalist justification condition, probabilistically construed, to solve the problem are dim, too. Furthermore, my belief in *The Lottery** may not be irrational, and, given the massive odds against winning, it may well be rational. So, placing a rationality condition on knowledge (or a non-irrationality condition for that matter) appears not to help. Finally, classical invariantists

might be tempted by a different modal condition on knowledge. For instance, they may be tempted by a sensitivity-condition on knowledge. The sensitivity condition takes roughly the following shape: One knows that *p* only if one's belief that *p* is sensitive. One's belief that *p* is sensitive if and only if were *P* to be false one would no longer believe that *p*. Or, given a standard possible world semantics of counterfactuals, one believes a proposition, *P*, sensitively if and only if at the closest world at which *P* is false, one does not believe that *p*. True, there is a wide range of cases in which appeal to the sensitivity will enable the classical invariantist to explain the difference between beliefs based on inductive grounds that intuitively qualify as knowledge and beliefs in lottery propositions that, intuitively, don't qualify as knowledge. Consider, for instance, a situation in which I am given a bottle and come to believe, on the basis of the same grounds as in *The Gift*, that I am holding a bottle of wine in my hands. My belief is sensitive. At the closest world at which I am not holding a bottle of wine in my hands, I will hold something else in my hands (such as a bunch of flowers or a box of chocolates). At that world, however, I won't believe that I am holding a bottle of wine in my hands. As opposed to that, my belief in *The Lottery** is not sensitive. After all, at the closest world at which the computer-breakdown ticket wins, I still believe that it doesn't. However, the sensitivity condition does not always succeed in separating cases in which belief based on inductive grounds qualifies as knowledge and cases in which it does not in accordance with our intuitions. For instance, in *The Gift*, the sensitivity condition will predict, contrary to intuition, that my belief that there is wine in the bottle does not qualify as knowledge. After all, given how the case is described, at the closest worlds at which there is no wine in the bottle, there is some other liquid in it.[24] At those worlds I continue to believe, falsely, that there is wine in the bottle. So, my belief is not sensitive. Moreover, there are further intuitive costs that sensitivity-based theories of knowledge incur that are as significant as they are well known. Most prominently, there is excellent reason to believe that sensitivity-based conceptions of knowledge won't be able to preserve closure under competent deduction

[cf. Hawthorne 2005] and there are also problems with knowledge of propositions about the future [cf. Hawthorne 2004: 45]. Since the intuitive costs for sensitivity-based conceptions of knowledge are significant as well, it does not appear to be an attractive option for the classical invariantist either.

5. Concluding Remarks

In conclusion, then, appeal to the anti-luck platitude or a safety-condition on knowledge construed along the lines of (K) won't solve the fallibilist puzzle. The intuitive costs that such solutions incur are too high to allow them to remain plausible. Accordingly, classical invariantists will have to look for an alternative solution to the fallibilist puzzle. However, there is at the very least no obvious viable such solution in sight. Placing an additional internalist-, probabilistically construed externalist-, rationality-, or non-irrationality condition on knowledge just won't make any progress towards solving the puzzle. At the same time, the intuitive costs of a solution that appeals to a sensitivity condition on knowledge are just as high as the intuitive costs of some of the solutions that appeal to the anti-luck platitude/safety-condition on knowledge. So long as no adequate solution is on the table, however, it is not even clear that classical invariantism provides a viable theory of knowledge.

References

Cohen, Stewart 1988. How to be a Fallibilist in *Philosophical Perspectives*, 2, ed. J. Tomberlin, Oxford: Blackwell: 91-123.

—— 1999. Contextualism, Skepticism, and the Structure of Reasons in *Philosophical Perspectives*, 13, ed. J. Tomberlin, Oxford: Blackwell: 57-90.

Dancy, Jonathan 1985. *Introduction to Contemporary Epistemology*, Oxford: Blackwell.

