

A CASE AGAINST CLOSURE

Doris Olin*

SÍNTESE – Este artigo examina a objeção ao fechamento [dedutivo] que surge no contexto de certos paradoxos epistêmicos, paradoxos cuja conclusão é que a crença justificada pode ser inconsistente. É universalmente aceito que, se essa conclusão é correta, o fechamento deve ser rejeitado, para que se evite a crença justificada em enunciados contraditórios (P , $\sim P$). Mas, mesmo que os argumentos desses paradoxos – o paradoxo da falibilidade (do prefácio) e o paradoxo da loteria – sejam mal-sucedidos, eles, ainda assim, sugerem a existência de evidência independente para uma objeção mais direta contra o fechamento. O exame do argumento da falibilidade revela uma exigência de modéstia epistêmica que viola o fechamento a partir de múltiplas premissas. A reflexão sobre o paradoxo da loteria nos confronta com um dilema em que cada alternativa fornece um contra-exemplo ao fechamento a partir de uma única premissa. Seja ou não possível a inconsistência racional, há uma objeção contra o fechamento.

PALAVRAS-CHAVE – Fechamento dedutivo. Falibilidade. Paradoxo da Loteria. Paradoxo do Prefácio. Justificação. Inconsistência.

ABSTRACT – This paper examines the case against closure that arises in the context of certain epistemic paradoxes, paradoxes whose conclusion is that it is possible for justified belief to be inconsistent. It is generally agreed that if this conclusion is correct, closure must be rejected in order to avoid justified belief in contradictory statements (P , $\sim P$). But even if the arguments of these paradoxes – the fallibility (preface) paradox and the lottery paradox – are unsuccessful, they nonetheless suggest independent grounds for a more direct case against closure. Examination of the fallibility argument reveals a requirement of epistemic modesty that violates multiple premise closure. Reflection on the lottery paradox presents us with a dilemma in which each alternative provides a counterexample to single premise closure. Whether or not rational inconsistency is possible, there is a case against closure.

KEY WORDS – Closure. Fallibility. Lottery paradox. Preface paradox. Justification. Inconsistency.

I

The principle of epistemic closure, according to which knowledge and justification of belief extend over known implication, is currently under attack from several quarters. Recent discussions of scepticism, for instance, have targeted the closure principle as the villain in the traditional sceptical argument, and have thus focused on discrediting the notion that knowledge is always closed under logical implica-

* Department of Philosophy – Glendon College, York University – Canada.

VERITAS	Porto Alegre	v. 50	n. 4	Dezembro 2005	p. 235-247
---------	--------------	-------	------	---------------	------------

tion.¹ Closure has also been under siege in connection with two epistemic paradoxes, paradoxes whose conclusion is that justified belief can be inconsistent. If this conclusion is correct, then we are faced with an apparent consequence that seems clearly unacceptable, namely, that belief in contradictory statements (P , $\sim P$) can be reasonable. To escape this consequence, many have rejected certain versions of the closure principle.²

In this paper, I examine only one aspect of the case against closure, specifically, that which arises in the context of the two epistemic paradoxes concerning inconsistency – the fallibility (preface) paradox and the lottery paradox. Even if their arguments for justified inconsistency are not successful – I think they are not – and therefore do not require the denial of closure to avoid justified belief in contradictions, they nonetheless suggest other considerations that provide a more direct case against closure.

The intuitive appeal of closure lies in the sense that in a valid argument, the epistemic status of the premises extends to the conclusion. Precise formulations of this intuitively plausible thesis in the literature vary along at least two parameters. One standard version of closure is formulated in terms of knowledge: If S knows that P , and P implies Q , then S knows that Q . Another version pertains to justification: If S is justified in believing P , and P implies Q , then S is justified in believing Q . My interest here is in the latter principle, which is arguably more fundamental. Both versions, however, require restriction, for there are logical consequences of your justified beliefs (knowledge) which are complex or distant, and of which you are quite unaware. Thus, a more defensible formulation in terms of justification is:

Closure: If S is justified in believing P , and knows that P implies Q , then S is justified in believing Q .

