



Philosophical Review

The Lottery Paradox, Knowledge, and Rationality

Author(s): Dana K. Nelkin

Source: *The Philosophical Review*, Vol. 109, No. 3 (Jul., 2000), pp. 373-409

Published by: [Duke University Press](#) on behalf of [Philosophical Review](#)

Stable URL: <http://www.jstor.org/stable/2693695>

Accessed: 21/10/2013 10:31

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

The Lottery Paradox, Knowledge, and Rationality

Dana K. Nelkin

1. Introduction

Jim buys a ticket in a million-ticket lottery. He knows it is a fair lottery, but, given the odds, he believes he will lose. When the winning ticket is chosen, it is not his. Did he know his ticket would lose? It seems that he did not. After all, if he *knew* his ticket would lose, why would he have bought it? Further, if he knew his ticket would lose, then, given that his ticket is no different in its chances of winning from any other ticket, it seems that by parity of reasoning he should also know that every other ticket would lose. But of course he doesn't know that; in fact, he knows that *not* every ticket will lose.

On the other hand, if Jim didn't know his ticket would lose, then can he know any empirical facts at all? If Jim does not know something that has an extremely high probability of being true (.9999) and is in fact true then what can he know?

This, in short, is one version of the lottery paradox, a version I call the "knowledge version."¹ The proposition that Jim knew that his ticket would lose has undesirable consequences, but so does the contrary proposition that Jim did not know that his ticket would lose. The paradox can be presented more formally as follows:

The Knowledge Version of the Lottery Paradox

- (1) Jim knows that his ticket (t1) will lose.
- (2) If Jim knows that his ticket (t1) will lose, then he knows

I would like to thank the participants in my graduate seminar on rationality given at Florida State University in the spring of 1998 for very helpful discussion when I first presented some of my ideas on the lottery paradox. I am also grateful to Tyler Burge, Simon Evnine, and Darryl Jung for stimulating discussions of related issues concerning rationality, and to Pat Matthews and an anonymous referee for the *Philosophical Review* for their very useful comments on previous drafts of this paper. Finally, I am indebted to Sam Rickless for his incisive questions, valuable comments, and encouragement throughout my work on this paper.

¹See Kyburg 1961, 197–99, for the original presentation of the paradox.

that t_2 will lose, he knows that t_3 will lose . . . and he knows that $t_1,000,000$ will lose.

So,

(3) Jim knows that t_1 will lose . . . and Jim knows that $t_1,000,000$ will lose. (1, 2)

(4) Jim knows that either t_1 will not lose or t_2 will not lose . . . or $t_1,000,000$ will not lose.

(5) Propositions of the following form comprise an inconsistent set: (a) p_1 . . . (n) p_n , (n+1) not p_1 or . . . not p_n .

So,

(6) Jim knows propositions that form an inconsistent set. (3, 4, 5)

(7) It is not possible to know propositions that form an inconsistent set.

So,

(8) (1), (2), (4), (5), or (7) is false.²

(5) states a logical truth. (7) follows from two extremely plausible assumptions: knowledge implies truth, and inconsistent things can't both be true. This leaves (1), (2), and (4) as candidates for the premise that should be rejected.

It is tempting to many to assume that (1) is false (and I count myself in this group). Jim simply does not know that his ticket will lose. It follows that there must be an answer to the skeptical question concerning empirical knowledge. However, two obstacles remain in the way of offering a satisfying solution to the paradox along these lines. The first, naturally, is the small task of producing an adequate answer to the skeptical question. The second is that there is a second version of the paradox, and it would be desirable if both versions could be solved in tandem.

I call the second version of the paradox the "rationality" version. It goes like this: Although Jim might not know that his ticket will lose, surely it is rational for him to believe that it will. But again, if it is rational for him to believe that his ticket will lose, then it would seem equally rational for him to believe that every other ticket will lose, too. But surely it is rational for Jim to believe that not every ticket will lose; after all, he knows that it is a fair lottery.

²I am assuming throughout that each premise is indexed to a single time.

It seems to follow that it is rational for Jim to hold inconsistent beliefs, and, in particular, beliefs Jim can easily recognize to be inconsistent. But it is not rational for Jim to hold recognized inconsistent beliefs. Contradiction.

More formally,

The Rationality Version of the Lottery Paradox

(1*) It is rational for Jim to believe that his ticket (t1) will lose.

(2*) If it is rational for Jim to believe that t1 will lose, then it is rational for Jim to believe that t2 will lose, it is rational for Jim to believe that t3 will lose . . . and it is rational for Jim to believe that t1,000,000 will lose.

So,

(3*) It is rational for Jim to believe that t1 will lose . . . and it is rational for Jim to believe that t1,000,000 will lose. (1*, 2*)

(4*) It is rational for Jim to believe that either t1 will not lose or t2 will not lose . . . or t1,000,000 will not lose.

(5*) Propositions of the following form comprise an inconsistent set: (a) p1 . . . (n) pn, (n+1) not p1 or . . . not pn.

(6*) Jim recognizes that the following propositions comprise an inconsistent set: (i) t1 will lose . . . (n) t1,000,000 will lose, either t1 will not lose . . . or t1,000,000 will not lose.

So,

(7*) It is rational for Jim to believe inconsistent things that he recognizes are inconsistent. (3*, 4*, 5*, 6*)

(8*) It cannot be rational to believe inconsistent things that one recognizes are inconsistent.

So,

(9*) (1*), (2*), (4*), (5*), (6*), or (8*) is false.

Here, too, we face a choice about which premise to give up.

It is notable that many who discuss the lottery paradox consider only one version of it. What approach is taken often varies with what version of the paradox is under discussion. For example, many of those who consider only the knowledge version of the paradox see rejecting (1) as the obvious (or almost obvious) choice.³ In contrast, those who consider the rationality version of

³See, for example, DeRose 1996.

the paradox tend to see the rejection of (1*) as either in need of extensive argumentation or as the wrong choice.⁴

As I mentioned earlier, I believe that it is desirable that the solution to each version be seen to fit well with the solution to the other. This does not entail that the two versions must be solved in the same way, but it does require that some explanation be given if they are not. One reason for this is that the two versions of the paradox are almost parallel in form. Second, the knowledge that one's ticket will lose, appealed to in the knowledge version of the paradox, appears to be based on an inference from the belief that the odds of one's winning are one in a million. And it is natural to suppose that there is a close connection between inferential knowledge and rationality. Thus, there are at least two *prima facie* reasons for treating both versions of the paradox in the same way.

In this paper, I will offer solutions to both versions of the paradox. My strategy will be to reject both (1) and (1*). In section 2 of this paper, I will offer some initial reasons for adopting this strategy, and explain why it is superior to its main rival. Taking the approach I do leaves me with the burden of answering the skeptical question concerning empirical knowledge. It also leaves me with a parallel question concerning rational belief: if it is not rational for Jim to believe that his ticket will lose the lottery, then how can there be any empirical beliefs that are rational for him to hold? These questions presuppose a pair of deeper questions: what is it about knowledge and rationality, respectively, that makes (1) and (1*) false? In sections 3 and 4, I examine answers to these questions that have been given by others. Although some of these attempts are on the right track, I believe that they do not fully explain why (1) and (1*) are false and, hence, that a complete rebuttal to the skeptical objections has not yet been given. In sections 5 and 6, I offer and defend my own account.

2. Initial Reasons for Denying Knowledge and Rational Belief in the Lottery Case

While Jim might say "I knew I would lose," after hearing the winner announced, I think many of us would view his words as not

⁴See, for example, Ryan (1996) who considers something like the rationality version. She argues for the rejection of something like (1*). See Kyburg 1961, Foley 1979 and 1993, and Klein 1995 for arguments that (8*) should be rejected rather than (1*).

strictly speaking true. Perhaps what Jim really meant (or should have meant) was that he knew he would almost certainly lose. But had he *known* that he would lose, then he would not have bought the ticket. Intuitively, (1) seems false. And this conclusion is supported by the plausibility of the remaining premises in the knowledge version of the reductio. What should we say about (1*)? Here, in contrast to (1), many find that (1*) seems true, at least on first reading. In section 5, I develop an explanation for the falsity of (1*) that helps defuse the initial appeal of (1*). But in this section, I will offer one reason for pursuing this line in the face of initial resistance. It is simply that the other premises in the reductio are themselves very plausible. If we were to retain (1*), which premise would we reject instead?

(2*) is hard to dispute. If it is rational for Jim to believe that his ticket will lose, it is on the basis of his belief about the probability of his ticket being chosen. But a similar ground is available for Jim's belief that t₂ will lose and for his belief that t₃ will lose, and so on. It would be unacceptably arbitrary to assume that it could be rational for Jim to believe that his ticket will lose, but not rational for Jim to believe that another ticket, similar in all relevant respects, will lose.

Despite the appeal of (2*), Harman (1986) offers an argument against it. He begins by claiming that one can rationally infer that ticket 1 will lose on the grounds that the odds are 999,999 to 1. One can then rationally infer that ticket 2 will lose on the grounds that the odds, relative to one's rational beliefs, are 999,998 to 1. One could continue to infer rationally that each ticket will lose until there is a small enough number left. For example, having inferred that 999,990 plus tickets will lose, one could not rationally infer that 999,999 will lose, because the odds, relative to one's rational beliefs are now only 10 to 1. As Harman puts it, "the order of inference *really* matters here" (71).

