American Philosophical Quarterly
Volume 42, Number 1, January 2005

THREE DOORS, TWO PLAYERS, AND SINGLE-CASE PROBABILITIES

Peter Baumann

Suppose Jack is the only player in the TV show "Let's Make a Deal!" Jack knows there is a prize he wants behind one of three doors and nothing behind the other two doors. He doesn't know behind which door the prize is. Jack knows he can pick one door and keep what is behind it. Jack also knows this: After he has picked a door, the host, Monty Hall, will open a door with nothing behind it (and not the one he chose); he will then have the choice between either sticking with his original pick, get what is behind the door and receive additional $100, or switch to the other remaining door and not get the $100 (but anything that is behind the door). All this is common knowledge. What should he do?[1]

Most people's untutored intuitions say one should stay with the original choice. But there is a decisive argument against it: If one plays a long enough series of such games one will win in one third of the cases if one follows the sticking strategy whereas one will win in two thirds of the cases—the remaining cases—if one follows the switch strategy. Hence, one's chances of winning the prize are 1/3 if one sticks with one's original choice, and 2/3 if one switches (cf., among others, Gillman 1992; Morgan et al. 1991; Bradley and Fittelson 2003). Application of Bayes's principle shows that the conditional probability of winning by switching, given the evidence (that Monty Hall has opened a particular door, etc.), is 2/3, whereas the conditional probability of winning by sticking to the originally chosen door, given the evidence, is 1/3.[2]

There should be no doubt that with respect to a long enough series of games the switching strategy is the better one. However, there is serious doubt that this implies anything about what is rational to do in an individual and isolated case. So, this adds to well known issues concerning the very idea of single-case probabilities. Moser and Mulder 1994 have argued for this point (see also Horgan 1995) and this view will be supported here with a different (and new) argument concerning two, not just one, players. It will be argued here that applying arguments like the above one to individual cases violates a very plausible non-arbitrariness-condition and leads to Moore-paradoxical incoherence.

I. Two Players

For the sake of the argument, one can modify the original assumptions. There are now 2 players. Both know this but when they make their initial choice of a door they don't know what the other player is choosing. If they choose the same door, everything is like in the original version of the game: Monty

Hall opens one of the other doors (and only a door with nothing behind it). If they choose different doors, then Monty Hall will open the unchosen door, even if it has a prize behind it. If the players see that there is a prize behind the opened door, they go for it, no matter what their strategy. If both players pick the door with the prize behind it, then they both get the full prize. Everybody involved is rational. All this is common knowledge. The strategies here are modified ones which are still quite similar to the strategies in the original 1-player scenario.

It is easy to see that the modified switching strategy as applied to a long enough series of games will lead to wins in 2/3 of the cases whereas the modified sticking strategy is (slightly) worse and will only lead to a win in 5/9 of the games. To see that, suppose door 1 wins (since the other two possibilities are analogous, one can disregard them here). Then there are 9 different possibilities of combined choice (equally probable, mutually exclusive and jointly exhaustive). For each combination of the choices of the two players there are 1 or two possible door openings by Monty Hall. The table below gives the outcomes for the two strategies for A (B's case is analogous).

A chooses	B chooses	M.H. opens door	switch	stick
1	1	2 or 3	L	W
1	2	3	L	W
1	3	2	L	W
2	1	3	W	L
2	2	3	W	L
2	3	1	W	W
3	1	2	W	L
3	2	1	W	W
3	3	2	W	L
			6/9 wins	5/9 wins

Now, interestingly the situation is the same for A and B. Should one then say that for both players the chances of winning by switching are 6/9 *in an individual game* (not just in a long enough series of games)?

This leads to serious problems. Consider the probabilities after Monty Hall has opened an empty door (the above nine cases thus reduce to seven). The chances of winning by switching are now 4/7 whereas the chances of winning by sticking are now 3/7. If one makes the further assumption that A and B have chosen different doors (but they don't know that they have, according to the rules of the game), then the objective or non-epistemic probabilities of winning would be 1/2 for both switching and sticking. However, since A and B don't know they have chosen different doors, their epistemic probabilities will remain 4/7 for winning if they switch and 3/7 for winning if they stick. Given that A has initially picked door 1 and B door 2, the following holds (again, given ignorance about the other contestant's initial pick):

A's epistemic chance of winning by finally choosing door 2 = 4/7

A's epistemic chance of winning by finally choosing door 1 = 3/7

B's epistemic chance of winning by finally choosing door 1 = 4/7

B's epistemic chance of winning by finally choosing door 2 = 3/7

In other words, just before the final choice of a door the following holds:

the chance of door 1 being the winning door
= 3/7 if one considers A but
= 4/7 if one considers B;

the chance of door 2 being the winning door
= 4/7 if one considers A but
= 3/7 if one considers B

This would be clearly absurd if one were talking about non-epistemic probabilities here: How could the probability that door 1 wins be both 3/7 and 4/7? If the probability that A wins if he chooses door 2 is 4/7 and the probability that B wins if he chooses door

1 is 4/7, too, then the probabilities would not add up to 1.

