

- Moore, G.E. 1942. A reply to my critics. In *The Philosophy of G.E. Moore*, ed. P. Schilpp. La Salle, Illinois: Open Court.
- Sutton, J. 2007. *Without Justification*. Cambridge, MA: MIT Press.
- Williams, J.N. 2004. Moore's paradoxes, Evans's principle, and self-knowledge. *Analysis* 64: 348–53.

How to solve the hardest logic puzzle ever in two questions

GABRIEL UZQUIANO

Rabern and Rabern (2008) have noted the need to modify 'the hardest logic puzzle ever' as presented in Boolos 1996 in order to avoid trivialization. Their paper ends with a two-question solution to the original puzzle, which does not carry over to the amended puzzle. The purpose of this note is to offer a two-question solution to the latter puzzle, which is, after all, the one with a claim to being the hardest logic puzzle ever.

Recall, first, Boolos's statement of the puzzle:

Three gods A, B and C are called, in some order, True, False and Random. True always speaks truly, False always speaks falsely, but whether Random speaks truly or falsely is a completely random matter. Your task is to determine the identities of A, B and C by asking three yes-no questions; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language, in which the words for 'yes' and 'no' are 'da' and 'ja', in some order. You don't know which word means which. (Boolos 1996: 62)

And remember his guidelines:

- It could be that some god gets asked more than one question.
- What the second question is, and to which god it's put, may depend on the answer to the first question.
- Whether Random speaks truly or not should be thought of as depending on the flip of a coin hidden in his brain: if the coin comes down heads, he speaks truly; if tails, falsely.
- Random will answer 'da' or 'ja' when asked any yes-no question. (Boolos 1996: 62)

Since the stipulation that Random should set out to speak truly or lie, albeit randomly, trivializes the puzzle, Rabern and Rabern (2008) have proposed

capturing the spirit of the puzzle by replacing the third stipulation above with the following:

- Whether Random answers ‘da’ or ‘ja’ should be thought of as depending on the flip of a coin hidden in his brain: if the coin comes down heads, he answers ‘ja’; if tails, ‘da’.¹

Rabern and Rabern (2008) provide a clever two-question solution to the original puzzle, which exploits the fact that there are both self-referential questions that no truth-teller can answer and self-referential questions that no liar can answer. So, if we assume that all of A, B and C set out to speak the truth or lie, albeit randomly in one case, then we can glean more information from their inability to answer certain self-referential yes-no questions. The question remains, however, whether we can give a two-question solution to the amended puzzle in which one of them answers ‘da’ or ‘ja’ randomly.²

In what follows, we offer a two-question solution to the amended puzzle. We proceed in two stages. First, we assume that A, B and C agree to answer our questions in plain English. We then show how to do without this assumption.

Assume the gods agree to answer our questions in plain English. Consider what happens when you ask True or False whether he and Random would give the same answer to the question of whether Dushanbe is in Kirghizia. Since Random would indeed answer in a truly random fashion, neither True nor False will be in a position to anticipate whether Random would answer ‘yes’ or ‘no’ to the question of whether Dushanbe is in Kirghizia – even if they are certain that Random will either answer ‘yes’ or ‘no’ to the question. But if they don’t know this, they cannot be sure of telling the truth or lying by means of a ‘yes’ or ‘no’ answer. So, they aren’t in a position to answer our question and should therefore remain silent. This is in contrast to Random, who will gladly answer ‘yes’ or ‘no’ to the question of whether he and True or False would give the same answer to the question of whether Dushanbe is in Kirghizia.³

- 1 This is (B3*) in Rabern and Rabern (2008: 107).
- 2 To belabour the obvious, we will not find a two-question solution to the amended puzzle unless we make allowance for questions that cannot be answered by some gods by means of ‘da’ or ‘ja’ consistently with their nature. We need to distinguish at least six possible combinations with respect to the identity of A, B and C – TFR, TRF, FTR, FRT, RTF and RFT – and a ‘yes’ or ‘no’ answer will not eliminate more than three combinations after the first question. Unfortunately, another ‘yes’ or ‘no’ question cannot be guaranteed to eliminate two of the last three live options.
- 3 What if we build into the puzzle the assumption that True and False have the oracular ability to anticipate Random’s answers even before they occur? Then, we could still use certain self-referential questions to obtain the same effect. More on this later.