This solution won't work. For it seems that *arbitrary* order should not matter to what inferences count as rational. And second, if Harman is correct, then an unwelcome consequence follows. Suppose, with Harman, that Jim can rationally infer that t₁ through t_{999,000}, say, will lose. Jim is also rational in believing that one of t₁ through 1,000,000 will not lose. It would seem to follow that Jim could rationally infer a logical consequence of these beliefs, namely, that one of t_{999,001}–t_{1,000,000} will not lose. But this is strongly

counterintuitive. If Jim were rational in believing that one of those 1,000 tickets will win, then, depending on the order of his inferences, he should either try to get his hands on one of those 1,000 tickets or feel fortunate to be holding one of them already! But there is no reason for Jim to do either of these things. Thus, Harman's argument fails, and (2*) remains secure.⁵

(4*) seems equally difficult to deny. In fact, we can imagine that it is rational for Jim to believe that one ticket will win on the basis of excellent empirical evidence (he supervises the lottery himself, for example). If one is not willing to say that it is rational for Jim to believe that one ticket will win, then it is unclear why one should be so concerned to preserve (1*). In other words, if one is struck by the fact that it seems rational for Jim to believe that his ticket will lose, it would appear at least equally rational for him to believe that some one ticket will win.

(5*) is a logical truth and cannot be the culprit.

(6*) is true by hypothesis, and it is a perfectly intelligible hypothesis, as our own recognition of the relevant inconsistency makes clear. This leaves (8*) as an alternative that might be rejected in place of (1*). Denying (8*) is a strategy that has been adopted by Richard Foley, among others. As I see it, there are two related lines of argument in favor of this strategy. The first is to replace (8*) with a weaker principle governing rational belief (I call it the Foley Principle, or (FP)):

(FP) It cannot be rational to believe a proposition that is internally inconsistent.

For example, it cannot be rational to believe a conjunction, the conjuncts of which cannot all be true. Applied to the lottery case, this means that while it is rational for Jim to believe inconsistent things, it is not rational for him to believe the conjunction of those inconsistent things. In other words, although it is rational for Jim to believe that t_1 will lose, it is rational for Jim to believe that t_2 will lose . . . and it is rational for Jim to believe that one of t_1 or t_2 . . . or $t_1, 000, 000$ will lose, it is not rational for him to believe that each of $t_1, t_2, \dots, t_1, 000, 000$ will lose and that one of t_1, t_2

⁵Ryan (1996) uses a similar argument against those who would reject (8*).

... t1,000,000 will not lose. In effect, this amounts to a denial of what Foley calls the “conjunction rule”:

- (CR) If it is rational for S to believe p at time t and it is rational for S to believe q at t, then it is rational for S to believe p and q at t.

This move softens the picture we might have had of Jim believing an outright contradiction while at the same time being forced to say that he is rational. Our reluctance to reject (8*) might then be mitigated by the knowledge that it can be replaced by the weaker (FP).⁶

This line of argument can be combined with a dialectical strategy that should now be familiar (since I have already used it myself): the other premises in the rationality version of the paradox are all very plausible. And if you are even tempted to reject (1*) instead of (8*), then you appear to be committed to an unacceptable skepticism.

In subsequent sections, I will try to show that this dialectical challenge can be met. But let me briefly address the first line of argument, concerning the replacement of (8*) with (FP). It is unclear that anything important is gained by denying (CR) and replacing (8*) with (FP). After all, one is still committed to the conclusion that it is rational for Jim to have beliefs that he recognizes are inconsistent. There is simply an added qualification: in order for Jim to remain rational, he must refrain from combining these into a single belief. Yet it is counterintuitive that one be rational as long as one is careful not to draw particular logical consequences from one's beliefs.⁷

Foley defends his view in part by explaining that it does not commit him to an “anything goes” prescription for reasoning and deliberation. Denying (CR) does not entail that it is never rational

⁶This is not a line of argument that Foley presents directly, because he sets out the paradox differently; thus, his dialectical emphases are somewhat different (see Foley 1979 and 1993, especially 143–44). According to Foley's rendition of the paradox, rather than reaching the (apparently absurd) conclusion that Jim rationally believes inconsistent things (as represented by (7*)), we reach the (apparently absurd) conclusion that Jim believes an outright contradiction. This paradox would be solved simply by embracing (FP) and rejecting (CR). Foley is explicit about his commitment to rejecting (8*) in Foley 1993, 163–64.

⁷Similar points are made by Ryan (1996) and Evnine (1999).

for one to believe conjunctions of one's beliefs; it simply entails that it is not always rational for one to do so. And the lottery paradox is one of an unusual set of cases in which it is not rational to conjoin one's beliefs.⁸

At this point, it might seem as though a stalemate has been reached. But I believe that we have come again to the crux of the matter. Foley's main reason for his denial of (CR) and (8*) is the lottery paradox (together with the related preface paradox). But if there is a plausible alternative solution, then we will not have to resort to saying that one can be rational just so long as one refrains from drawing logical consequences from one's beliefs. In the remainder of this paper, I will explore an alternative solution to the lottery paradox.⁹

3. The Skeptical Objections and "One False Belief" Accounts

My proposal, then, is to reject (1) and (1*). But this solution, by itself, leaves the skeptical questions unanswered. If (1) and (1*)

⁸Foley also appeals to the preface paradox in defending the rejection of (8*). The preface paradox is the case in which it seems it can be rational for an author to believe each sentence she writes in a book, but also rational for her to believe, as she writes in the preface, that at least one of her claims is false. Here, too, the paradox dissolves if one rejects the thesis that recognized inconsistency is irrational. A full evaluation of the preface paradox is beyond the scope of this paper. See Ryan 1991 for an excellent discussion and a solution that does not require rejecting the thesis captured by (8*).

⁹An alternative solution I do not consider in any detail in this paper is one that rejects the entire framework within which the lottery paradox is constructed. On this view, we should reject the idea of a "rationality of belief" altogether, and replace it with a "rationality of degrees of belief." In other words, on this sort of picture, it does not make sense to say that it is rational to believe *p* simpliciter. Rather, it is appropriate to say that it is rational to believe *p* to degree *X*. On this view, several of the premises of the rationality version of the lottery paradox should be rejected. Throughout this paper, I will assume the basic framework of the lottery paradox and assume that there is a "rationality of belief." I believe that its existence is intuitively plausible, but others have given arguments in favor of it, as well. See, for example, Harman 1986, 21–27. Also see Foley 1993 for a discussion of the pros and cons of rejecting a rationality of belief. In the end, Foley accepts a rationality of belief, although he incorporates elements of a rationality of degrees of belief into his view. In particular, he argues that a belief is rational (simpliciter) if and only if one is rational in having a degree of confidence in the belief that exceeds one's minimum threshold for belief.

are both false, then does it follow that there are no empirical facts we can know, or even rationally believe? And, if so, is that an acceptable consequence?

I agree with opponents of this solution that it would be an unacceptable consequence that we are unable to know or even rationally believe any empirical facts. So the alleged consequence must not really follow from rejecting (1) and (1*). The problem then becomes how to distinguish between cases in which the object of knowledge or rational belief is a proposition like

(L) My lottery ticket will lose

and cases in which the object of knowledge or rational belief is a proposition like

(F) The room I just left still has furniture in it.¹⁰

We usually assume that we know and rationally believe propositions like that expressed by (F). But there might have been a super-efficient furniture burglar who just completed his job in that room, or an automated furniture rental company using the latest in military technology to pick up the furniture at the end of its lease period, or an evil demon trying to deceive me . . .

We cannot be certain of either (F) or (L). What then is the relevant difference between the lottery case and the furniture case? In this section I will explore a strategy for distinguishing between the cases that focuses on the fact that in the lottery case, unlike the furniture case, one knows (or rationally believes) that one of a relevantly similar set of (potential) beliefs is false (henceforth, the “One False Belief” strategy). This special feature of the lottery situation is what distinguishes it from the vast majority of our other empirical beliefs and explains why it is that we cannot have knowledge (or rational belief) that one’s lottery ticket will not win.

The general strategy can be found in both Bonjour 1985 and Ryan 1996. Bonjour is primarily concerned with the knowledge version of the paradox, and Ryan with a variant of the rationality version.

¹⁰See Foley 1987, 245 for this example.

