**Why the hardest logic puzzle ever cannot be solved in less than three questions**

Gregory Wheeler and Pedro Barahona

Rabern and Rabern (2008) and Uzquiano (2010) have each presented increasingly harder versions of 'the hardest logic puzzle ever' (Boolos 1996), and each has provided a two-question solution to his predecessor's puzzle. But Uzquiano's puzzle is different from the original and different from Rabern and Rabern's in at least one important respect: it cannot be solved in less than three questions. In this note we solve Uzquiano's puzzle in three questions and show why there is no solution in two. Finally, to cement a tradition, we introduce a puzzle of our own.

Recall Uzquiano's puzzle and his guidelines for solving it.

> 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 manner. 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 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. (Uzquiano 2010: 44)

The guidelines:

1. It could be that some god gets asked more than one question.
2. What the second question is, and to which god it is put, may depend on the answer to the first question.
3. 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'.

What distinguishes these puzzles from each other are three different specifications for Random's behavior. Boolos (1996) allows for Random to speak truly or lie, albeit randomly, whereas Rabern and Rabern (2008) stipulate that Random answers 'da' and 'ja' randomly. Both Uzquiano's solution strategy as well as Rabern and Rabern's exploit a common trait in the first two puzzles, which is that there are yes/no questions that True cannot answer and yes/no questions that False cannot answer, but no question that Random will fail to answer. Uzquiano (2010) eliminates this particular asymmetry from his version of the puzzle by granting Random the option of remaining silent. This modification, need it be said, is what makes Uzquiano’s the hardest logic puzzle ever.

There are three parts to our solution. First, we assume that A, B, and C agree to answer our questions in English, and we show how to solve the puzzle in three questions. Next, we show how to solve the puzzle without this assumption. It turns out that most scenarios covered by our solution resolve the identities of the gods in two questions, but there is one case where information from a third question is necessary. Our final step is to show that there is always at least one stray case to scupper any two-question solution strategy.

***

We begin by observing that there are 6 state descriptions that correspond to the possible identities of the three gods.

| | | | |
|---|---|---|---|
| (P1) | A-True | B-False | C-Random |
| (P2) | A-True | B-Random | C-False |
| (P3) | A-False | B-True | C-Random |
| (P4) | A-False | B-Random | C-True |
| (P5) | A-Random | B-True | C-False |
| (P6) | A-Random | B-False | C-True |

The puzzle is solved just when one possibility remains, revealing the true identities of all three gods. Here is the first of our three questions.

(Q1) Directed to god A: Would you and B give the same answer to the question of whether Lisbon is south of London?

If B is Random, then god A must be either True or False, in which case A cannot answer and will remain silent. If A is Random, he can either answer or remain silent. If A is True and B is False, then A will answer ‘no’. Finally, if A is False and B is True, then A will answer ‘yes’. In more detail, we have the following possibilities.

If **A answers ‘yes’ to Q1**, then three possibilities remain. Either A is False, B is True, and C is Random, or A is Random.

| | | | |
|---|---|---|---|
| (P3) | A-False | B-True | C-Random |
| (P5) | A-Random | B-True | C-False |
| (P6) | A-Random | B-False | C-True |

If **A answers 'no' to Q1**, then three possibilities remain. Either A is True, B is False, and C is Random, or A is Random.

| | | | |
|---|---|---|---|
| (P1) | A-True | B-False | C-Random |
| (P5) | A-Random | B-True | C-False |
| (P6) | A-Random | B-False | C-True |

If **A gives no answer to Q1**, then four possibilities remain: Either A is Random or B is Random.

| | | | |
|---|---|---|---|
| (P2) | A-True | B-Random | C-False |
| (P4) | A-False | B-Random | C-True |
| (P5) | A-Random | B-True | C-False |
| (P6) | A-Random | B-False | C-True |

So, whatever the outcome of the first question, we may identify a god who is not Random. If A answers either 'yes' or 'no', then B is not Random, and if A remains silent, then C is not Random. Now, turn to the second question.

(Q2) Put to B or C we now know not to be Random: Would you and the god not questioned thus far give the same answer to the question of whether Lisbon is south of London?

The possible answers to (Q2) depend on how A answered (Q1).

If **A answered 'yes' to (Q1)**, then, given the open possibilities in this case, we know that B is not Random. So, the second question put to B is this: Would you and C give the same answer to the question of whether Lisbon is south of London? However B responds, we may determine the identify of each god since B answers 'yes' if and only if (P6) is actual, B answers 'no' if and only if (P5) is actual, and B is silent if and only if (P3) is actual.