DeRose, Keith 1995. Solving the Skeptical Problem, *Philosophical Review* 104: 1-52.

—— 1999. Contextualism: An Explanation and Defence in *Epistemology*, eds. J. Greco and E. Sosa, Oxford: Blackwell.

Gettier, Edmund 1963. Is Justified True Belief Knowledge?, *Analysis* 23: 121-23.

Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: OUP.

—— 2005. The Case for Closure in *Contemporary Debates in Epistemology*, eds. M. Steup and E. Sosa, Oxford: Blackwell.

Lewis, David 1996. Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67.

MacFarlane, John 2005. The Assessment Sensitivity of Knowledge-Attributions in *The Oxford Guide to Epistemology*, eds. T. Gendler-Szabó and J. Hawthorne, Oxford: OUP.

Morillo, Carolyn 1984. Epistemic Luck, Naturalistic Epistemology, and the Ecology of Knowledge, *Philosophical Studies* 46: 109-29.

Richard, Mark 2004. Contextualism and Relativism, *Philosophical Studies* 119: 215-41.

Pritchard, Duncan 2004a. The Psychology and Philosophy of Luck, *New Ideas in Psychology* 22: 1-28.

—— 2004b. Epistemic Luck, *Journal of Philosophical Research* 29: 193-222.

—— 2005. *Epistemic Luck*, Oxford: OUP.

—— forthcoming. Anti-Luck Epistemology, *Synthese*.

Schaffer, Jonathan 2004. From Contextualism to Contrastivism, *Philosophical Studies* 119: 73-103.

—— 2005. Contrastive Knowledge in *The Oxford Guide to Epistemology*, eds. T. Gendler-Szabó, and J. Hawthorne, Oxford: OUP.

Slote, Michael 1979. Assertion and Belief in *Papers on Language and Logic*, ed. J. Dancy, Keele: Keele University Press.

Sosa, Ernest 1999. How to Defeat Opposition to Moore in *Philosophical Perspectives*, 13, ed.

J. Tomberlin, Oxford: Blackwell: 141-54.

—— 2002. Tracking, Competence, and Knowledge in *The Oxford Handbook of Epistemology*, ed. P.K. Moser, Oxford: OUP.

Stanley, Jason 2005. *Knowledge and Practical Interest*, Oxford: OUP.

Unger, Peter 1968. An Analysis of Factual Knowledge, *Journal of Philosophy* 65: 157-70.

—— 1975. *Ignorance: A Case for Scepticism*, Oxford: Clarendon.

Williamson, Timothy 1996. Knowing and Asserting, *Philosophical Review* 105: 489-523.

—— 2000. *Knowledge and Its Limits*, Oxford: OUP.

Notes

[1] Notice that the definition of fallibilism is supposed to be *stipulative* and must not be confused with the view, also often labelled 'fallibilism', that it is not the case that in order to know one must be able to eliminate *every* error possibility. Whether the two types of fallibilism turn out to be equivalent will depend, ultimately, on how the notion of *eliminating an error possibility* is construed. Since, notoriously, there is a variety of ways in which this notion has been construed, I urge the reader not to run the two views together and stick to the stipulative definition.

[2] The main defences of contextualism can be found in [Cohen 1988, 1999], [DeRose 1995, 1999], and [Lewis 1996].

[3] The main defences of this type of view are [Hawthorne 2004] and [Stanley 2005].

[4] Contrastivist conception of knowledge have been proposed e.g. in [Schaffer 2004, 2005].

[5] For two relativist proposals see [Richard 2004] and [MacFarlane 2005].

[6] The first one to develop a conception of knowledge that places something similar to the anti-luck platitude at the centre of attention is, of course, Peter Unger in [Unger 1968]. For further useful discussion of the anti-luck platitude see [Pritchard 2004b, 2005].

[7] Similar versions of the safety principle can be found e.g. in [Sosa 1999] and [Pritchard 2005].

