

The lottery paradox

The “lottery paradox” is a kind of skeptical argument: that is, it is a kind of argument designed to show that we do not know many of the things we ordinarily take ourselves to know. One way of presenting the paradox is based on the following plausible claim:

If I know that P, and know that if P, then Q, I am thereby in a position to know that Q.

We generate cases of the paradox by substituting in for “P” some claim which we ordinarily take ourselves to know, and substitute in for “Q” some claim which follows from the claim substituted in for “P” which we take ourselves not to be in a position to know.

Here are some examples from Hawthorne’s *Knowledge and Lotteries*:

“... many normal people of modest means will be willing, under normal circumstances, to judge that they know that they will not have enough money to go on an African safari in the near future. And under normal circumstances, their conversational partners will be willing to accept that judgement as correct.

However ... [w]e do not suppose that people know in advance of a lottery drawing whether they will win or lose. But what is going on here? The proposition that the person will not have enough money to go on an African safari this year entails that he will not win a major prize in a lottery. If the person knows the former, then isn’t he at least in position to know the latter by performing a simple deduction?”

Here, we have the following claims filled in to the schema above:

P = I will not have enough money to go on an African safari next year.

Q = I will not win a major prize in a lottery for the rest of this year.

We are inclined to say that someone can know the truth of P, but that this knowledge does not put him in a position to know the truth of Q. But this is puzzling, since Q follows from P — and the person might know this.

Other examples of the same sort are easy to generate, and needn’t involve lotteries:

The “lottery paradox” is a kind of skeptical argument: that is, it is a kind of argument designed to show that we do not know many of the things we ordinarily take ourselves to know. One way of presenting the paradox is based on the following plausible claim:

If I know that P, and know that if P, then Q, I am thereby in a position to know that Q.

We generate cases of the paradox by substituting in for “P” some claim which we ordinarily take ourselves to know, and substitute in for “Q” some claim which follows from the claim substituted in for “P” which we take ourselves not to be in a position to know.

Other examples of the same sort are easy to generate, and needn’t involve lotteries:

“I am inclined to think that I know that I will be living in Syracuse for part of this summer. But once the question arises, I am not inclined to think that I know whether or not I will be one of the unlucky people who, despite being apparently healthy, will suffer a fatal heart attack in the next week.”

“I am inclined to think that I know where my car is parked right now. But once the question arises, I am not inclined to think that I know whether or not I am one of the unlucky people whose car has been stolen during the last few hours.”

In these examples, we have some proposition — following Hawthorne, let’s call it an **ordinary proposition** — which is some proposition of the kind of which we usually take ourselves to have unproblematic knowledge, and some other proposition — the **lottery proposition** — which is entailed by the ordinary proposition but which we do **not** usually take ourselves to know.

So, in responding to these cases, it looks like we have three choices:

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

The problem — and the reason why this is an example of a paradox — is that none of these options seems very appealing.

In these examples, we have some proposition — following Hawthorne, let's call it an **ordinary proposition** — which is some proposition of the kind of which we usually take ourselves to have unproblematic knowledge, and some other proposition — the **lottery proposition** — which is entailed by the ordinary proposition but which we do **not** usually take ourselves to know.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

The problem — and the reason why this is an example of a paradox — is that none of these options seems very appealing.

Let's consider the first option — which is a kind of skepticism — first.

It follows from this option that, almost always, when we say that we know something, what we say is false. In reality, we know hardly anything — after all, can't we usually come up with some sort of relevant lottery proposition?

This is quite hard to believe. It also seems to conflict with certain plausible-seeming claims about knowledge, like the following claim about assertion:

One should assert only what one knows.

It would follow from this + skepticism that we should hardly ever assert anything. But surely this is not right.

Consider also:

In deliberating about what to do, you should only use known propositions as premises.

Surely the paradox should not convince us that almost all of our practical reasoning is flawed.

It seems as though skepticism should only be an option of last resort.

The simplest alternative to skepticism is a simple anti-skepticism which says that the lottery paradox just shows more than we thought we knew: in general, according to this sort of view — which we can call **dogmatism** — we know the lottery proposition.