If **A answers 'no' to (Q1)**, then, given the open possibilities in this case, we know that B is not Random. So, the same question is put to B, namely, Would you and C give the same answer to the question of whether Lisbon is south of London? Here again B answers 'yes' if and only if (P6) is actual, B answers 'no' if and only if (P5) is actual, and B is silent if and only if (P1) is actual.

If **A gives no answer to (Q1)**, then, given the open possibilities in this case, we know that C is not random. So, the second question is directed to C: Would you and B give the same answer to the question of whether Lisbon is south of London? Here C answers 'yes' if and only if (P5) is actual and C answers 'no' if and only if (P6) is actual, but if C is silent then either (P2) is actual or (P4) is actual.

So, two questions suffice to solve the puzzle *unless* one fails to elicit an answer to both (Q1) and (Q2). To resolve the uncertainty between (P2) and (P4) in this case, a third question is required.

(Q3) Put to A: Would you and C give the same answer to the question of whether Lisbon is south of London?

Since we know that B is Random, A answers 'yes' if and only if state (P4) is actual and A answers 'no' if and only if state (P2) is actual.

To turn this argument into a solution to Uzquiano's puzzle, where A, B, and C will only answer 'da' or 'ja' to our questions, if they answer at all, we appeal to the embedded question lemma in Rabern and Rabern (2008: 108):

(EQL) Would you answer 'ja' to question Q?

True or False will answer 'ja' only if the correct answer to Q is 'yes', and they will answer 'da' only if the correct answer to Q is 'no'. We then may replace (Q1), (Q2), and (Q3) by the following:

(Q1′) Directed to god 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 of whether Lisbon is south of London?

(Q2′) Put to one of B or C 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 the god not addressed by (Q1) give the same answer to the question of whether Lisbon is south of London?

(Q3′) Put 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 C give the same answer to the question of whether Lisbon is south of London?

Finally, suppose both True and False have the ability to predict Random's answer even before the coin lands in Random's brain. Fortunately, Uzquiano's construction for this scenario can be applied to our solution, too. Following Uzquiano (2010), we make use of the following observation:

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

Neither True nor False will answer (L), since each is required to answer ‘ja’ if and only if their answer is ‘da’, which is prohibited by their natures. However, whether Random answers ‘ja’ or ‘da’, or whether Random answers at all, is entirely a random matter. Now consider the following modification to our first question.

(Q1*) Put to A: Would you answer ‘ja’ to the question whether either:
(a) B is not Random and you are False, or
(b) B is Random and you would answer ‘da’ to (Q1*)?

A answers (Q1*) if and only if B is not Random. For if B is Random, then (a) is false by (EQL) and (b) is false by an analogue of (UL), since A is either True or False and thus is unable to answer ‘da’ to (Q1*). And if B is not Random, then there are two possibilities to consider.

- A is Random, in which case he answers ‘ja’, ‘da’, or remains silent, each with probability 1/3.
- A is not Random, so by (EQL), A will answer ‘ja’ only if the correct answer to (Q1*) is affirmative, and answer ‘da’ only if the correct answer to (Q1*) is negative. Since, by hypothesis, B is not Random, (b) is 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 (Q1*) will be affirmative. If A is True, then the second conjunct of (a) will be false and the correct answer to (Q1*) will be negative, since both (a) and (b) are false. So, A will answer ‘ja’ to (Q1*) if and only if either A is Random or A is False and B is True, and A will answer ‘da’ to (Q1*) if and only if either A is Random or A is True and B is False.

In sum, if A cannot answer (Q1*), then either A is Random or B is Random. If A answers at all, then either A is Random or C is Random. Either way, we have identified one of B or C as a god who is not Random. Turn to our second question.

(Q2*) Put to one of B or C we now know not to be Random: Would you answer ‘ja’ to the question whether either:
(c) the god not questioned thus far is not Random and you are False, or
(d) the god not questioned thus far is Random and you would answer ‘da’ to (Q1*)?