Bonjour writes that what distinguishes the lottery case from others is

the presence of a large number of relevantly similar, alternative possibilities, all individually very unlikely, but such that the person in question *knows* that at least one of them will in fact be realized, but does not know which one. In such a case, since there is no relevant way of distinguishing among these possibilities, the person cannot believe with adequate justification and cannot know that any particular possibility will not be realized, even though the probability that it will not may be made as high as one likes simply by increasing the total number of possibilities. What rules out knowledge in such a case is not merely that the probability of truth relative to the person's justification is less than certainty but also that the person *knows* that one of these highly probable propositions is false (and does not know which one). It is, I submit, this further knowledge, and not merely the lack of certainty, that prevents one from knowing the proposition in question. There are obviously, however, very many cases in which the justification which a person has for a belief fails to make it certain that the belief is true, but in which further knowledge of this sort is not present. (1985, 236 n. 21)

In a similar spirit, Ryan argues that one is not epistemically justified (or, in other words, epistemically rational) in believing that one's lottery ticket will not win. The reason is that "[one] is in a peculiar situation of knowing one of [one's] beliefs is false but not having any idea which belief is false" (1996, 130). In particular, Ryan appeals to a principle she calls the "Avoid Falsity Principle" or AFP:

(AFP) For any set L of competing statements, if (i) a person S has good reason to believe each member of L is true and (ii) either S has good reason to believe at least one member of L is false or S is justified in suspending judgement about whether at least one member of L is false, then S is not epistemically justified in believing any of the competing individual statements of L. (1996, 130)¹¹

By "competing statements" Ryan means "all statements that are individually reasonable but are called into question by the intro-

¹¹This formulation differs in one notable way from the informal statement quoted in the text above. There, the key idea is that one *knows* that one of a set of statements or beliefs is false. In the more formal statement, the key idea is that one is *justified* in believing that one of a set of statements or beliefs is false. Since Ryan takes justification as a necessary condition for knowledge, there is no real conflict here.

duction of reasonable counterevidence.” By “having good reason to believe *p*” Ryan means “one’s total evidence supports *p* more than it supports not-*p*.” Finally, “being justified in suspending judgement about *p*” means that one is not justified in believing *p* and not justified in believing not-*p*.

In simpler terms, then, (AFP) says that if one has a set of beliefs, each of which is better supported by one’s total evidence than its negation, but which are collectively called into question by a rational belief that one of the set is false, then one is not rational in believing any of the set.^{12,13}

(AFP) seems to give the right result in the lottery case. Although it seems that for each ticket, Jim’s evidence better supports the conclusion that it will lose rather than that it will not lose, it is clearly rational for Jim to believe that one ticket will not lose. Thus, according to (AFP), it is not rational for Jim to believe that his ticket will lose or that any other ticket will lose. (AFP) also seems to give the right result in other cases like the furniture case. There is no competing set of statements that includes (F) and that are collectively called into question by a justified belief that one of those statements is false. For example, I don’t have a justified belief that on one day this year, though I know not which, the furniture will disappear from a room shortly after I leave it. Thus, (AFP) allows us to preserve the intuition that I am rational in believing that the furniture is still in the room I just left.

However, there are problems for (AFP). The first is that it ap-

¹²For the sake of simplicity, I am here leaving out the clause that if one is merely justified in suspending judgement about whether one member of the set is false, then one is not epistemically justified in believing any of the set. I do not believe this will affect the subsequent discussion.

¹³It is necessary to make one clarification before turning to an assessment of Bonjour’s and Ryan’s attempts to meet the skeptical challenges. Although Ryan explicitly uses the terms ‘epistemically justified’ and ‘epistemically rational’ interchangeably throughout her paper, it becomes clear toward the end of her paper that she thinks of epistemic justification as the “justification condition” for knowledge, understood as justified true belief (together with a fourth anti-Gettier case condition). Thus, one might argue that Ryan’s version of the paradox is distinct from the rationality version in which ‘rational’ has a sense independent from our concept of knowledge. In what follows, I will treat Ryan’s answer to the skeptical challenges as applying to the rationality version of the paradox, but I believe that what I say could be transposed to Ryan’s version, if it is indeed a distinct version.

pears vulnerable to a certain kind of counterexample. Suppose, to vary an example of Ryan's, I am missing a book that I was reading just before six colleagues, Alice, Bonnie, Carol, Delia, Elizabeth, and Fiona entered my office. I question each of them about it, and all deny having taken it. Suppose that, for each colleague, my total evidence supports the conclusion that she did not take the book over the conclusion that she did. Still, suppose that it is rational for me to believe that one of them took it. In this case, (AFP) says that it is not rational for me to believe that Alice did not take it, or that Bonnie did not take it, and so on. Perhaps this seems like the right result: it is rational for me to suspend judgment. But now let us add to the case. Suppose that I have known Fiona all of my life, I have always known her to be fiercely honest, and I have never known her to do anything absentmindedly. I have known the others for a few weeks, and while they have all been friendly, I do not know much about their personal qualities. (AFP) still says that it is not rational for me to believe that Fiona did not take the book. But this seems counterintuitive.

(AFP) can be modified in a way that takes account of the 'Fiona' case by requiring that each member of the competing set of beliefs be epistemically indistinguishable. In other words, if each member of the competing set is equally well supported by one's total evidence, and that set is called into question by a rational belief that one member of the set is false, then it is not rational to believe any members of the set. More formally, a modified version of (AFP) (call it (AFP')) is the same as the original (AFP) with condition (i) replaced by

- (i') a person S has good *and equal* reason to believe each member of L is true.

(AFP') is immune to counterexamples like the Fiona case.

But a challenge remains even for (AFP'). The Fiona case highlights the fact that there are very few real cases in which (AFP') applies. It is not usually the case that one has a competing set of beliefs each of whose members is equally well supported by one's evidence. Further reflection reveals that (AFP') is so tailored to solve the rationality version of the lottery paradox that it does not by itself do a lot of explanatory work for us. Perhaps this (or something like it) is the best we can do. But a defender of the "reject (8*)" solution to the paradox might reasonably argue that (AFP')

does not go very much further than a restatement of the position that (1*) rather than (8*) should be rejected. For, in essence, (AFP') asserts that it cannot be rational to hold a particular configuration of inconsistent beliefs, namely, a set of equally well supported beliefs and a belief that one of those beliefs is false. But someone who thinks that it can be rational to hold inconsistent beliefs will be unconvinced, and can legitimately press for a further explanation. Thus, we have reason to seek an alternative explanation for why (1*) should be rejected.

Before turning to alternatives, let us assess Bonjour's One False Belief account.¹⁴ According to Bonjour, if one knows that one statement in a competing set of statements is false, then one does not know any of the members of the set. This principle seems to divide the cases in the right way. For example, since Jim knows that one ticket will not lose, but has no relevant way of distinguishing the tickets from each other, he does not know that his ticket will lose. But there is no competing set of statements, similar in all relevant respects, that includes (F) and about which Jim knows that one member is false.

However, it is important to note that the One False Belief strategy applied to the knowledge version seems specially tailored to cover lottery-like cases. Although the burden of proof tends to be considered lighter for rejecting (1) in the knowledge version than for rejecting (1*) in the rationality version, it is reasonable to ask in this case, too, whether there is a deeper explanation for why (1) should be rejected.

In sum, the main problem with the One False Belief accounts is that they are so finely tailored to lottery-like cases that they are limited in their ability to explain what is really responsible for our lack of knowledge or rational belief in those cases.

¹⁴Bonjour's strategy differs from Ryan's in two main respects: First, it is employed in the service of solving the knowledge version of the paradox. And second, what is essential for explaining why (1) should be rejected is that Jim has a particular piece of knowledge, namely, that one of a "competing set" of statements is false. In contrast, what matters in Ryan's official explanation for why (1*) should be rejected is that Jim has another *rational belief*, not a piece of *knowledge*. There is nothing essential to the strategy about pairing rational belief with the rationality version and knowledge with the knowledge version. But I do not believe that changing the pairings makes a significant difference.

4. Knowledge and the Subjunctive Conditionals Account

Keith DeRose, addressing the knowledge version of the lottery paradox, agrees that (1) must be rejected. Or, rather, he agrees that at the very least it *seems* that (1) must be rejected. (The significance of this qualification will become apparent shortly.)

Yet DeRose offers an explanation that is more general than Bonjour's One False Belief account. His idea is this:

According to SCA [the Subjunctive Conditionals Account], the reason we judge that you don't know you've lost the lottery is that (a) although you believe you are a loser, we realize that you would believe this even if it were false (even if you were the winner), and (b) we tend to judge that S doesn't know that P when we think that S would believe that P even if P were false. (1996, 569)

SCA is not so finely tailored to the lottery situation as is the One False Belief account. And in fact, it has independent plausibility, as attested to by its association with well-known analyses of knowledge.¹⁵ At the same time, it seems to give the right result in the lottery case. Finally, it seems to give the right result in everyday cases of empirical belief. Consider the following:

(N) The Bulls beat the Knicks last night.

If you read that the Bulls beat the Knicks in the newspaper this morning, then under ordinary circumstances, we would not hesitate to say that you know that the Bulls beat the Knicks. SCA allows for this: we do not judge that you would believe that the Bulls beat the Knicks even if that were false.

Now DeRose is careful to say that he only means to explain why it *seems* to us that (1) is false, and not why (1) is actually false. And I believe that there is good reason for this. To see why, let us convert SCA to a general explanation of why Jim lacks the knowledge that he will lose. Let SCA' be the view that the reason Jim does not know that his ticket will lose is that (i) if his belief were false, he would still believe it, and (ii) S doesn't know that P when S would believe that P even if P were false. SCA' will not succeed in explaining Jim's lack of knowledge because (ii) is false.¹⁶

¹⁵See, for example, Nozick 1981.

¹⁶DeRose is aware of this; see 568.

First, consider the furniture case. What does SCA' say about (F)? It is plausible that I would believe (F) (that the furniture is still in the room I just left) even if it were not. According to SCA', then, I do not know that the furniture is still in the room I just left. This is counterintuitive.