A second dimension along which closure principles may vary concerns the number of premises allowed in the deduction. The version just enunciated allows just one premise P ; call this 'Single Premise Closure' (SPC). A stronger variant is Multiple Premise Closure (MPC):

Multiple Premise Closure: If S is justified in believing P_1, \dots, P_n , and knows that P_1, \dots, P_n jointly imply Q , then S is justified in believing Q .

(MPC), of course, implies (SPC), but not conversely. Note that both (SPC) and (MPC) are expressed in terms of what the subject is justified in believing, as opposed to what she justifiably believes.³ To say that a subject S is justified in believing P does not imply that S in fact believes P and, therefore, does not imply that she justifiably believes P . 'S is justified in believing P ', as here understood, implies only that S is evidentially in a position to justifiably believe P .

¹ See, for example, Fred Dretske, 'Epistemic Operators', *Journal of Philosophy* 69 (1970), 1007-1022; and his more recent 'The Case Against Closure', in E. Sosa and M. Steup (eds.), *Contemporary Debates in Epistemology* (Malden, Ma: Blackwell Publishing, 2005), 24-49.

² See, for example, Richard Foley, 'Justified Inconsistent Beliefs', *American Philosophical Quarterly* 16 (1970), 247-57; and Peter Klein 'The Virtues of Inconsistency', *The Monist* 68 (1985), 105-135. Specifically, these authors reject the principles (MPC) and (CP), introduced below.

³ Peter Klein does the same; see his *Certainty: A Refutation of Scepticism* (Minneapolis: University of Minnesota Press, 1981).

This formulation of the closure principles reflects the intuition that what one may justifiably believe extends over known implication. But there are situations in which recognition of what is entailed by your justified belief P should, rationally, lead you to give up the belief in P. This does not, however, necessitate any revision or qualification of our closure principles. For the advocate of closure can maintain that if S comes to see that P implies Q, and Q is clearly unacceptable given S's total evidence, then *P is no longer justified for S*. It is understood that each of the relevant clauses refers to the same time, that is: if S knows at t that P implies Q, and, at t, Q is clearly unacceptable given S's total evidence, then P is not justified for S at t. Similarly, each clause in a closure principle must be understood to refer to the same time.

Finally, it is instructive to compare our version of closure with two other closely related principles. Consider:

Transmission of Evidential Support (TES): If E is evidence for P, and P implies Q, then E is evidence for Q.

Those who accept (TES) will likely find the following plausible:

Transmission of Justification (TJ): If S is justified in believing P, and knows that P implies Q, then S is justified in believing Q on the basis of P.⁴

(TJ) makes a claim as to *how* Q is justified; Q is justified on the basis of P, which is itself justified. This is in contrast to the closure principles above, which leave it an open question how Q is justified.

Neither of the two Transmission principles, however, stands up to scrutiny. Suppose I justifiably believe

P: Cynthia will watch television the entire evening on Wednesday.

based on her announced intention to do so; and it is clear to me that P implies

Q: There will not be a power failure on Wednesday evening.

Now consider, first, the Transmission of Evidential Support principle. Note that my evidence that Cynthia will watch television on Wednesday evening, Cynthia's announced intention, is clearly not evidence that there will not be a power failure on Wednesday evening. So the evidence for P is not evidence for Q. Turning to the Transmission of Justification, my belief that Q can hardly be justified on the basis of my belief that P. For one thing, there is the fact that the evidence for P is not evidence for Q. But there is a more decisive consideration. That Cynthia's expression of intent provides adequate evidence to justify me in believing P is contingent on my being entitled to take for granted the truth of Q, or to ignore the possibility that Q is false. That being so, I can hardly appeal to P to justify a belief that Q.

⁴ This is the terminology of Crispin Wright, who uses the term 'transmission' for a principle like (TJ) in 'Anti-sceptics simple and subtle: Moore and McDowell', *Philosophy and Phenomenological Research* 65, 2002, 330-48.

That is, P cannot serve as the evidence which justifies me in believing Q.⁵ There are relations of epistemic priority or ordering that are not captured in either the transmission or closure principles, but which pose a problem only for transmission. More intriguing examples of this will emerge later in the discussion.