What we learn from (Q2*) depends on A's response to (Q1*). If A answered (Q1*), then we know that (Q2*) is directed to B. If B answers 'ja' to (Q2*), then (i) B answers 'ja' if and only if A is Random, B is False, and C is True; (ii) B answers 'da' if and only if A is Random, B is True, and C is False; and (iii) B is silent if and only if A is True, B is False, and C is Random. If instead A gives no answer to (Q1*), then we know that (Q2*) is directed to C and there are four possibilities to distinguish. If C answers 'ja', then A is Random, B is True, and C is False. If C answers 'da', then A is Random, B is False, and C is True. However, if C fails to answer, then we only know that neither A nor C is Random. A third question is needed in this case.

(Q3*) Put to A: Would you answer 'ja' to the question whether you are not Random and you are False?

Since we know that B is Random, A answers 'ja' if A is True and C is False, and A answers 'da' if A is False and C is True.

Even though our solution requires three questions, how do we know that there isn't a clever two-question solution waiting in the wings? To resolve this worry, we appeal to a lemma from Information Theory:

(QL) If a question has N possible answers, these N answers cannot distinguish $M > N$ different possibilities.

Initially, any question put to any god has three possible answers, 'ja', 'da', and no response, and there are six different possibilities to distinguish, (P1) through (P6). So, by (QL), it follows trivially that the puzzle cannot be solved in one question. To have a shot at solving the puzzle in two questions, it is necessary for there to be a first question that reduces the number of possibilities to no more than three. Our first question fails this condition precisely when there is no response, since this leaves four possibilities to consider. To see that this limitation is a feature of any solution and not merely our own, observe that whatever first question is posed to whichever god, say to A, we cannot exclude the possibility that A is Random, since Random answers 'ja', 'da', or remains silent with equal probability. But, if A were Random, there would be no information in his response that could be used to distinguish between states (P5) and (P6). So, for each of the three possible answers to the first question, states (P5) and (P6) will not be eliminated. So, the best we can do is to separate the six initial possibilities by the three possible answers, yielding a single case corresponding to 'ja' and a single case corresponding to 'da', which leaves two cases for no answer to distinguish. Thus, there is always a scenario in which at least four possibilities remain after the first question. Because any question has three distinct answers, 'ja', 'da', and no answer, the scenario in which there are four possibilities cannot be distinguished by a single question. For this case a third question is necessary.[1]

Our solution to the hardest logic puzzle ever trades on identifying the pair of non-Random gods by relying on their reliable behavior in answering and remaining silent. Alas, this suggests an even harder variation on the puzzle if we replace the god False with another named Devious:

> Three gods, A, B, and C are called in some order, True, Random, and Devious. True always speaks truly, and whether Random speaks truly or falsely or whether Random speaks at all is a completely random manner. Devious always speaks falsely, if he is certain he can; but if he is unable to lie with certainty, he responds like Random. Your task is to determine the identities of A, B, and C by asking the minimum number of yes-no questions; each question must be put to exactly one god. The gods understand English, but will answer 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 here are our guidelines:

1. It could be that some god gets asked more than one question.
2. What the second question is, and to which god it is put, may depend on the answer to the first question.
3. 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'.
4. When Devious is able to lie he does so; but if Devious cannot be sure of telling a lie, then rather than remain silent, he responds randomly like Random, i.e., there is a fair three-sided dice in Devious' brain that is tossed when he is not

---

[1] More formally, (QL) says that an answer can decrease the initial entropy of an agent's knowledge about the identity of the gods by the maximum quantity of information an answer may convey. When there are four equally probable states, the entropy of the agent is $H = -4(1/4) \times \log2(1/4) = 2$ bits. The maximum information contained in a 3-valued answer is $I = -3(1/3) \times \log2(1/3) = 1.585$ bits, so the entropy can only decrease to 2-1.585 = 0.415 bits. Thus, our agent is not guaranteed certain knowledge after two questions, because that result requires probability 1 or 0 to every state *j*, such that *j*'s contribution to the whole entropy is $H_j = p_j \times \log2(p_j)$ is 0.

certain to lie. If the dice comes down 1, he doesn't answer at all; if the dice comes down 2, he answers 'da'; if 3, 'ja'.

Although we do not have a solution to this puzzle, we are certain that it cannot be solved in less than three questions.[2]

*CENTRIA - Centre for Artificial Intelligence*
*The New University of Lisbon, FCT*
*2829-516 Caparica, Portugal*
*{grw,pb}@fct.unl.pt*

*Draft of July 23, 2010*

*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.

Uzquiano, G. 2010. How to solve the hardest logic puzzle ever in two questions. *Analysis* 70: 39—44.



[2] Our thanks to …