A second case highlights the general difficulty with counterfactual analyses. Imagine that you are a Knicks fan. Now imagine that your neighbor wants you to suffer. If the newspaper did not report a Bulls victory over the Knicks, your neighbor was prepared to have your paper professionally altered so as to include an authentic-looking report of a Bulls win. Suppose that the Bulls did in fact win, and the newspaper reports the score correctly. Your sadistic neighbor never interferes. Intuitively, it seems that you know that the Bulls won. But you would have believed that the Bulls won even if they did not. So SCA' says that this case is just like the lottery case: you don't know that the Bulls won.¹⁷

These cases show that SCA' fails to distinguish the lottery case from everyday cases of empirical belief in the right way. Nevertheless, I think that SCA' is on the right track. There is something right and relevant about the idea that Jim would believe he had lost the lottery even if he had won. His belief is not connected to the truth in an important way. This insight suggests that we again look deeper—in this case into what *grounds* the subjunctive conditional about Jim. Doing so will allow us to avoid using the subjunctive conditional—with all the controversy and difficulties surrounding its truth conditions—as a precise measuring tool, while instead using it as a rough guide to why we lack knowledge in lottery-type cases. This is an idea to which I will return in the next section.

A final reason for questioning SCA' is that it has no natural correlate for explaining why we lack *rational belief* in lottery-type cases. In other words, it is not clear how it would fit with a solution to the rationality version of the paradox.

5. The Statistical Support Accounts

So far we have seen two attempts to explain Jim's lack of knowledge in the lottery situation. The One False Belief account focuses on

¹⁷Similar counterexamples have been offered to Nozick's counterfactual analysis of knowledge. For example, see Klein 1987.

a special feature of lottery situations: that one knows that one of a range of relevantly similar statements is false. The Subjunctive Conditionals Account and SCA' focus on a feature of lottery situations that is shared by many other kinds of situations: that one would have a belief whether or not it is true. In this section, I would like to focus on a feature of the lottery situation that is special, but not unique: Jim's belief that his ticket will lose is based on his knowledge (or belief) about the statistical probability that his ticket will lose. That is, it is based on an inference of the following form (henceforth, "P-inference"):

- (P) p has a statistical probability of n [where n is a very high number] $\rightarrow p$

In what follows, I will develop an account that takes as its starting point the premise that what prevents Jim from knowing that he will lose the lottery is the existence of a P-inference. I call this the Statistical Support account for the knowledge version of the lottery paradox.

If the fact that Jim's belief that he will lose is based on a P-inference is the ultimate—or intermediate—explanation of why one cannot have knowledge in the lottery case, then it must generalize to non-lottery situations in which we base beliefs on probabilities.

Suppose that I have a computer program that chooses a background color for my computer screen by randomly selecting a number between one and one million each time I turn on my computer. Associated with one of the numbers is a red screen. Associated with 999,999 numbers is a blue screen. Suppose my screen will be blue when I turn on my computer tomorrow. Do I know that my screen will be blue tomorrow? I believe the answer is no. I know that it is highly likely to be blue, but I do not know that it will be blue.

Notice that it is possible that the single "red" number will never be chosen by my computer program. Thus, this case differs from the lottery situation in which one knows that some ticket will not lose. In the lottery situation, there is a set of statements (ticket1 will lose, ticket2 will lose, and so on) of which I know one to be false. A parallel set of statements in the computer case is: tomorrow the screen will be blue, the next day the screen will be blue, the

day after that it will be blue, and so on. But I do not know that one member of this set is false. Thus, we cannot appeal to the One False Belief account in this case. Rather, what accounts for my failure to know in this case is the fact that my belief is based merely on statistical evidence. This suggests that we should avail ourselves of a similar account in the lottery situation: the obstacle to Jim's knowledge in the lottery situation is that his belief that his ticket will lose is based merely on statistical evidence. There is no need to appeal to the further piece of knowledge that one ticket will win in order to explain his lack of knowledge.¹⁸

The thesis that P-inferences cannot yield knowledge has the important virtue of explaining in an independently plausible way why (1) should be rejected. And in doing so, it makes this solution to the paradox that much more plausible. For there are non-lottery cases in which one's belief that *p* is based on a P-inference and in which we would deny that one has knowledge.¹⁹

At the same time, this account allows room for our intuitions about other empirical beliefs. Consider (F). We can know (F) because it is not based on statistical evidence. It is difficult to say what it is based on, but it seems to be based on a variety of perceptual beliefs and on other beliefs that together form a theory of the kinds of things that happen in my environment. Similarly for (N). My belief about the Bulls victory does not depend on a belief about statistical evidence. Even if I know newspapers sometimes make mistakes, my belief is not based on a statistical probability. It is

¹⁸See Cohen 1988 for another view that identifies the statistical evidence in the lottery case as the obstacle to Jim's knowledge. Cohen's view is embedded in a contextualist view of knowledge, according to which whether we know something depends on the context. In certain contexts, it will be necessary to rule out certain "relevant alternatives" to what one believes; in other contexts, different alternatives will be relevant and require ruling out. The fact that Jim makes a P-inference to his belief that he will lose makes his winning a "salient" alternative that he cannot rule out. Therefore, Jim cannot be said to know that he will lose. In what follows, I offer a very different reason for why P-inferences cannot yield knowledge. See also Williamson 1996 for an identification of the statistical nature of the evidence as the obstacle to knowledge in the lottery case. Williamson's main aim in that paper is to defend a connection between assertion and knowledge, and he does not elaborate on this point.

¹⁹As DeRose points out, there are also lottery drawings that do not necessarily have winners. Even in these cases, it seems that one does not know that one will lose, despite not knowing that anyone else will win.

based on my belief about what the newspaper says, its coherence with other things I believe, and perhaps on my belief that the newspaper is reporting correctly. Thus, identifying the probabilistic nature of the evidence in the lottery case as the epistemically relevant feature makes it possible to distinguish the lottery case from other cases of everyday empirical belief.

At this point, the following question naturally arises: why should P-inferences pose an obstacle to knowledge? In answering this question, a return to an insight of the previous section will be useful. There we encountered the idea that Jim's belief that his ticket will lose seems not to be connected with the truth in the right way. And now we can add to this idea: there is a way in which true beliefs based merely on statistical probabilities are not connected to the facts that make them true. In general terms, the fact that makes Jim's belief true is not something that bears a causal (or, more generally, explanatory) connection to his belief.²⁰

In particular, the fact that Jim's ticket will lose does not cause or explain Jim's beliefs. Nor is there some other fact that is causally or explanatorily related in some way to both Jim's belief and the fact that makes it true. There are cases in which we can know things about the world, not because they cause or explain our beliefs, but

²⁰While his paper receives little attention in the literature on the lottery paradox, Alan Goldman (1984) argues for something like this point in the context of introducing an "explanatory" analysis of knowledge. In particular, he argues (i) that a lack of explanatory connection is what is responsible for a lack of knowledge in the lottery case and (ii) that "explanatory connection" should be cashed out in counterfactuals (albeit more elaborate ones than those in SCA'). Roughly, A explains B when $\text{prob } B/A > \text{prob } B$. Probabilities in turn are understood as follows: $\text{prob } X$ is high if in a large proportion of nearby possible worlds, X obtains. In addition to difficulties in assessing the relevant counterfactuals, Goldman's analysis is vulnerable to counterexamples like the "sadistic neighbor" case described in section 4. Since Goldman offers his analysis as providing a set of necessary and sufficient conditions for knowledge, the analysis is also vulnerable to counterexamples from the "sufficiency" direction. (For example, a medium who really does have extrasensory perception, but who has reason to believe that she does not, would count as having knowledge on Goldman's view.) Despite these problems with Goldman's view, I am in full agreement with the fundamental point that an explanatory connection is what is missing in the lottery situation. Perhaps it is because of its contingent association with a counterfactual analysis of knowledge that this fundamental point has not received sufficient attention in the literature on the lottery paradox.

because there is a common cause or explanation, for example. I know that I will stop writing soon because I have formed the intention to do so. My intention explains both why I believe that I will stop and the fact that I will stop. But Jim is not in a situation of this kind. Because his belief is based on a belief about the probability of his losing, and because that belief does not bear a causal or explanatory relation to the fact of his actually losing, there are no causal or explanatory connections between his belief that he will lose and the fact that makes it true.

These considerations point the way to an explanation of Jim's lack of knowledge: a causal or explanatory connection is missing in the lottery situation, and indeed it is missing because Jim's belief is based on mere probabilities. This account dovetails with certain well-known externalist analyses of knowledge, an association that engenders mutual support for both the account and certain forms of externalism.

Externalism, as I will understand it here, is, to borrow the words of Armstrong, the view that "what makes a true . . . belief a case of knowledge is some natural relation which holds between the belief state and the situation which makes the belief true" (1973, 157). To see how externalism and the present account of the knowledge version of the lottery paradox help to support each other, consider one classic form of externalism: Alvin Goldman's causal account. Despite the fact that there are a number of more recent externalist accounts that have been offered as improvements on this one (including Goldman's own sophisticated reliabilist accounts), I think Goldman's causal account is most useful for our purposes. There are three reasons for this: it is simple, we can easily see what motivates its externalist component, and, finally, it is easy to see how it applies to the lottery paradox.

Goldman is concerned to address an alleged counterexample to the traditional analysis of knowledge as justified true belief put forward by Gettier (1963). Very briefly, suppose Smith believes with justification that Jones owns a Ford on the basis of very good evidence (for example, seeing Jones drive a Ford, hearing Jones say he owns a Ford). It seems to follow that Smith is justified in believing that either Jones owns a Ford or Brown is in Barcelona, even if he has no evidence whatsoever concerning Brown's whereabouts. But now suppose that contrary to Smith's evidence, Jones does not own a Ford and, by an amazing coincidence, Brown is in Barcelona.