We are now in a position to solve the amended puzzle in two questions.

(1) Directed to A:

Would you and B give the same answer to the question of whether Dushanbe is in Kirghizia?

There are three options:

- (a) A cannot answer the question. This will be the case if and only if A is either True or False and B is Random.
- (b) A answers 'yes' to the question. This is compatible with two cases:
 - (i) A is Random and has answered randomly.
 - (ii) A and B are True and False, in some order, in which case the answer 'yes' tells us that A is False and B is True. (True and False would answer differently to the question of whether Dushanbe is in Kirghizia. True would answer 'no' to this last question because Dushanbe is in Tajikistan and not Kirghizia; False would therefore answer 'yes' to the question.)
- (c) A answers 'no' to the question. This is compatible with two cases:
 - (i) A is Random and has answered randomly.
 - (ii) A and B are True and False, in some order, in which case the answer 'no' tells us that A is True and B is False.

Whatever the outcome of our first question, we have identified a god who isn't Random. If there is no answer, then B is Random; if A answers 'yes' or 'no', then we know that B isn't Random.

(2) Directed to one of A or B we now know not to be Random:

Would you and C give the same answer to the question of whether Dushanbe is in Kirghizia?

There are again three options:

- (a) There is no answer. This will be the case if and only if C is Random. To unveil the identities of the other two, note that since we now know that C is Random, we now know that A and B are True and False in some order. To find out who is who, we refer back to A's answer to question (1), which will by now tell us who of A and B is True and who is False. (This is because (b)(i) and (c)(i) above are no longer live options.)
- (b) The answer is 'yes'. This tells us that C isn't Random, and that our interlocutor is False. Therefore, C is True and the other god is Random.
- (c) The answer is 'no'. This tells us that C isn't Random, and that our interlocutor is True. Therefore, C is False and the other god is Random.

To turn our solution into a two-question solution to the amended puzzle, where A, B and C will only answer ‘da’ or ‘ja’ to our questions, replace (1) and (2) by the following:

(1') Directed to A:

Would you answer ‘ja’ to the question of whether you would answer with a word that means ‘yes’ in your language to the question of whether you and B would give the same answer to the question whether Dushanbe is in Kirghizia?

(2') Directed to one of A or B we now know not to be Random:

Would you answer ‘ja’ to the question of whether you would answer with a word that means ‘yes’ in your language to the question of whether you and C would give the same answer to the question whether Dushanbe is in Kirghizia?

It is straightforward to check that these questions fit the bill by exploiting the embedded question lemma in Rabern and Rabern (2008: 108), whereby we can often extract the correct answer to a question q by asking the following question:

Would you answer ‘ja’ to question q ?

True or False will answer ‘ja’ only if the correct answer to q is affirmative; and they will answer ‘da’ only if the correct answer to q is negative.

There is still a loose end. What if we stipulate that True and False have the oracular ability to predict Random’s answers even before the coin flip in Random’s brain? No matter. We can achieve similar effects by means of self-referential questions of the style employed in Rabern and Rabern 2008. To that end, we need only make the following observation. Consider the question:

(L) Would you answer ‘ja’ to the question whether you would answer ‘da’ to (L)?

Neither True nor False is in a position to answer this question consistently with his nature. For whatever their identity, they will be forced to answer ‘ja’ if and only if their answer is ‘da’, which they can’t do. This is of course in contrast to Random, who will randomly answer ‘ja’ or ‘da’ to (L).

But now note what happens when we direct the following question to A:

(1*) Would you answer ‘ja’ to the question whether either:

- (a) B isn’t Random and you are False, or
- (b) B is Random and you would answer ‘da’ to (1*)?

A will be able to answer question (1*) if and only if B isn’t Random. For if B is Random, then (1*) will collapse into the question of whether A would

answer ‘ja’ to the question of whether he would answer ‘da’ to (1*), which he cannot do if he is True or False. But A must be True or False if B is Random.