In these examples, we have some proposition — following Hawthorne, let's call it an **ordinary proposition** — which is some proposition of the kind of which we usually take ourselves to have unproblematic knowledge, and some other proposition — the **lottery proposition** — which is entailed by the ordinary proposition but which we do **not** usually take ourselves to know.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

The simplest alternative to skepticism is a simple anti-skepticism which says that the lottery paradox just shows more than we thought we knew: in general, according to this sort of view — which we can call **dogmatism** — we know the lottery proposition.

This faces serious problems. If we say that we know that, for example, person 1 will not win the lottery, then it seems plausible, by parity of reasoning, to say the same thing about person 2, person 3, etc. But then it seems that we can know of an arbitrarily large percentage of the ticket holders that they will not win the lottery. But this seems absurd.

Further absurdities result if we endorse the following variant on the closure principle discussed above:

Multi-premise closure (MPC)

Knowing propositions P_1, P_2, \dots, P_n , and knowing that $P_1 \& P_2 \& \dots P_n$ implies Q, is enough to know Q

Sometimes the Indiana Lottery has no winner. If dogmatism is correct, then parity reasoning shows that before such a lottery I can know of each ticket-holder that she will not win. But then (if I also know that those are all of the ticket holders) it follows from MPC that I know that no one will win. But that is surely crazy. What if the Lottery has at least one winner 99% of the time?

However, the dogmatist might try to dodge these especially bad consequences by denying MPC; and a plausible argument can be made that this is not such a bad idea.

Consider the following principle, which seems closely related to MPC: If I am justified (rational) in believing propositions P_1, P_2, \dots, P_n , and knowing that $P_1 \& P_2 \& \dots P_n$ implies Q, then I am justified (rational) in believing Q.

This principle is false, as is shown by the paradox of the preface:

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

The simplest alternative to skepticism is a simple anti-skepticism which says that the lottery paradox just shows more than we thought we knew: in general, according to this sort of view — which we can call **dogmatism** — we know the lottery proposition.

Multi-premise closure (MPC)

Knowing propositions P1, P2, ... Pn, and knowing that P1 & P2 & ... Pn implies Q, is enough to know Q

Sometimes the Indiana Lottery has no winner. If dogmatism is correct, then parity reasoning shows that before such a lottery I can know of each ticket-holder that she will not win. But then (if I also know that those are all of the ticket holders) it follows from MPC that I know that no one will win. But that is surely crazy. What if the Lottery has at least one winner 99% of the time?

However, the dogmatist might try to dodge these especially bad consequences by denying MPC; and a plausible argument can be made that this is not such a bad idea.

Consider the following principle, which seems closely related to MPC: If I am justified (rational) in believing propositions P1, P2, ... Pn, and knowing that P1 & P2 & ... Pn implies Q, then I am justified (rational) in believing Q.

This principle is false, as is shown by the paradox of the preface:

The paradox of the preface

I have just finished writing my masterpiece. Given my intense desire to write a masterpiece, I have checked and re-checked every sentence of the book; I am justified in believing every one of them. Nonetheless, I have an appropriate level of humility, and realize that it is very unlikely that there would be no mistakes in a 10,000 page book about paradoxes. So, in the preface of the book, I write: "Though many are to be thanked for the insights contained in this book, only the author is to be blamed for its mistakes." This sentence expresses, among other things, my justified belief in the claim that there is at least one false sentence in the book.

This seems to be a situation in which I am justified in believing every member of a set of propositions which are nonetheless jointly inconsistent. But this shows that the multi-premise closure principle to do with justified belief must be false.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

Multi-premise closure (MPC)

Knowing propositions P_1, P_2, \dots, P_n , and knowing that $P_1 \& P_2 \& \dots P_n$ implies Q, is enough to know Q

Sometimes the Indiana Lottery has no winner. If dogmatism is correct, then parity reasoning shows that before such a lottery I can know of each ticket-holder that she will not win. But then (if I also know that those are all of the ticket holders) it follows from MPC that I know that no one will win. But that is surely crazy. What if the Lottery has at least one winner 99% of the time?

Consider the following principle, which seems closely related to MPC: If I am justified (rational) in believing propositions P_1, P_2, \dots, P_n , and knowing that $P_1 \& P_2 \& \dots P_n$ implies Q, then I am justified (rational) in believing Q.