However, all this is about epistemic probabilities. Shouldn't one in this case relativize to players: From A's perspective choosing door 2 has a chance of winning of 4/7 (in this individual case) whereas from B's perspective choosing door 2 has a chance of winning of 3/7 (in this individual case)? But what could "perspective" mean and why would it make such a difference? What it could mean is this: A and B do not share the same knowledge or information; hence, their subjective or epistemic probabilities can differ.[3] One knows what the other doesn't know: A knows that he has picked door 1 but doesn't know that B has picked door 2 whereas B knows the latter but not the former. Sure, this epistemic difference between A and B is hard to deny.

But why should it make a difference for the probabilities (in an individual case)? Why should it constitute a *relevant* difference of informational states? It is hard to see why—as will be shown in a moment. If this is true, then to allow for a difference between A's and B's subjective probabilities concerning doors and prizes would violate a very plausible principle of non-arbitrariness:

(N) If

a.) two rational persons' (A, B) assignments of a particular probability that p (that a particular door has a prize behind it) are not arbitrary but backed by evidence, and if

b.) they both share the same relevant information, then their subjective probabilities that p are the same.

This needs to be explained. In a way it is unproblematic to arbitrarily "pick" one's (prior) probabilities. But this is only unproblematic because and insofar as in the long run differences between prior probabilities will even out under the impact of conditionalization. If one updates one's probabilities on future evidence, then they will converge (at least according to classical Bayesianism which will not be discussed here). However, it is essential to the case considered here that there will be no "long run" but just a single case. Apart from that, in the modified Monty Hall-scenario persented above both players are assumed to have good reasons for their probability assignments. Hence, condition (N-a) is met.

What about condition (N-b)? What counts as irrelevant information? If Ernie has information that he is in Cairo but has no clue where Bert is, and if Bert has information that he is in Budapest but has no clue where Ernie is, then they do not have the same information. But they might still have the same information relevant to the question whether Britain will give up monarchy within the next five years. In the modified Monty Hall case the two players have different information but—like in the Ernie-Bert case—it is hard to see how that could be relevant with respect to the relevant question: which door—door 1 or door 2—will have what probability of being the winning door in this particular game. Hence, (N-b) is met, too.

There is a very tempting objection to this last step which in the end turns out to be subtly misleading. Suppose A has originally picked door 1 and he knows this of course. This information plus the information that Monty has opened door 3 plus some probabilistic reasoning about switching to the "other" door seems to suggest that he should make door 2 his final choice. He couldn't arrive at this conclusion without the information that he has originally picked door 1. Hence that information seems clearly relevant. Furthermore, it is information B (who has originally picked door 2) does not have. It seems to follow that (N-b) is not met.

However, as will be argued in the next section, an expression like "the probability that the other door will win" and an expression like "the probability that door 2 will win" do have different meanings, and this difference matters here. All the relevant information the chooser has can be expressed in terms of "the original door" and "the other door." There is no relevant information having to do with the fact that one door is, say, door 1. Hence, A and B share the same relevant information.

One might object that the information B has and A lacks—that B has initially picked door 2—is very relevant for A: Were A to receive this information he could easily figure out that he and B have picked different doors. Obviously, A's subjective probabilities of winning by switching and winning by sticking would both change to 1/2 (given the assumption that the players are rational). So, how could information about the initial pick not be relevant? This objection is misled. The information that B has initially picked door 2 is relevant for A only given the further knowledge that he himself (A) has initially picked door 1. In other words, the relevant information would be expressed by "B has initially picked another door than I have"; nothing hinges on what particular door B has picked (see the next section for the difference between "the probability that the other door will win" and "the probability that door 2 will win").