II

The context for this examination of closure is the possibility of justified inconsistent belief. Two main arguments for the rationality of inconsistency can be distinguished in the literature: (i) the fallibility paradox (which is a more general version of the preface paradox) and (ii) the lottery paradox. Let us begin with the fallibility paradox. The key fact on which it is based is that we are fallible in that we have often discovered error in our justified beliefs, beliefs based on the data of perception, testimony, memory and so on. If I have often discovered error in my justified beliefs formed by method M in the past, then it seems plausible to infer that there is error in my present justified beliefs formed by the same method. But my present justified beliefs are P_1, \dots, P_n ; so I am also entitled to believe $\sim(P_1 \& \dots \& P_n)$.^{6,7} Thus I am justified in believing inconsistent statements.

Note that in order to have justified belief in an inconsistency, I must work through the steps in the argument; for instance, I must be justified in believing that there was error in my past justified beliefs, and must make the inference to error in my present justified beliefs. In particular, I must be justified in believing that my present justified beliefs are P_1, \dots, P_n . This is a strong and unrealistic requirement; it is arguable that none of us has the appropriate justified belief. Nevertheless, it seems *epistemically possible* that a rational being should have such a belief. So the argument can be construed as establishing at least that it is epistemically possible to have justified inconsistent belief.

What is the impact of the fallibility argument on closure? Grant that a set of inconsistent statements can be shown to imply contradictory statements, that is, statements of the form: $P, \sim P$. Thus, if justified belief in inconsistent statements is possible, so, it would seem, is justified belief in a contradiction. But for the vast majority of philosophers, the notion of contradictory justified beliefs is anathema. To escape from this apparent embarrassing consequence, the proponent of justified inconsistency is forced to deny Multiple Premise Closure (MPC). Nor is this the only cost of sanctioning inconsistency. If Single Premise Closure (SPC) is to be retained, we must also reject:

⁵ Wright offers a similar analysis, *op. cit.* Stewart Cohen agrees with the general point in 'How To Be a Fallibilist', *Philosophical Perspectives*, 2, 1988, 118; John Hawthorne disagrees in *Knowledge and Lotteries* (Oxford: Clarendon Press, 2004), 34.

⁶ The argument should be more carefully stated in terms of 'first order justified beliefs' so as to avoid the problems created by self-reference. If (Z) (There is error in my present set of justified beliefs) were included in the beliefs referred to in (Z), then (Z) could not be false (if it were false, it would be true). Further, if all the other statements in the belief set were true, it would then be case that (Z) was false if and only if it was true. I opt first for simplicity, and then introduce the qualifier later in the discussion.

⁷ Note that the preface paradox is the same argument applied to an author's belief in the statements in her book.

Conjunction Principle (CP): If S is justified in believing P, and S is justified in believing Q, then S is justified in believing P&Q.

For (CP) combined with (SPC) has the logical force of (MPC).⁸ Thus, the cost of granting the fallibility argument for inconsistency is that two closure principles, (MPC) and (CP), must be rejected. (SPC), however, may be retained.

I believe that the fallibility argument is flawed.⁹ Very briefly: The fallibility argument moves from errors in past beliefs to errors in the present set of beliefs formed by the same method. That a method of belief formation has led to error in the past indicates that the evidence employed is not perfectly reliable. This means that an air-tight connection with truth is absent; but it does not follow that there is a reliable, projectible connection with falsity. For instance, reading tea leaves is a method of belief formation that has certainly led to error in the past. Suppose that present tea leaf readings have led, via a standard method of interpretation, to a set of statements about the future, R_1, \dots, R_n . If the fact of the tea leaf readings constitutes reason to think there is error in the set of R_i s, that is, to think $\sim(R_1 \& \dots \& R_n)$ is true, then it appears that tea leaf reading can provide reason for beliefs about the future (since the negation of the conjunction of the R_i s is a statement about the future).¹⁰ But this seems absurd. If, on the other hand, facts about tea leaf reading are simply irrelevant to the truth or falsity of the R_i s, as we normally think, then the key inference in the fallibility argument seems illegitimate.

Still, even if defective, the argument invites us to consider what each of us should believe about her first order beliefs.¹¹ We are admittedly fallible creatures; our beliefs frequently contain error, not just our careless or irresponsible beliefs, but even our perfectly reasonable beliefs. Recognition of our fallibility, in this sense, surely warrants some degree of epistemic modesty. Perhaps it is going too far to say that we should believe that there is error in our present first order justified beliefs. But at the least, we ought not to believe that all our first order justified beliefs are true, that is, we should withhold belief on the matter. This seems the minimal requirement of epistemic modesty.