What if B isn’t Random? We need to distinguish two cases:

- A is Random, in which case he will answer ‘ja’ or ‘da’ randomly.
- A isn’t Random, in which case, by the Embedded Question Lemma, A will answer ‘ja’ only if the correct answer to (1*) is affirmative; and ‘da’ only if the correct answer to (1*) is negative. Since B is, by assumption, not Random, (b) will be false no matter who A is. If A is False, then, since B isn’t Random, (a) will be true and the correct answer to (1*) will be affirmative. If A is True, then the second conjunct of (a) will be false and the correct answer to (1*) will be negative – as both (a) and (b) are false. It follows that A will answer ‘ja’ to (1*) if and only if either (i) A is Random or (ii) A is False and B is True; and A will answer ‘da’ to (1*) if and only if either (i’) A is Random or (ii’) A is True and B is False.

We can now solve the hardest logic puzzle ever in two questions:
(1*) Directed to A:

Would you answer ‘ja’ to the question of whether either:

- (a) B isn’t Random and you are False, or
- (b) B is Random and you would answer ‘da’ to (1*)?

If A cannot answer (1*), then B is Random. Otherwise, if A answers ‘ja’, then either (i) A is Random or (ii) A is False and B is True; and if A answers ‘da’, then either (i’) A is Random or (ii’) A is True and B is False. Either way, we have identified one of A or B as a god who isn’t Random. On to our second question:

(2*) Directed to one of A or B we now know not to be Random:

Would you answer ‘ja’ to the question of whether either:

- (a) C isn’t Random and you are False, or
- (b) C is Random and you would answer ‘da’ to (2*)?

If there is no answer, then C is Random and we can read the identities of A and B off A’s answer to (1*). If our god answers ‘ja’ to (2*), then C isn’t Random and, since we know our god isn’t Random, he must be False; therefore, C is True and the other god Random. But if our god answers ‘da’ to (2*), then C isn’t Random and, since we know our god isn’t Random, he must be True; therefore, C is False and the other god Random.

One final comment: The key to two-question solutions to the puzzle is to learn how to glean information from a certain god’s inability to answer a question. The difference with respect to the two-question solution given by Rabern and Rabern 2008 to the original puzzle is that we don’t assume

Random to either tell the truth or lie, albeit randomly, in response to our questions; indeed, our solution trades on the completely random nature of his answers. Unfortunately, the general strategy outlined here suggests an even harder variation on the hardest puzzle ever:

- Three gods A, B and C are called, in some order, True, False and Random. True always speaks truly, False always speaks falsely, but whether Random speaks truly or falsely *or whether Random speaks at all* is a completely random matter. Your task is to determine the identities of A, B and C by asking three yes–no questions; each question must be put to exactly one god. The gods understand English, but will answer all questions in their own language, in which the words for ‘yes’ and ‘no’ are ‘da’ and ‘ja’, in some order. You don’t know which word means which.

And, furthermore:

- Whether Random answers ‘da’ or ‘ja’ or whether Random answers at all should be thought of as depending on the toss of a fair three-sided dice hidden in his brain: if the dice comes down 1, he doesn’t answer at all; if the dice comes down 2, he answers ‘da’; if 3, ‘ja’.

I, for one, don’t know how to solve this puzzle in two questions.⁴

Pembroke College
Oxford OX1 1DW, UK
gabriel.uzquiano@philosophy.ox.ac.uk

References

- Boolos, G. 1996. The hardest logic puzzle ever. *The Harvard Review of Philosophy* 6: 62–65.
- Rabern, B. and L. Rabern. 2008. A simple solution to the hardest logic puzzle ever. *Analysis* 68: 105–12.

4 I would like to thank Agustín Rayo and an anonymous referee for helpful comments. Thanks to Landon Rabern from whom I have learned that he and Brian Rabern have used a similar idea in an unpublished typescript. Special thanks to my father, the late Marcelino Uzquiano, for all his patience and encouragement.