This principle is false, as is shown by the paradox of the preface:

The paradox of the preface

I have just finished writing my masterpiece. Given my intense desire to write a masterpiece, I have checked and re-checked every sentence of the book; I am justified in believing every one of them. Nonetheless, I have an appropriate level of humility, and realize that it is very unlikely that there would be no mistakes in a 10,000 page book about paradoxes. So, in the preface of the book, I write: "Though many are to be thanked for the insights contained in this book, only the author is to be blamed for its mistakes." This sentence expresses, among other things, my justified belief in the claim that there is at least one false sentence in the book.

This seems to be a situation in which I am justified in believing every member of a set of propositions which are nonetheless jointly inconsistent. But this shows that the multi-premise closure principle to do with justified belief must be false.

This argument cannot be straightforwardly adapted to show the falsity of MPC - why?

Nonetheless, one might think that the Preface provides some reason for doubting MPC, since knowledge seems to be closely related to justified, rational belief.

Is this enough to make dogmatism a plausible position? Is it really true that we can know of each of the 999 out of 1000 people who will not win the lottery, that they will not win?

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

Multi-premise closure (MPC)

Knowing propositions P1, P2, ... Pn, and knowing that P1 & P2 & ... Pn implies Q, is enough to know Q

Nonetheless, one might think that the Preface provides some reason for doubting MPC, since knowledge seems to be closely related to justified, rational belief.

Is this enough to make dogmatism a plausible position? Is it really true that we can know of each of the 999 out of 1000 people who will not win the lottery, that they will not win?

Other problems for the dogmatist stem from the following plausible claims about knowledge:

- (1) If I know P, then I can assert P.
- (2) If I know P, then I can use P as a premise in practical reasoning - reasoning about what I should do.

But now consider lottery propositions. If someone says, "There is a 1/1000 chance that my ticket will win", can I simply reply by asserting: "Your ticket will not win."

And if someone offers me 1¢ for my lottery ticket, it certainly does not seem that I can reason as follows: "I know that I will not win; hence, any amount of money for this ticket is better than the \$0 I will win; so I should sell the ticket for 1¢."

So dogmatism, it seems, faces serious problems.

One might, on the contrary, think that our doubts about MPC suggest a different solution to the lottery paradox than that suggested by the dogmatist: perhaps they suggest that we should pursue the third option above, and deny the (single premise) closure principle.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

One might, on the contrary, think that our doubts about MPC suggest a different solution to the lottery paradox than that suggested by the dogmatist: perhaps they suggest that we should pursue the third option above, and deny the (single premise) closure principle.

This way of going seems to let us say exactly what we want: we can say that we know the ordinary proposition, and avoid saying that we know the lottery proposition; and the only cost is that we have to say that sometimes you can know P and know that it implies Q without saying that we also know Q. What is so bad about that?

One cost is that it really seems as though we can extend the reach of our knowledge by seeing the logical consequences of things we know. Denying closure is just denying that this is always a way in which we can extend the reach of our knowledge; it is to say that sometimes this method of belief formation, even if based on a genuine piece of knowledge, takes us to a belief which falls short of knowledge.

A second problem with this response to the lottery paradox is that it seems to follow from some other plausible principles:

Equivalence

If I know P, and know that P and Q are equivalent - in the sense that, necessarily, one is true if and only if the other is - then I can also know Q.

Distribution of knowledge over conjunction

If I know that P and Q, then I also know P, and know Q.

To show that these imply closure — if we assume that the relevant knower has a minimum of logical sophistication — we suppose that I know P and know that P implies Q, and argue from these assumptions to the conclusion that I also know Q, as follows:

If P implies Q, then P is equivalent to the conjunction P & Q. If I know that P implies Q, I'll therefore be in a position to know that P and P & Q are equivalent. Then, by Equivalence, it follows that I know P & Q. Then, by the distribution of knowledge over conjunction, it follows that I know Q.

So it is hard to deny closure without also denying one of Equivalence and Distribution — and these seem pretty tough to deny.

Let's now consider a different response to the paradox which does not fit neatly into any of the three categories of response listed above.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

Let's now consider a different response to the paradox which does not fit neatly into any of the three categories of response listed above.