Even this moderate response, however, has repercussions for closure. Consider the statement that *all* beliefs of a certain sort are true; it seems clear that this is not equivalent, as is sometimes assumed, to the conjunction of the beliefs of

⁸ To see this, note first that (CP) can be employed to join any number of statements in a conjunction as follows. The first two statements are conjoined; then the resultant conjunction is conjoined with the third statement; and so on. Assuming we are justified in believing each of an initial set of statements, we are, by repeated applications of (CP), justified in believing their conjunction; and then by one application of (SPC), justified in believing whatever follows from the initial set. Thus, we have all the effects of (MPC).

⁹ I argued this in *Paradox* (Chesham: Acumen Publishing, 2003), 68-69. There the argument is intricate and apt to be mis interpreted. I now believe this much simpler version is adequate.

¹⁰ Note that the issue is not whether there is reason for thinking $\sim(R_1 \& \dots \& R_n)$ is true, but whether the tea leaf readings provide such a reason.

¹¹ Here I introduce the qualifier 'first order' to avoid issues of self reference.

that sort.¹² If A, B and C are all the beliefs of type Φ , I may believe the conjunction A&B&C without believing that all beliefs of type Φ are true; for I may not realize that A, B and C are all the type Φ beliefs. The following, however, seems a valid argument:

- Argument α*
- (1) A, B, C
 - (2) The beliefs of type Φ are: A, B, C.
 - (3) So all the beliefs of type Φ are true.

(Note that (1) expresses three distinct premises.) Now suppose that P_1, \dots, P_n are T's present first order justified beliefs. If argument α is valid, then so is argument β as follows:

- Argument β*
- (1) P_1, \dots, P_n
 - (2) T's present first order justified beliefs are: P_1, \dots, P_n .
 - (3) Therefore, all T's present first order justified beliefs are true.

Given the validity of argument β , if T is justified in believing the premises of the argument, and sees that the premises imply the conclusion, she should, by (MPC), be justified in believing the conclusion. It is a given that she is justified in believing the statements P_1, \dots, P_n . As noted earlier, one is not normally justified in believing a statement such as (2), but it is epistemically possible that one should be. Assuming that T is so justified, and sees the validity of the argument, (MPC) warrants her in accepting (3); that is, (MPC) warrants her in holding a belief that epistemic modesty prohibits. If our sense of modesty prevails, then a counterexample to (MPC) has been described. A situation that is epistemically possible (although perhaps rarely or never actual) provides an instance of a failure of (MPC).¹³ Thus, (MPC) cannot be considered a sound epistemic principle.

There are other putative counterexamples to (MPC) in the literature, of course. Many have a structure similar to that of the preface/fallibility scenario. First, a list of n statements, each of which one is apparently justified in believing, is specified; there seems no reason to prefer any one statement on the list to any other. Then counterevidence is provided that there is (at least) one false statement on the list, with no indication as to which statement it is that is false. This provides the basis for an apparent counterexample to (MPC) in a variety of ways. For instance, the set consisting of the first $n-1$ statements on the list, plus the claim that one of the statements on the list is false, implies that the n th statement is false; but it seems absurd to suppose that one is justified in believing that the n th statement is false.¹⁴

¹² See Foley, *op. cit.*, and Klein, *op. cit.* This is the reason that, as noted above, the fallibility argument cannot move directly from the claim that there is error in my present justified beliefs to the assertion that $\sim(P_1 \& \dots \& P_n)$. Simon Evnine concurs with the logical point in 'Believing Conjunctions', *Synthese* 118, 1999, 203.

¹³ The same sort of argument could of course be constructed in terms of the preface example, although it is not clear there is any advantage to so doing. For one thing, there is some controversy as to whether an author is justified in believing every statement in her book. The fallibility argument, on the other hand, is concerned with one's present justified beliefs, whatever they may be.

¹⁴ See Hawthorne, *op. cit.*, 183-84.