[8] The safety principle is usually stated as a biconditional. For instance, Pritchard construes the safety principle along the following lines:

(SP*)   *S*'s belief is safe *iff* in nearly all (if not all) near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true. [Pritchard forthcoming: 7]

The problem with Pritchard's biconditional version of the safety principle is that it threatens to be to restrictive. Consider, for instance, a case in which I form a belief that there is a goldfinch over there on the basis of identifying a set of features—F_1, F_2, and F_3—which I take to be characteristic of goldfinches. In fact, however, they are not characteristic of goldfinches.

25

They are shared by finches of type X. Now suppose that in the environment in which I grew up there are lots of finches. Suppose there is a finch of type Y which I identify on the basis of a set of features—F_4, F_5, and F_6—which I take to be characteristic of finches of type Y. In fact, however, they are not characteristic of finches of type Y. For they are also shared by goldfinches. Now suppose that in my present environment the only type of finch that is around is the goldfinch. Moreover, there could not easily be any other kind of finch around. If so, at no nearby world there is any other kind of finch over there. If so, then my belief that there is a goldfinch over there acquired via identifying features F_1, F_2, and F_3 is safe. There is no nearby world at which I form a belief about the proposition that there is a goldfinch over there in this way at which my belief turns out to be false. After all, in my present environment there aren't any finches of type X (which share those features and would lead me to form false beliefs). However, it would seem that I acquire false beliefs at some nearby worlds. At some nearby worlds I identify features F_4, F_5, and F_6 and on that basis come to believe, falsely, that there is a finch of type Y over there. Intuitively, my belief that there is a goldfinch over there does not qualify as knowledge. After all, in a salient sense, it is not safe from error. Accordingly, we would expect the safety-principle to explain this intuition. However, on the present biconditional construal, the safety principle is unable to do so. Since the weaker conditional construal will suffice for present purposes, I shall stick to it.

[9] Notice that safe belief entails true belief. Thus the truth condition on knowledge is entailed by the safety condition.

[10] Safety conditions on knowledge have been defended by Sosa—e.g. [Sosa 1999, 2002]—, Pritchard—e.g. [Pritchard 2005, forthcoming]—, and Williamson [Williamson 2000].

[11] Notice, however, that my ambitions are much more moderate than Pritchard's. For reasons given above (fn. 8) I remain unconvinced that Pritchard's safety principle is plausible. Since in the problematic case, I take it that the anti-luck platitude will predict ignorance, it also follows that Pritchard's safety principle does not capture the sense in which knowledge excludes luck. Although it may be that some version of the safety principle captures the sense in which knowledge excludes luck, I do not want to get into these issues here. So I shall remain content with the weaker claim that the safety principle as characterised in (K) captures part of the sense in which knowledge excludes luck.

[12] In the source from which I am quoting Pritchard actually uses the term 'nearest' rather than 'nearby'. However, in a later paper, Pritchard changes and uses the term 'nearby' at this junction. I have chosen the more recent terminology because I take it that it reflects Pritchard's current beliefs more accurately. I also believe that it reflects the truth more accurately. However, I shall not argue for what I believe here.

[13] Notice that Pritchard provides an alternative motivation for strengthening (K) in this way: He points out that given the "substantive cognitive achievement" that knowledge is, in order to qualify as knowledge one's true belief need not only be non-luckily true. Rather, every trace of luck must be eliminated.

[14] Another reason why one must not expect there to be an argument from the fact that the computer-breakdown ticket is most likely to win to the conclusion that now the computer breakdown occurs at some nearby world is that although the computer breakdown ticket is the most likely to win vis-à-vis each of the other possible tickets taken individually, it may not be most likely that the computer-breakdown ticket will win vis-à-vis the possibility of some other ticket winning the lottery. That is, it may still be much more probable that *some* ticket will be drawn than that the computer breaks down. These considerations also substantiate the idea that the closest worlds at which the computer breaks down are far-off. They are very different than the ones in which some ticket will be drawn.