Even though Jones has a true justified belief, we do not think he knows that either Jones owns a Ford or Brown is in Barcelona. Thus, it appears that a condition must be added to truth, justification, and belief to yield sufficient conditions for knowledge.

Goldman observes that if Smith had received a letter from Brown postmarked in Barcelona, we would say that Smith has knowledge of the disjunction. Alternatively, if Jones really owned a Ford, and his owning it was manifested in his behavior, which in turn caused Smith's belief, then we would say that Smith has knowledge. Generalizing from these observations, Goldman writes, "one thing that seems to be missing in this example is a causal connection between the fact that makes p true and Smith's belief of p ." He goes on, "The requirement of such a *causal connection* is what I wish to add to the traditional analysis" (1967, 145). In further support of his view, Goldman points to a variety of cases of empirical knowledge in which some such causal condition seems to be satisfied. Perceptual knowledge and knowledge based on memory are paradigm cases in which we seem to require a causal connection between what is believed and what makes that belief true.²¹

Thus, this causal version of externalism derives support by its providing a solution to a Gettier problem and by the existence of a causal connection in paradigm cases of knowledge. Now notice that it also has the resources to solve the knowledge version of the lottery paradox in something like the way sketched above: Jim's belief that he will lose is not causally connected in any way to the fact that he will lose. The fact that he will lose does not cause Jim's belief, nor is there any common cause or chain of causes. Thus, Jim does not know that he will lose. This explanation of why Jim does not know that he will lose allows us to solve the knowledge version of the lottery paradox by rejecting (1), while preserving the idea that many cases of everyday belief can constitute knowledge (beliefs based on perception and memory among them).^{22, 23}

²¹See Goldman 1967 (1999, 148–51) for a discussion of beliefs about the future and other kinds of belief.

²²Dretske also argues that his analysis of knowledge, the information-theoretic account, accounts for Jim's lack of knowledge in the lottery situation (1981, 99–102). Very briefly, according to Dretske's account, one knows that s is F if and only if one's belief that s is F is caused (or causally sustained) by the information that s is F . The information that s is F causes a belief when a signal, r , carries the information that s is F in virtue of r 's being G and r 's being G causes the belief. Finally, r 's being G carries the

information that s is F only if the conditional probability of s being F , given that r is G , is 1. On Dretske's account, Jim lacks knowledge because he does not have access to the information that he will lose. That is, he has no signal such that its having a certain property entails that he will lose. This account shares with mine the idea that the belief that constitutes knowledge must be connected with the truth in some way, and that Jim's belief that he will lose fails to be so connected. It differs from my account in what it takes that connection to be. Ultimately, for Dretske, what is essential is that one's belief bear a causal connection to something whose existence guarantees—for one reason or another—the truth of the belief. Thus, on Dretske's view, in contrast to mine, it is not necessary that one's belief bear either a causal or explanatory connection to what makes the belief true. Further, there is some reason to think that the accounts will differ in the cases that they count as knowledge. For there can be a causal connection between one's belief that s is F and the fact that s is F , even if that connection does not proceed via something whose existence guarantees that s is F . Thus, one could fail to satisfy a necessary condition of Dretske's account of knowledge without failing to satisfy the externalist condition defended in the text. Also, since r 's being G could *entail* that s is F without r 's being G being caused by s 's being F , one could satisfy Dretske's conditions for knowledge while failing to satisfy the externalist condition of knowledge defended in the text.

²³Defeasibility views of knowledge were introduced in order to solve the Gettier problem, and, as an anonymous referee suggested, can also naturally be applied to the lottery paradox. For example, consider Ginet's (1975) account, offering necessary and sufficient conditions for S 's knowing the proposition that p : (i) p is true, (ii) S is confident that p , (iii) S 's being so is supported by a disinterested justification for p , and (iv) there is not true some proposition r such that were S to have justification for believing r and to retain all the properties entailed by his justification for claiming to know that p , then those properties plus the justification for believing r would be very far from justifying him in confidence that p (73–74). (If there is a true proposition r that meets the conditions specified in (iv), then I call it a “defeater” of S 's justification for p .) [See also Lehrer and Paxson 1969 and Ginet 1988 for related accounts.] This account can be applied to the lottery case in the following way: Jim lacks knowledge that his ticket will lose because there is a true proposition, T , such that if he had justification in believing it, he would be very far from justified in believing that his ticket will lose. T is the proposition that either Jim's ticket or ticket _{t} will win (where ticket _{t} is the actual winning ticket). The main problem with this solution is that the defeasibility account is itself vulnerable to counterexamples in which the defeater is “misleading.” For example, consider S 's belief that Tom Grabit stole a book, where S bases this belief on his having seen someone who looked just like Tom steal a book. In addition, S 's justification is defeated by the true proposition that Tom's mother explained that Tom has an identical twin brother who actually stole the book. But if it turns out that Tom's mother is a pathological liar and has completely fabricated this excuse for Tom, then the defeater is misleading and its existence does not deprive S of the knowledge that Tom stole a book. (See Lehrer and Paxson 1969 and Sosa 1970.)

Although Ginet does not claim that condition (iv) above explains Jim's

It will be useful to sum up the dialectic thus far. Externalist accounts of knowledge have been independently motivated by the need for a response to Gettier-type cases, among other things, and by the existence of causal connections in paradigm cases of knowledge.²⁴ Since certain forms of externalism offer promise in explaining why one lacks knowledge in lottery situations, they provide support for the view that (1) should be rejected. In particular, they support the idea that the rejection of (1) should be based on the fact that Jim's belief is probability-based, and so ultimately "uncon-

lack of knowledge in the lottery case, he offers two related principles that, in his view, do provide an explanation:

- (G1) In order to be justified in believing p on the basis of q , q cannot support an inference to the belief that there is a defeater. (1975, 78)
- (G2) In order to be justified in believing p on the basis of q , S must therein be justified in being confident that there is no defeater. (1975, 79)

According to Ginet, Jim meets neither condition because the very basis of Jim's belief that his ticket will lose supports an inference to the conclusion that there is a defeater for it. For Jim's belief is based on the proposition that "there will be a perfectly fair drawing of a single winning ticket from a lottery of one million tickets" (77). From this proposition, Jim could infer that there is a true proposition, namely, one of the same form as T , that defeats Jim's justification for his belief (78). There are two main problems for this account. First, the account explains why Jim lacks knowledge, but it cannot explain why Joan who plays a lottery with similar odds but with no guaranteed winner lacks knowledge. Yet, intuitively, if Jim lacks knowledge, so does Joan, and for the same reasons. Second, both (G1) and (G2) should themselves be rejected. These principles are questionable on the grounds that they appear to require the ability to employ sophisticated concepts like that of a defeater in order to be justified in any belief. But the principles also inherit the challenge to the original defeasibility account of knowledge. (G2) requires that one's justification for a proposition provide one with a justification for believing that there are no defeaters, misleading or otherwise. That this is too strong can be seen by returning to the Grabit case. For it is intuitive that S is justified in believing that Tom stole a book even if he is not justified in believing that there are no defeaters of the kind involving Tom's (insane) mother. Something similar applies to (G1): even if one's basis for a proposition provides one with justification for believing that there is a defeater, this should not necessarily preclude one from having justification, since the defeater might be misleading. One might even have justification for believing that the defeater in question is misleading, in which case one's claim to justification clearly remains unthreatened, despite the fact that one fails to satisfy (G1).

²⁴Other motivations include the desire to provide a response to certain kinds of skeptical challenges. See, for example, Armstrong 1973.

nected” with the fact that makes it true. At the same time, the fact that externalism could be used to ultimately ground an explanation of why beliefs based on P-inferences cannot constitute knowledge bolsters externalism in turn because it reveals yet more explanatory power for the theory.²⁵

It would be nice if we could leave things here. But, in addition

²⁵Armstrong appeals to a variant of the lottery paradox in defending a key part of his own externalist theory: the claim that in order for a believer to have noninferential knowledge, there must be a specification of the believer that ensures with “absolute reliability” that his belief is true (1973, 184–90). The argument is very different from the one presented in the text, and I do not believe the argument succeeds. It goes like this. First, assume the very plausible Conjunctivity of Knowledge Principle (CKP): if one knows that *p* and one knows that *q*, and one knows that *r*, and . . . then it is rational to believe that *p* and *q* and *r*, and . . .

Now imagine that when Bob is in a certain experiential state, there is a 99% chance that there is a sound in his immediate environment. Imagine further that he is in this state 200 times, and each time there is in fact a sound. Assume that Bob meets all necessary conditions for knowledge that there is a sound in his environment each time (he believes there is a sound, and so on) except for a “reliability” condition. In other words, leave it as an open question whether knowledge requires perfect reliability or whether 99% reliability will suffice for knowledge. Now assume for reductio that the first time Bob is in the relevant experiential state, he knows that there is a sound in his environment. Armstrong argues that the assumption that Bob knows (with only 99% reliability) that there is a sound in his environment will force us to deny (CKP). For if Bob knew that there was a sound the first time, then he knew there was a sound each time. By (CKP), it follows that it is rational for Bob to believe that there was a sound each time. But, intuitively, it is not rational for Bob to believe this. Contradiction. We should reject the hypothesis that Bob knew that there was a sound in his environment if there was only a 99% chance of his belief being true, given his state.