A general difficulty with examples of this sort, examples having the preface/fallibility structure, is that they involve two distinct lines of evidence: the evidence for the statements on the list, and the counterevidence that (at least) one statement on the list is false. Thus there is need to adjudicate competing claims: perhaps the counterevidence outweighs the evidence for the statements on the list, so that none of those statements is justified; or perhaps the evidence for the statements on the list is sufficiently strong to outweigh the counterevidence. Some further argument as to how these distinct bodies of evidence combine is necessary to achieve a clear counterexample to (MPC). This is not an issue with the epistemic modesty example, of course, since it has an entirely different structure.

Clearly, (MPC) is the stronger of the two closure principles, and is thus more vulnerable; (SPC) is untouched by this discussion. But a more sweeping attack on closure is found in the context of the lottery paradox, to which I now turn.

III

The lottery paradox rests on one key philosophical premise, the Principle of High Probability, which can be stated as follows:

(HP) There is a number n ($n < 1$) such that if P has probability of at least n for S , then S is justified in believing P .

Suppose a fair lottery with exactly 1,000 tickets, of which one will be drawn. The probability that ticket 1 will lose is .999. Given (HP), it seems reasonable to think that a probability of .999 is sufficient to warrant belief. So I am entitled to believe T_1 (Ticket 1 will lose). Similarly for T_2 , T_3 , and so on; I am entitled to believe each T_i . But I am also entitled to believe that exactly one ticket will win. So I am justified in believing each of:

$$T_1, \dots, T_{1,000}, \sim(T_1 \& T_2 \dots \& T_{1,000})$$

But this is an inconsistent set. Thus I am justified in holding inconsistent beliefs.

The success or failure of this second argument for rational inconsistency rides on the issue whether the T 's are justified or not. Of course, I can believe that it is *highly likely* that ticket x won't win; but am I entitled to believe, simply, ticket x won't win? Note that there is an apparent difference in my cognitive attitude to ticket x losing in lottery L_1 , after I read in the newspaper that some other ticket won, as compared to my attitude to ticket x losing in another lottery L_2 , before the draw has taken place. But in what can this difference consist? One plausible account is that I believe ticket x is a loser in the first case; while in the second, I believe just that it is highly probable that ticket x is a loser.¹⁵

¹⁵ It is generally conceded that one does not know lottery statements on the basis of the size of the lottery. So another suggestion is that I know in the first case, but only believe in the second. But the difference between knowing and believing does not appear to be a difference in *cognitive attitude*. Note also that the example can be structured so that the probability of ticket x losing is the same in each case (newspapers are, after all, fallible); thus the difference noted is not a function of different probabilities.

While suggestive, the above is hardly definitive. Recently, however, there has been a flurry of debate on the epistemic status of lottery statements, in the course of which examples have been produced that may have significance for the paradox and for closure.¹⁶ The focus of the debate has been knowledge and lottery statements; here it will be transposed to justification and lottery statements.

Suppose that Bernard owes a considerable sum of money. He has no investments, and no sources of income other than his salary. After a thorough review of his financial situation, he believes, justifiably, that he will not be able to pay off his debt in the next month. Surely, we all have beliefs of this sort, beliefs that are fully justified. As it happens, Bernard also owns a ticket in a Super-Lotto in which the draw has yet to occur. Bernard is aware that winning the prize in the Super-Lotto would make him a fabulously wealthy man; paying off his debt next month would be a triviality. Does it not follow that Bernard must be warranted in believing that he will not win the lottery? Do we not have to grant this in order to make sense of the fact that Bernard is justified in believing that he will not be able to pay off his debt next month? But if Bernard is entitled to believe he will not win the lottery, then surely the T_i s in any lottery are warranted.

This and related examples might seem to clinch the case for rational belief in the T_i s, and thereby for the possibility of rational inconsistency, since it is unquestionable that we sometimes justifiably believe that we will not be able to pay off a debt (or afford a trip to Bermuda) despite knowing that we have a ticket in a lottery in which the draw has not yet occurred. It would be entirely premature, however, to concede the point; for it is possible to make a similar case that belief in the T_i s is *not* warranted.