This sort of response begins with a view about the word “knows” called **contextualism**. This is the view that “knows” is, in an important respect, like words like “I”, “now”, and “here”: it is a word which refers to different things depending on the context, or situation, in which it is used.

In the case of “I”, how this works is pretty obvious: it is used to refer to whatever is the speaker of the context of utterance. How might this idea work for “knows”? Contextualists differ in their answers to this question, but one plausible view of this is that someone knows P if and only if they believe P, P is true, and their evidence is **good enough for the context**.

What does it mean for your evidence to be good enough for the context? One thought is that **your evidence must rule out all the scenarios which would both make the belief false and which are relevant to the conversation of the context**.

So, for example, imagine that you ask me if I know where my car is parked, and I say: “Yes, it is in front of Legends.” Here what I say is true, because the possibilities relevant to our conversation are ones in which I forget where I left the car — and I do (let's suppose) have evidence which rules these possibilities out.

But now imagine that someone else joins our conversation, and says: “You know, it might not be there — cars are stolen from that lot all of the time. What they have done is that they have changed the scenarios which are relevant to the conversation: more specifically, they have added some relevant scenarios, in which my car has been moved from the place where I left it; and I, not having been back to my car since the morning, have no evidence which rules these out. Hence after this remark I can no longer truly say, “I know where my car is.”

This is why I cannot in general know the lottery proposition: for simply introducing that proposition into the discussion makes relevant possibilities which I cannot rule out, and these possibilities take away my knowledge of the ordinary proposition as well — hence we are able to say that I (usually) know the ordinary propositions, that I never know the lottery propositions, but that all of this is consistent with closure.

This all sounds pretty nice; and it is one of the many reasons that many philosophers are attracted to contextualism.

But this view also has some unattractive aspects.

- ➡ Deny that we know the ordinary proposition.
- ➡ Concede that we know the lottery proposition.
- ➡ Deny **closure**: that knowing P, and knowing that P implies Q, is enough to know Q.

This sort of response begins with a view about the word “knows” called **contextualism**. This is the view that “knows” is, in an important respect, like words like “I”, “now”, and “here”: it is a word which refers to different things depending on the context, or situation, in which it is used.

In the case of “I”, how this works is pretty obvious: it is used to refer to whatever is the speaker of the context of utterance. How might this idea work for “knows”? Contextualists differ in their answers to this question, but one plausible view of this is that someone knows P if and only if they believe P, P is true, and their evidence is **good enough for the context**.

What does it mean for your evidence to be good enough for the context? One thought is that **your evidence must rule out all the scenarios which would both make the belief false and which are relevant to the conversation of the context**.

This is why I cannot in general know the lottery proposition: for simply introducing that proposition into the discussion makes relevant possibilities which I cannot rule out, and these possibilities take away my knowledge of the ordinary proposition as well — hence we are able to say that I (usually) know the ordinary propositions, that I never know the lottery propositions, but that all of this is consistent with closure.

This all sounds pretty nice; and it is one of the many reasons that many philosophers are attracted to contextualism.

But this view also has some unattractive aspects.

One is that it seems to conflict with the principles about assertion and practical reasoning mentioned above. It follows from contextualism that what you are licensed to use as a premise in practical reasoning — reasoning about what you ought to do — depends on which possibilities are salient in the conversation. But think about the example of the car in the parking lot: should I reason any differently about how to get home after that possibility is introduced in the conversation?

A second, more basic worry about contextualism is that it makes knowledge of a proposition P dependent on factors which, intuitively, have nothing to do with factors relating to the truth of P: plainly, the truth of the claim that my car is in the Legends parking lot has nothing to do with whether or not someone raised the spectre of car theft in a conversation in Malloy Hall — nor need such a topic being raised give me reason to raise, or lower, the probability I assign to the proposition (though of course sometimes it may — if, for example, I did not know about the rate of car theft around here). So how can whether I know this proposition depend on such factors

It is clear that the lottery paradox provides an important challenge to a wide range of views about knowledge. One thing you may want to think about is: how could one of the views about knowledge we’ve discussed be modified to avoid the problems we’ve found with each?