Armstrong’s reasoning fails because it is doubtful that it is not rational for Bob to believe there was a sound every time. We are to suppose that “all other conditions for knowledge” have been met by Bob, and these might include a rationality condition. As far as the story goes, Bob may have no awareness of his fallibility. There is no reason to assume that Bob is not rational in believing he heard a sound each time. If we do assume this, it might be because we also assume that Bob *believes* that the odds of his belief being true are 99%. But in that case it is not obvious that it is the *actual* probability of his being right that is getting in the way of Bob’s having knowledge, or Bob’s belief about it, or both.

Despite the failure of Armstrong’s argument, he is correct that a pure reliabilist account of knowledge will not be able to deny (1) if it assumes that (i) knowledge is constituted by belief based on a reliable method and (ii) reliability is measured entirely by the probability (less than 100%) that a method leads to true belief.

to objections that remain to be considered in the next section, there is the problem of the rationality version of the paradox. Placing external requirements on knowledge is widely accepted, but external requirements on rationality sound fishy. Surely whether a belief is rational or not does not depend on its truth (at least in most cases), and *a fortiori* on whether it is connected to the truth in right way.

Where should we go from here? I believe that it continues to be important, when thinking about the rationality version of the lottery case, to think of how beliefs and truth are connected, even if only in Jim's conception of things. Consider the fact that Jim does not (and cannot) see any causal or explanatory connection between his belief and the fact that makes the belief true. The belief on which he bases (or would base) his belief that his ticket will lose is not itself connected causally or, more broadly, explanatorily with the fact that his ticket will lose, and he can see this.

In the rest of this section, I will offer a Statistical Support account for the rationality version of the paradox. My proposal is that (1*) should be rejected, and that what explains the fact that it is not rational for Jim to believe that his ticket will lose is that he cannot see a causal or explanatory connection between his belief and the fact that makes it true. In fact, given the nature of his evidence, he can see that there is no such connection to be found. I believe that there is an internalized version of the "connection" requirement that holds for rational empirical belief, at least that based on inference.²⁶

²⁶My view is similar in spirit to the view put forward by Harman in his 1968 and 1973. Harman there argues that all (rational) inference to empirical facts must rely on an inference to the best explanation in order to constitute knowledge or justified belief, and thus, that "pure probabilistic inferences" [P-inferences] should be rejected. The object of the evidential belief(s) must be (seen to be) explained by the facts that make the belief constituting knowledge true, or the facts that make the knowledge true must be seen to explain the object of the evidential beliefs.

An important difference between my view and Harman's is the way in which we argue for them. In addition to the lottery paradox itself, Harman's main argument against the idea that inferences to 'p' from 'p has a high statistical probability' are rational is what I call the "false lemma" argument. It goes like this. Consider another case of Gettier's: Mary's friend, Mr. Nogot, tells her he owns a Ford, picks her up in a Ford, and so on. On the basis of this evidence, Mary concludes that one of her friends owns a Ford. However, Nogot's Ford has just been repossessed. Mary has another friend, Mr. Havit, who does own a Ford, although Mary does not

In other words, it must be possible for one to recognize one's commitment to the idea that one believes something *because* it is true. This does not mean that one must be able to represent to oneself the details of a causal or, more broadly, explanatory, connection between one's belief and its object. But one must be able to suppose that there is such a connection. In some cases, one can judge that there is a causal connection. For example, in the case of Smith's belief that Jones owns a Ford (see 391–92 above), Smith might be able to recognize that he takes there to be a causal connection between the fact of Jones's owning a Ford, his evidence, and his belief. In other cases, one can judge that there is a more general explanatory connection. For example, in the case of general beliefs, it might not be possible for there to be a causal connection between one's belief and its object. Consider my belief that all objects on the earth are affected by its gravitational pull. There is no obvious causal route from the fact that all objects have this property to my belief about it. Nevertheless, there is still an explanatory route. This fact explains other facts that have been perceived by people who have theorized about the nature of gravity

believe that he does. Mary does not have knowledge that her friend owns a Ford (despite the fact that her belief is true and justified). The reason, according to Harman, is that her belief is based on a false premise, or lemma (namely, that Nogot owns a Ford). Thus, a necessary condition for knowledge is that it not be based on a false lemma. But if purely probabilistic inferences can be rational, then Mary could have inferred her final conclusion from her original evidence, which is true, and which makes her conclusion highly probable. There would be no need to first infer that Nogot owns a Ford. Thus, if there could be pure probabilistic rules of evidence, then we could not account for the intuition that Mary lacks knowledge, or for the explanation of this lack of knowledge in terms of false lemmas. So, there cannot be pure probabilistic rules of evidence.

One problem with this reasoning is that it is at best an argument that belief based on pure probabilistic rules cannot constitute *knowledge*. Only if we make the additional assumption that true, rational belief based on true premises is sufficient for knowledge do we have an argument that pure probabilistic inferences are not rational. But even if we were to make this assumption, it is possible to distinguish the kind of probabilistic inference made in the hypothetical Mary case from the lottery case. An opponent might argue that the latter kind of inference, which is based entirely on statistical evidence, and where no other evidence is available, is rational, where the former is not.

It is also worth noting that Harman (1986) seems to reject his earlier view that P-inferences are not rational. I consider his reasons in section 6. See also section 2 above.

and from whose testimony I have ultimately benefited. I can represent to myself an explanatory connection (which includes causal elements) between the fact in question and my belief about it. Again, the key idea is that in each case I can take myself to believe something *because* it is true. Thus, while many explanatory connections are causal, what is essential is that one be able to posit the existence of an explanatory connection between one's belief and the object of one's belief. One need not have any clear idea about what that connection is, but on reflection one must be able to recognize that one is committed to there being such a connection.

Let us see how this proposal sorts cases. Consider ordinary cases like (F) and (N), in which we think that our beliefs are rational. In the case of (F), we are rational in believing the furniture is still there because we assume that we saw the furniture there moments earlier and we have not heard any unusual sounds, and so on. What best explains the experiential evidence on which we base our belief is that nothing out of the ordinary has happened there and that, consequently, the furniture is still there. Similarly for the case of (N). Our belief that the Bulls won is rational because we believe (or would, on reflection) that someone at the newspaper received a report of the conclusion of the game in which the Bulls won (either by watching it or by its being reported by someone who did see it), and subsequently entered the results to be printed in the paper. It is possible for us to represent to ourselves that there is a causal (and hence explanatory) connection running from the Bulls game, through our evidence for it, and ultimately to our belief about it. Suppose instead that we believed that the sports editor simply picked scores at random to report, or that a freak snowstorm had cut off all communication from inside the basketball arena from half-time of the game until now. Then it seems it would not be rational for me to believe that the Bulls won on the basis of the newspaper report. Thus, in paradigm cases of rational belief, it seems it must be possible for me to see a connection—causal or explanatory—between my belief and the fact that makes it true. I need not be able to see any details of the connection, but I must be able to take there to be such a connection. In contrast to the furniture and newspaper cases, it is clear from the nature of the evidence in the lottery case that there could be no causal or explanatory connection between my evidence and the facts.

Thus, this criterion for rational belief supports the rejection of (1*) as a solution to the rationality version of the paradox.

At this point an issue arises that calls for clarification. I have argued that one must be able to postulate a causal or more broadly explanatory connection between one's belief and the object of belief in order for one's inferred (nondeductive) belief to be rational. But one might worry that this is so general as to allow that someone might irrationally postulate a causal connection between any beliefs and their objects. In particular, Jim might irrationally take there to be such a connection between his belief that he will lose the lottery and the fact that he will lose. So perhaps the criterion does not properly sort the cases after all.

In reply, let me note that if Jim were to irrationally postulate a causal connection between his belief that he will lose the lottery and the fact that he will lose, he is thereby guilty of irrationality. So this case does not show that the criterion might count Jim as rational in believing (L). Thus, the criterion remains unthreatened by this sort of case. Nevertheless, cases of this sort show that Jim's ability to represent a causal or explanatory connection must itself be rational in order for his belief to be rational.

The "internalized connection" criterion seems to sort cases appropriately. Nevertheless, more needs to be said in its favor. For it is compatible with my analysis of the cases of (F) and (N) that there is simply more than one kind of rational support for beliefs, and that P-inferences are rational. Why think that there is an "internalized connection" requirement for all rational belief? This would be the place to say something deep about the nature of rational belief. Instead, I'll try to shift the burden of proof, and make some remarks on rational belief that I hope will be suggestive.

First, let me try to defuse the sense that there is a need to recognize P-inferences as a type of rational inference. I believe that the feeling that there is such a need comes from the fact that claims like (1*) are intuitively appealing. Can the appeal of (1*) be explained in a way that does not commit us to accepting (1*)? Let me offer such an explanation. Strongly associated with the appeal of (1*) is the conviction that Jim should not make plans contingent on his winning the lottery, and, in fact, that he should make plans based on the assumption that he is not going to win. Similarly, we think that he should not "get his hopes up," that he

should not devote excessive energy to creating mental images of his winning, and that, under ordinary circumstances, he should not buy a lottery ticket if the price is high enough to be important to him. In fact, I believe that the appeal of (1*) derives largely from the fact that Jim's belief that he will lose rationalizes all of these behaviors that we find rational and normative for Jim.