The argument might run as follows. Suppose a terrorist group decides that a certain public official must be assassinated. One member of the group is to do the deed, but there are compelling reasons why the agent should be selected randomly from within the group. A lottery, whose results will be kept secret, is held to determine which person will be assigned the task; the person holding the 'winning' ticket will do the deed. The assassination, we may suppose, takes place according to plan. Now consider Max, a typical member of the terrorist group. Surely, we are not entitled to believe, without independent evidence, that Max is not the assassin.¹⁷ But if Max did not win the lottery, then Max is not the assassin. We know that losing the lottery is a sufficient condition for not being the killer. Does it not follow, then, that we are not entitled to believe that Max lost the lottery? If so, surely the T_i s in any given lottery are not warranted on purely statistical grounds.

¹⁶ See John Hawthorne, *op. cit.*, Keith DeRose, 'Knowledge, Assertion and Lotteries', *Australasian Journal of Philosophy* 74 (1996), 568-80; Jonathan Vogel, 'Are There Counter-Examples to the Closure Principle?' in *Doubting*, M.D. Roth and G. Ross (eds.), (Dordrecht: Kluwer, 1990), 13-27.

¹⁷ I assume throughout that there is no such independent evidence.

So here are two highly intuitive, structurally similar arguments for opposing conclusions. Are we, or are we not, justified in believing the T_i s on purely statistical grounds? We need to examine the two cases more closely.¹⁸

The key statements in the first example are:

- (A) (1) Bernard will not be able to pay off his debt next month.
- (2) If Bernard wins the lottery, then he will be able to pay off his debt next month.
- (3) Bernard will not win the lottery.

The corresponding statements in the second case are:

- (B) (1*) Max did not win the lottery.
- (2*) If Max did not win the lottery, then he was not the assassin.
- (3*) Max was not the assassin.

Those who accept the argument of the lottery paradox (call them ‘the reformers’) can argue that since we are entitled to accept (1), we must also be entitled to accept (3). Those who do not sanction rational inconsistency (‘the traditionalists’), on the other hand, can maintain that since (3*) is not warranted, neither is (1*).

Consider (A) and (B) as arguments, with (3) and (3*) as conclusions, respectively. A key difference between the two arguments is that one has a lottery statement as conclusion, while the other has a lottery statement as a premise. But, clearly, both arguments are valid. The reformer argues that since (A) is valid, and the premises are justified, then so must be the conclusion. But this presupposes a closure principle, specifically (MPC). The traditionalist claims, on the other hand, that since argument (B) is valid, and (2*) is warranted while (3*) is not, it must be the case that (1*) is not warranted. Again, (MPC) is assumed.

But it cannot be that both the reformer and the traditionalist are right. It cannot be that (3) is warranted, while (1*) is not, for we have exactly the same sort of statistical evidence for each (we can even make the size of the lotteries the same). Here it is important to be clear on a point discussed earlier: That we are justified in believing statements that imply P does not guarantee that those statements can serve as our justification for P. For instance, we believe that (i) Barbara will be in New York on Wednesday (based on her announced intention to travel there for an interview); we also believe that (ii) if Barbara were to win the lottery, she would stay in Toronto on Wednesday to receive the prize. But we cannot, in this situation, appeal to (i) and (ii) to justify (iii) Barbara will not win the lottery. Similarly, in argument (A), premises (1) and (2) do not, and cannot, constitute our evidence for (3); we cannot justify our belief that Bernard won’t win the lottery on the basis of our belief that he won’t be able to pay his debt next month, and winning the lottery would enable him to do so. For one thing, our evidence for (1) and (2), facts about Bernard’s current financial situation, and about the size of the lottery award, do not constitute evidence that Bernard will not win the lottery. But a more decisive consideration is that our being justified in believing Bernard won’t be able to pay his debt next month is

¹⁸ A version of the argument that follows can also be found in *Paradox*; here it is given a fuller treatment, and the final conclusion is not exactly the same.

contingent on the epistemic status of (3) for us: we must either be justified in believing (3), or justified in ignoring the possibility that (3) is false. So we can hardly appeal to (1) and (2) to justify (3). Consequently, the only evidence we have which might serve to justify (3) is the statistical data concerning the lottery.