[15] For a proposal along these lines see [Pritchard 2004b].

[16] Notice that there may be some weighted lotteries at which he can't allow this.

[17] Gettier cases have initially been advanced by Edmund Gettier in order to refute the then received view that knowledge is justified true belief [cf. Gettier 1963]. It is now widely believed that the reason Gettier cases work is that a justified true belief may still be too luckily true to qualify as knowledge [cf. Dancy 1985: 134] and [Pritchard 2005: 148]. Accordingly, one would expect Gettier cases to be dealt with by the anti-luck platitude. Given that (K) is still a development of a version of that platitude, there is still reason to believe that defenders of the present proposal can deal with Gettier cases. For further discussion see [Pritchard 2004b].

[18] In [Pritchard 2004b: 219, n. 23], Pritchard deals with lottery propositions in this way. Also Carolyn Morillo maintains that we can know in advance that a given ticket in a fair lottery won't win [cf. Morillo 1984: 119-23].

[19] For further discussion of the closure principle see [Hawthorne 2004, 2005].

[20] It is noteworthy that having to deny closure is particularly problematic for defenders of safety. After all, preserving closure is an integral part of safety-based conceptions of knowledge. This is indicated by the fact that safety-based conceptions of knowledge usually come in a pack with a neo-Moorean solution to the problem of scepticism, a solution, that is, according to which we can know the denials of sceptical hypotheses by competently deducing them from known ordinary empirical propositions. If closure under competent deduction has to go, it is no longer clear that defenders of safety-based conceptions of knowledge can motivate the neo-Moorean response to scepticism. In this way, if defenders of safety-based conceptions of knowledge have to ditch closure, it would seem they will also have to make major readjustments in the general direction of their theory of knowledge.

[21] For defences of this relation see [Unger 1975], [Slote 1979], and [Williamson 1996, 2000].

[22] One might object that the these relations aren't as plausible as they might seem. If I *believe* (although fail to know) of each (though not of all) of the other propositions that they are false, and I believe them to be false with the same or greater conviction than I believe P_1 to be false, then it might not be irrational for me to opt for a course of action that will succeed only if P_1 is true. Even if the objection is correct, surely, the following will certainly be a case of irrational action: I *rationally believe* that I *know* that P_1 is false, and although I *believe* of each proposition, $P_2 \dots P_n$, that it is false, I also *rationally believe* (I even *know*) that I do not and, given my evidence, cannot *know* any of $P_2 \dots P_n$ to be false. Given what I rationally believe, then, I can rule out P_1: P_1 is not an open possibility any more. At the same time, I cannot rule out any of $P_2 \dots P_n$: they are still open possibilities. Now, in this situation, if I want to succeed in my action, it is surely irrational for me to choose a course of action that will be successful only if the only possibility obtains which is not an open possibility for me, viz. P_1. I will explain how the case presented below can be modified to fit these considerations concerning the relations between rational belief, knowledge, and rational action in footnote 18.

[23] The case can also be construed in such a way that I am rationally committed to believing that I know that the computer-breakdown ticket won't win and I am also rationally committed to believing (I even know) that I don't know of any of the other tickets that they will win. All that has to be done is to suppose that I am a defender of the present escape route. Now, if I refuse to follow through my rational commitments I will fall foul of irrationality. The present escape route can avoid the problem only on pain of irrationality. If, on the other hand, I do follow them through, I will rationally believe that I know that the computer-breakdown ticket won't win and I will also rationally believe (I even know) that I do not, and given my evidence, cannot know of any of the number tickets that they won't win. If so, since, by supposition, I want to win the lottery, the relations between knowledge, rational belief, and

rational action outlined above predict that it is irrational for me to choose the computer-breakdown ticket—contrary to what intuition dictates.

[24] Notice that the case is described such that I see the liquid in the bottle. This fact about the description of the case will ensure that the closest world at which there isn't any wine in the bottle, there is some other liquid in it.