But all of our beliefs about what it is rational for Jim to do could be true and yet (1*) be false. For compatible with the falsity of (1*) is the truth of

(1**) Jim is rational in believing he will very probably lose.

All of the behaviors just mentioned (making plans on the assumption that he will lose, not getting his hopes up, and so on) are sufficiently rationalized by Jim's belief that he will very probably lose. (1*) seems unnecessary.

Further, if most or all of our intuitions can be preserved by accepting (1**) in place of (1*), then there is dialectical reason to take this route. For doing so allows us to preserve (8*).

Finally, although I can do no more than gesture here, I believe that it is plausible that rational (nondeductive) inferences are essentially ones in which one can conceive some causal or explanatory connection between one's belief and the fact believed. In very general (and grand) terms, the role of rationality in one's belief-forming activities is to guide one's beliefs toward the truth, via reasons. There is a way in which this role is fulfilled in inferences to beliefs like (N), but not in the case of P-inferences.

In the everyday case like (N), as we saw above, the rationality of one's belief depends on one's being (rationally) committed to the proposition that there is an explanatory connection between one's belief and what the belief is about. But now suppose that I learn that my belief is false, and that the Bulls in fact lost their game against the Knicks. In that case, I would now reject something to which I was previously committed: that the newspaper reported the score accurately, or that there was a causal connection between the reporting of the score and the game itself, or something else in my total evidential package. In recognizing the falsity of my belief, I realize that my evidence (although perhaps quite reliable) failed to connect to the truth in some way. Either the evidence was inaccurate, or it was not connected with the truth in the right way.

Importantly, I do not need to see myself as having been irrational or at fault in any way. But once I realize that my belief is false, I must reject at least part of what gave me reason. In other words, in seeing that my reasons failed to guide me to the truth, I must reject something in my total “package” of reasons.

Notice, in contrast, that this is not the case for P-inferences. If I believe that p (say, my ticket will lose) on the basis of a high statistical probability for p , and I find out that not- p (I won!), then there is nothing at all in my reasons to reject. I still believe the same odds were in effect, and I still believe that they made my losing extremely probable. I have no reason to think that my evidence failed to bear a connection to my conclusion that I previously thought it did. Learning that my belief is false puts no pressure on me to find some problem in my reasons. Thus, there is a way in which they are not “sensitive” to the truth, or at least to what I conceive of as the truth. It seems to me that, given the role of rationality as a guide toward truth, this lack of sensitivity to the truth in the case of P-inferences might help to explain why such inferences are not rational.

Admittedly, this is more of a suggestion than a defense, and I do not mean to rest my case on it. Thus, even if one rejects the suggestion, I hope to have done enough to shift the burden of proof onto those who would prefer an alternative solution to the rationality version of the paradox. I have tried to do this in two ways: by defusing a natural resistance to rejecting (1*), and by showing that the main rival to this solution has problems that cannot be satisfactorily resolved.

Before turning to objections to the Statistical Support accounts presented in this section, let me briefly address the question of how the two versions of the paradox are related. In addressing the knowledge version, I have offered a Statistical Support account that explains why (1) should be rejected in terms of the probabilistic inference on which Jim’s belief is based, and his consequent failure to meet an externalist condition for knowledge. I have also offered a Statistical Support account that explains why (1*) should be rejected. It, too, begins with the probabilistic nature of the inference on which Jim’s belief is based. But it continues by pointing out that such inferences do not satisfy an *internalized* version of the requirement that one’s beliefs be appropriately connected to their objects.

If one accepts that rational belief is a necessary condition for

inferential knowledge, then one could accept the Statistical Support account of the rationality version of the paradox as a solution to the knowledge version as well. The same fact that makes (1*) false also makes (1) false. In that case, one could either dispense with the independent “externalist” Statistical Support account as a solution to the knowledge version of the paradox, or one could accept both and argue that it is overdetermined that Jim does not know that his ticket will lose.

Although it would be simpler to accept a single solution for both paradoxes, it is not necessary.²⁷ Even if both Statistical Support accounts are accepted, it is important to see that they share a common starting point. According to both, what ultimately deprives Jim of both knowledge and rational belief that his ticket will lose is the probabilistic inference on which his belief is based.

6. Objections and Replies

6.1 *Bonjour on Externalism and the Knowledge Version of the Lottery Paradox*

The only criticism of an externalist solution to the knowledge version of the lottery paradox of which I am aware appears in Bonjour 1985. Bonjour suggests that the ability of externalism to solve the

²⁷Here I disagree with Ryan, who argues that on plausible assumptions, “the only option open to explain why [Jim’s] belief fails to be an instance of knowledge is that it fails to satisfy the justification condition for knowledge” (1996, 136). Given that the justification condition is a rationality condition for Ryan, her conclusion strongly suggests that a single solution to both versions of the paradox should be provided. Her “What Else Could it Be?” argument goes like this: Jim does not know that his ticket will lose. There are four possible explanations for this: he does not believe it, he is not justified in his belief, his belief is false, or he is in a “Gettier” situation (that is, his belief is based on defective evidence). But, by hypothesis, Jim believes truly that he will lose, and his evidence is in no way defective. What else could it be? The only remaining alternative is that he is not justified (or rational), at least to the degree required for knowledge. The argument is clever, but goes wrong in assuming that there are only four possible ways that a belief could fail to constitute knowledge. For the correct analysis of knowledge might require an external condition in addition to good and nondefective evidence. For example, Goldman’s causal account of knowledge includes a causal requirement in addition to justified true belief. In that case, the obstacle to knowledge in the lottery case *could* be something else, namely, that Jim’s belief is not connected with the truth in the right way.

lottery paradox would count as a positive argument for externalism. And Bonjour acknowledges that externalism, or, at least, one form of it, appears promising at first glance. In his view, the lottery paradox appears problematic because it seems to require absolute certainty for a belief to count as knowledge. Yet, the skeptic would have an easy time showing that, with this criterion for knowledge, almost nothing that we ordinarily take to be knowledge really counts as knowledge. Bonjour conceives of externalism as solving the problem by providing a truth guaranteeing property of beliefs without that guarantee of truth being certainty (or, in other words, without the guarantee of truth being within the cognitive grasp of the believer). For, according to one version of externalism (Armstrong's), there is some description of the believer on which the belief is connected by a law of nature to the fact believed. Thus, a belief can be guaranteed to be true without the believer's having to know this. Further, it is reasonable to suppose that there is a relevant description and law in a number of cases that we normally count as knowledge.

Bonjour criticizes this account on the following grounds. Imagine a situation in which Agatha knows that she is one of 100 people who are participating in a philosophical experiment conducted by a Cartesian demon. Each is seated at her desk and believes herself to be perceiving a cup. But, although 99 of the subjects will be perceiving a cup in the normal way, the last one will be caused by the demon to have a complete hallucination of a nonexistent cup. Agatha knows all of this, but does not know whether she is the one being deceived. Bonjour argues that Armstrong's externalism will give the wrong result here. Suppose that Agatha is indeed perceiving a cup. Intuitively, Agatha does not know that she is perceiving a cup. At the same time, Armstrong's externalism implies that she does. Thus, externalism cannot provide a solution to the paradox.

In responding to this objection, I should note first that Bonjour is criticizing only one form of externalism. Goldman's causal externalism makes no claims to provide a truth-guarantee (although, as we saw above, it can be employed to solve the paradox by providing *another* kind of distinction between the lottery case and cases of knowledge). Nevertheless, Bonjour's argument might be adapted to target Goldman's causal account. Would it give the wrong result in Agatha's case?

It would if it implied that the cup's causing Agatha's perceptual

experience and belief were sufficient for knowledge. But Goldman's causal requirement is meant only as a necessary condition for knowledge. The problem is that Bonjour's argument against externalism works only on the assumption that externalism holds an externalist condition (along with true belief) to be sufficient for knowledge. But in order to provide a solution to the lottery paradox along the lines sketched in section 5 above, externalism need only provide a necessary condition for knowledge that is externalist. Thus, as far as Bonjour's argument goes, it leaves even Armstrong's form of externalism as an available option for solving the lottery paradox (assuming it is taken only as a necessary condition).²⁸

6.2 Harman's Challenges

Although Harman favors a solution similar in spirit to the one I offer for the rationality version of the paradox in his 1968 and 1973, he rejects it in his 1986. He also raises serious challenges even for the rejection of (1) in the knowledge version of the paradox.

Although Harman accepts that all rational inference is inference either from or to the best explanation, he also accepts that P-inferences are rational. The reason he gives for this is that inferring something based on its statistical probability is of the form: conclusion *because* evidence. In other words, the evidence (statistical probability) explains the conclusion.

On the face of it, this seems the wrong way to describe such an inference. Why think there is any explanatory link at all between the statistical probability of p and p itself? I believe Harman's reasoning goes something like this. One can infer that a die is "load-

²⁸Further, it is not clear that Bonjour's case succeeds as a counterexample to Armstrong's account, even if it is taken as providing a sufficient condition for knowledge. For in the case in which Agatha's belief might have been caused by an evil demon, it seems that the kind of nomological requirement Armstrong had in mind is not satisfied. For Armstrong claims that there must be a law connecting the belief and the truth such that having the belief ensures its truth.