(An aside: It is worth noting that this feature of the examples does not persist through all variations. If we change the Barbara/New York example so that the relevant events are in the past, then we can have a different sort of evidence for the first premise. Thus we might have: (i*) Barbara was in New York on Wednesday (based on our having seen her there) (ii*) if Barbara had won the lottery, she would have stayed in Toronto on Wednesday (iii*) Barbara didn't win the lottery. Here justified belief in (i*) and (ii*) can provide the grounds for (iii*). For in the past tense version in which we see Barbara in New York, the justification for (i*) is not contingent on the epistemic status of (iii*.)

To return to the argument. In the situation envisioned, there is the same evidence for (3) and (1*), that is, the statistical information. So either both (3) and (1*) are warranted, or neither is. Let us consider each side of the dilemma. If (3) and (1*) are both warranted, it still seems clear that (3*) is not; and the only way to make room for this fact is to deny (MPC). Suppose, on the other hand, that neither (3) nor (1*) is warranted. Then the only way to accommodate the fact that (3) is not warranted is to deny (MPC). The conclusion that (MPC) does not hold seems inescapable.

What is worse, (SPC) also seems vulnerable to this argument. Start with the dilemma: either (3) and (1*) are both justified, or neither is. Suppose the first possibility holds; then (1*) and (2*) are each individually justified. If the conjunction of (1*) and (2*) is also justified, as seems highly plausible, then argument (B) can be restated as one-premise argument and, as such, provides a counter-example to (SPC). Suppose, on the other hand, that neither (3) nor (1*) is individually warranted. In argument (A), (1) and (2) are each individually justified. If the conjunction of (1) and (2) is warranted, then argument (A) can be restated as a one-premise argument in which the premise is warranted, and the conclusion is not. Again, there is an apparent counter-example to (SPC).

Whether the argument has repercussions for (SPC) depends on whether it is plausible to maintain that one of the conjunctions is not justified, despite the fact that each of its conjuncts is. Clearly, this is not guaranteed even assuming (CP) fails. The most plausible counterexamples to (CP) are lengthy, informative and highly detailed conjunctions that might be regarded as having low probability. But to protect (SPC) from the argument above requires that we deny that a very modest conjunction is justified: either the conjunction (1)&(2) (if the lottery statements are not justified) or (1*)&(2*) (if the lottery statements are justified). It is hard to see why such modest conjunctions would not be justified. It appears, then, that if the argument succeeds against (MPC), it also overturns (SPC).

The key step in the argument is the dilemma:

Either (3) and (1*) are both justified or neither is.

But does this dilemma really hold? One might object to it on the grounds that justification is not straightforwardly a function of the strength of the evidence; rather, justification is context-sensitive in that the degree of evidence required for justification varies with the context. One possibility worth considering is that the standards of evidence are a function of the seriousness or significance of the conclusion inferred from the evidence. If this is so, then it can plausibly be argued that the standards of justification will differ in the two cases, the debtor and the assassin; specifically, more will be required in the case of (1*) because the question of whether Max won the lottery is known to be directly relevant to whether Max is the assassin – a highly significant issue.

Is it possible to escape the dilemma in this way? It has been suggested that (3) can be warranted while (1*) is not, despite our having the same evidence for each. The grounds for (3) are that exactly one of the n tickets in the lottery will win; so (3) can be warranted only if such purely statistical information can constitute good evidence. But if this is so, the larger the lottery, the stronger the evidence for a given lottery statement. Consequently, the strength of the evidence for (1*) can be increased just by increasing the size of the lottery. So we should be able to achieve whatever level of evidence is necessary to justify (1*) simply by adjusting the size of the lottery appropriately.¹⁹ We can maintain the parallel between (3) and (1*), then, even though we grant that (1*) requires stronger evidence, just by increasing the size of the lottery in the assassin case. Thus, once the size of the lotteries has been adjusted, there is no longer any basis for maintaining that (3) can be justified while (1*) is not. Finally, note that increasing the size of the lottery has no significant impact on the Max example; it remains clear that (3*) is not warranted.

There is still a further move, however, for those who try to evade the dilemma on the grounds that justification is context-sensitive. The proposal we have been considering, it might now be suggested, must be further refined as follows: (1*) requires stronger evidence in the sense that it requires a *better kind of evidence* than purely statistical evidence. (The testimony of a reliable eyewitness, for instance, might be a better kind of evidence.) The suggestion is, then, that certain kinds of evidence can, in principle, provide a level of confirmation that purely statistical data, however strong, cannot.

The full significance of this proposal (call it '(RP)') emerges more clearly if we distinguish explicitly between epistemic probability, understood as the degree of confirmation of a statement, and statistical probability, which is a matter of the relative frequency with which a given event occurs in a certain reference class. Suppose an epistemic probability of n is sufficient to justify (1*) in the Max story. According to (RP), purely statistical data, however strong, does not suffice; so a *statistical probability* of n does not yield an epistemic probability of n .

¹⁹ This is denied in the next suggestion as to how to overturn the dilemma.

The argument of the lottery paradox, however, moves implicitly from a statistical probability of α to an epistemic probability of α . According to the current proposal, this move is illegitimate. So a consequence of the revised proposal, then, is that the implicit reasoning of the lottery is fallacious, and the lottery argument, in its usual formulation, does not go through.

What is the impact of the revised proposal on the case against closure? On the negative side, (RP) seems *ad hoc*, seems specifically designed to stop the anti-closure argument. Why should it be the case that statistical probability is a source of epistemic probability, but yet there is an upper limit on the degree of confirmation it can yield? If statistical data can never yield a degree of confirmation of n , perhaps this is better explained by maintaining that statistical data alone can never be a source of epistemic probability²⁰ – in which case, neither (3) nor (1*) is warranted, and the dilemma holds.

I am sceptical about (RP). But I do not see a clear refutation of the proposal; nor can I demonstrate that a theory according to which statistical probability is simply not a source of confirmation is preferable to (RP). In fact, there is some plausibility to the notion that one kind of evidence can, in principle, achieve higher degrees of confirmation than another.²¹ Some might take this to be true, for instance, of direct perceptual evidence as compared to the evidence of memory. Nor does the proposal conflict with the idea that increasingly higher statistical probability yields increasingly higher epistemic probability; for it may be that the degree of confirmation provided by statistical data continues to increase, but approaches a limit m ($m < n$, where n is the degree of confirmation required by (1*)).

The upshot is that, insofar as (RP) is a live option, the present argument against closure must be revised so that there is no disparity between the two cases, so that they involve issues of roughly equal significance. There are two options here: make the example in (A) weightier, or make the example in (B) less weighty. The first option is the simplest: we add to the Bernard story that he will be hunted by the mob if he cannot pay his debts in the next month. Now both (3) and (1*) are connected to highly significant issues; and the premises of (A) are still warranted. To take the second option, we can revise the Max story so that the issues are not as momentous. Suppose Max belongs to a group of teenagers who want to establish a reputation for being cool. They decide to commit a minor act of vandalism: scratch the door of a car with a key. A lottery is held to determine who in the group will do it, and the ‘winner’ of the lottery does the deed. Now neither (3) nor (1*) is connected to momentous issues; and it seems clear that the conclusion of (B) is not warranted. Whichever of these two options for revising the examples is chosen, the dilemma seems to hold.

²⁰ Dana Nelkin, ‘The Lottery Paradox, Knowledge and Rationality’, *The Philosophical Review* 109, 2002, 373-409, makes a case for the view that high statistical probability does not suffice for warranted belief.

²¹ Of course, ‘kind of evidence’ must be construed appropriately.

I have considered above the objections to the dilemma that seem most plausible. Others may find further avenues to explore in assessing the dilemma and its implications for closure. Two points, however, should be kept in mind: First, any approach designed to save closure must grant that lottery statements can be justified on the basis of purely statistical evidence; otherwise (A) provides a counterexample to closure. Second, as long as the only ground for belief is the statistical information concerning the lottery, (3*) does not seem to be justified.

In sum: Even if the arguments for the rationality of inconsistency fail, they suggest independent grounds for the denial of closure. Consideration of the preface and fallibility arguments exposes a requirement of epistemic modesty that conflicts with the dictates of (MPC). Reflection on the lottery scenario, on the other hand, gives rise to a dilemma in which each alternative seems to provide a counterexample to (SPC). Whether or not rational inconsistency is possible, there is a case against closure.²²

²² A version of this paper was read at the Fifth European Congress for Analytic Philosophy in Lisbon, August 2005.