Also, it is worth noting that Armstrong intends the account described in the text as an account of *noninferential* knowledge. Armstrong's account of inferential knowledge, which would presumably be required to explain the lottery situation, is more complicated. See Armstrong 1973, 198–211.

ed” so as to favor side six on the basis of the frequency with which six comes up when the die is tossed. As Harman puts it,

[t]his is to reason from the observed evidence to a statistical explanation of the evidence. One concludes that the best explanation of the observed evidence is that the probability of getting a six on tossing this die is greater than one out of six. One infers that the observed frequency of sixes has occurred because the die is loaded. (1986, 70)

Thus, one can rationally infer a statistical explanation from an outcome. But there can be rational inferences both to and from explanations. Hence, it follows that one can also rationally infer an outcome from a statistical explanation.

For Harman’s argument to go through, he must be making two assumptions. The first is that if a proposition of type A explains a proposition of type B in one case, then propositions of type A explain propositions of type B in another case. The second is that rational inference “goes both ways”; that is, if it is rational to infer *from* an explanation, it is rational to infer *to* an explanation, and vice versa. Both assumptions are questionable.

The first assumption is needed in order to show the relevance of the die case to the lottery case. The problem is that the assumption leaves it open that the types in question are very general. And Harman requires that they be quite general. In particular, Harman needs it that if antecedent probabilities explain outcomes in one case, then antecedent probabilities explain outcomes in other cases. But this is questionable. In the die case, there is a series of outcomes (or serial outcome), while in the lottery case there is a single outcome. This difference is important because the antecedent probabilities in the lottery case cannot explain the outcome. In that case, the outcome is that t59,008, say, won and that the other tickets did not win. But antecedent probabilities do not explain these facts. Thus, Harman fails to show the relevance of the die case to the lottery case.

The other problem with the reasoning is that rational inferences do not necessarily go both from and to explanations in every case. Even if we grant that in the die case the outcomes are explained by the antecedent probabilities and that it is rational to infer from the outcomes to the antecedent probabilities, it does not follow that it is rational to infer from antecedent probabilities to outcomes. Suppose that we know that a die is loaded so as to make

the likelihood of a six coming up $1/2$. We cannot rationally conclude what a particular set of tosses of that die will be. Thus, the die case cannot be used to show that P-inferences are rational.

Nevertheless, Harman argues in a similar way when he raises a problem for the Statistical Support account for the knowledge version of the paradox. He writes that it “can seem wrong” that one knows that one’s ticket will lose in a lottery. But

an inference in the other direction, from observed evidence to a statistical explanation of the evidence, can give one knowledge . . . one can come to know the probability of six on this die is closer to $1/2$ than to $1/6$ Why is knowledge clearly possible in the one case and not clearly possible in the other? (1986, 71)

Harman does not have an answer to this question. But again it seems to me that he has ample resources to answer it if he returns to the important idea that rational inference must be explanatory. For it is plausible to argue that in the die case, the statistical probabilities provide an explanation for the evidence whereas in the lottery case, the statistical probabilities do not provide an explanation for the inferred outcome. The cases are different in relevant respects (for example, there is a series of outcomes in the die case and not in the lottery case). Further, reflection on the die case itself casts doubt on the idea that it is always rational to infer in one direction if one can rationally infer in the other. If inferential knowledge requires that it be based on rational inference, then there is no reason for doubting that Jim lacks knowledge as well as rational belief that his ticket will lose.

Finally, Harman offers an additional reason to question the Statistical Support account of the knowledge version of the lottery paradox (and any other account that rejects (1)). Together with the worry just mentioned, it leads Harman to conclude that “he has no idea how to account for our reluctance to attribute knowledge in cases of this sort” (1986, 72). Here is the problem:

Suppose Bill wants to know where Mary will be tomorrow. Bill knows that Mary intends to be in New York. Bill also knows that if Mary’s ticket is the winning ticket, she will instead be in Trenton for the award ceremony. But there is only one chance in a million of that. Can’t Bill conclude that Mary will be in New York tomorrow and in that way come to know where Mary will be tomorrow? That seems possible. But

doesn't it involve knowing her lottery ticket is not going to be a winning ticket? (1986, 71)

This is a challenging case. A first response is to deny that Bill's knowledge is based on an inference from the statistical probability of Mary's being in New York. Nothing in the case forces us to say that it is, and yet, my reason for denying that we have knowledge in lottery situations is that they involve a belief based on an inference from statistical probabilities. Nevertheless, the case seems to pose a problem for denying (1) in the lottery case. To see why, we need to look more closely at the structure of the case. The case is meant to reveal that (i) one could know 'p' (that is, that Mary will be in New York), (ii) know that 'p entails q' (that is, that her being in New York entails that she will not win the lottery), but (iii) not know 'q'. The problem is that (i) and (ii), together with a very plausible assumption that one (can) know the logical implications of one's beliefs, entails that (iii) is false. But (iii) derives support from the same considerations that lead us to say one lacks knowledge in the lottery case.

Of course, we could continue to maintain that Bill has independent reason to believe that Mary's ticket will lose, and that he does not *base* his belief that she will lose on the statistical probability that she will. So there is no real conflict with denying that, ordinarily, we lack knowledge in lottery situations. However, it does not seem that Bill really does have independent reason to believe that Mary's ticket will lose. So this response does not really relieve the sense that something has gone wrong somewhere.

I think the best way to respond to the case is to deny that Bill can know that Mary will be in New York. If Bill's basis for believing Mary will be in New York is that she intends to be there and that there is a 99.99% chance that she will be where she intends to be, then Bill doesn't know that she will be in New York. He knows that she will be in New York or Trenton, he knows that she will very likely be in New York, he should make plans on the assumption that she will be in New York, and so on. But if he knows that she would be in Trenton for the awards ceremony were she to win, then he does not know that she will be in New York. Thus, there are no obvious consequences for Bill that hinge on whether he knows that Mary will be in New York or only knows it extremely likely. Further, it is important to note that we are not in Bill's sit-

uation very often. This means that it remains open that we often know where people will be (and not just where they are very likely to be). Thus, although it might seem at first that denying that Bill knows that Mary will be in New York forces us to give up a number of intuitions, reflection shows that this is not the case. Harman's worries, though natural, are ultimately misplaced.

7. Conclusion

In both of its versions, the lottery paradox poses serious challenges to our assumptions about knowledge and rational belief. A satisfying solution should not only answer these challenges, but also explain both why our initial assumptions might have been wrong and why we had them in the first place. I believe that the Statistical Support accounts go a long way in this direction. The lottery situation is unusual because of the probabilistic nature of our evidence. Surprisingly, this fact is often de-emphasized in discussions of the lottery paradox. Yet, taking this feature as our starting point allows us to ask why beliefs based on probabilistic inferences are problematic. And we discover an answer: such beliefs fail to be connected with the truth in the appropriate ways, both from an objective point of view and from the point of view of the believer.

Florida State University

References

- Armstrong, David. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Bonjour, Lawrence. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- Cohen, Stewart. 1988. "How To Be a Fallibilist." In *Philosophical Perspectives*, vol. 2, ed. James Tomberlin, 91–123. Atascadero, Calif.: Ridgeview Press.
- DeRose, Keith. 1996. "Knowledge, Assertion, and Lotteries." *Australasian Journal of Philosophy* 74:568–80.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Evnine, Simon. 1999. "Believing Conjunctions." *Synthese* 118:201–27.
- Foley, Richard. 1979. "Justified Inconsistent Beliefs." *American Philosophical Quarterly* 16:247–57.
- . 1987. *The Theory of Epistemic Rationality*. Cambridge: Harvard University Press.

THE LOTTERY PARADOX

- . 1993. *Working Without a Net*. Oxford: Oxford University Press.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–23.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. Boston: D. Reidel.
- . 1988. "The Fourth Condition." In *Philosophical Analysis: A Defense By Example*, ed. D.F. Austin, 105–17. Dordrecht: Kluwer Academic Publishers.
- Goldman, Alan. 1984. "An Explanatory Analysis of Knowledge." *American Philosophical Quarterly* 21:101–08.
- Goldman, Alvin. 1967. "A Causal Theory of Knowledge." *Journal of Philosophy* 64:355–72. Reprinted in *The Theory of Knowledge*, ed. L. Pojman, 144–153 (Boston: Wadsworth, 1999).
- Harman, Gilbert. 1968. "Knowledge, Inference, and Explanation." *American Philosophical Quarterly* 5:164–73.
- . 1973. *Thought*. Princeton: Princeton University Press.
- . 1986. *Change in View*. Cambridge: MIT Press.
- Klein, Peter. 1987. "On Behalf of the Skeptic." In *The Possibility of Knowledge*, ed. Steven Luper-Foy, 267–81. Totowa, N.J.: Rowman & Littlefield.
- . 1995. "The Virtues of Inconsistency." *Monist* 68:105–35.
- Kyburg, Henry. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Lehrer, Keith, and Thomas Paxson Jr. 1969. "Knowledge: Undefeated Justified True Belief." *Journal of Philosophy* 66:225–37.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge: Harvard University Press.
- Ryan, Sharon. 1991. "The Preface Paradox." *Philosophical Studies* 64:293–307.
- . 1996. "The Epistemic Virtues of Consistency." *Synthese* 109:121–41.
- Sosa, Ernest. 1970. "Two Conceptions of Knowledge." *Journal of Philosophy* 67:59–66.