One can also make this point in the following way. Suppose A and B both forget what their initial pick was. However, both know that if they answer "Switch!" (or "Stick!" for that matter) when asked for their final answer the host will follow instructions and, knowing the contestants initial picks, give them what is behind the respective (the other or the same) door. It seems obvious that this variation does not introduce any relevant differences into the game. However, the difference in A's and B's information concerning their initial picks drops out of the picture. Hence, the information about initial picks cannot be relevant. In other words, it does not seem to make any difference for one's final choice whether one remembers which particular door one had picked initially. Hence, that information is not relevant.

One final objection against the irrelevance thesis goes like this. A's and B's final choices differ if they're both switchers (or if they are both rational). Since their final choices are based on good reasons, and since the only information they do not share is the one concerning their initial picks, the latter information must be relevant. However, this objection seems to clearly beg the question: It assumes what is being argued against here, namely that information about an individual case matters here. To put it more positively: A's and B's final choices do not differ insofar as they both switch. Their respective reasons for switching do not differ. So, if there are no different choices, then there is no point in identifying presumably different relevant information which would explain a difference in the final choice. Hence, it all boils down to this question: Do A's and B's final choices differ or not? There cannot be disagreement that they differ in some respects. For instance, A's choice is his and not B's choice (and vc. vs.). The real question is whether there is a relevant difference here. Suppose both A and B have voted for the same Party (P). In a sense they have made the same choice but in another sense they haven't. They might both vote for P thinking "P is better for me!" but "me" then refers to different persons; on the other hand they have both voted for P. Whether one says that A and B have made the same choice or not depends on what one thinks is relevant for qualifying choices as the same or different. The objection above assumes that information about the particular case at hand matters. This begs the question because the arguments presented here go against that thesis.

Hence, what has been said here suggests that the information that one has originally picked door 1 (or 2) is not relevant for probabilistic considerations about switching or sticking.[4] If that is true, the above objection collapses and it is quite correct to say that (N-b), too, is met by A and B.

If one accepts (N), then one has to conclude that the subjective probabilities of the two players are the same if they are both rational. Hence, as soon as one is dealing with an individual case and not with probabilities in the long run, the argument for switching collapses. This does not mean that the argument for sticking wins. As has been pointed out at the beginning, switching is a better strategy than sticking when it comes to a sample of many games. How then, could the probabilities differ for an individual game? The problem is more fundamental: One cannot assign non-arbitrary single-case probabilities of winning to different final choices of doors. This throws some sceptical light on the idea of single-case probabilities in general. It offers another argument against the idea that there could be single-case probabilities.

II. Rigidity

There is a semantical side to all this. It remains true to say that switching to *the other door* is the better strategy. However, there is no good argument that would recommend to A that he switches to *door 2* (in the case discussed above). This is no contradiction because different things are mentioned here. But isn't A's "other" door = door 2? How could he then have a reason to switch to *the other door* but no reason to switch to *door 2*?

This has to do with the fact that different kinds of singular terms are being used here. Both A and B (or an observer) are able to refer to the doors in two different ways: by using (rigidly referring) proper names like "door 1" and "door 2" or by using (non-rigidly referring) definite descriptions like "the door I originally picked" or "the other door." When one considers samples of many different cases, one can attribute probabilities to options only if they are being described in terms of definite descriptions like "sticking to the originally chosen door." One cannot—as long as one considers samples of many individual cases—assign probabilities if the options are being referred to in terms of proper names like "door 1" or "door 2": It is, of course, very misleading even to raise the question whether, in the long run, one should rather choose door 1 than door 2.

The situation is quite different when it comes to single-case probabilities. In a single case like the one above both players know which door they have originally picked. They can refer to the different doors both by definite descriptions and by proper names. The same is true for an observer. One can, again, consider cases in which Monty Hall has opened an empty door (and A has initially picked door 1 whereas B has initially picked door 2). Now, one might be tempted to argue that if switching to "the other door" gives A a 4/7 chance of winning and if the other door = door 2, then his final choice of door 2 gives him a 4/7 chance of winning. Isn't it hard to see how the latter chance could differ from the former, given the identity of the "other" door and door 2? And how then could the probabilities differ between A and B?

However all this is only true for objective probabilities, not for subjective or epistemic probabilities. Under different possible circumstances, the person does not have the information that the other door = door 2. Probability assignments are intensional (cf. Horgan 2000: 589–590) when it comes to subjective or epistemic probability.[5] So A's probability that the other door will win need not be identical with A's probability that door 2 will win.

But this is not yet the whole picture; there is something right about the argument just mentioned. To be sure, the prior probabilities of "the other door wins" and "door 2 wins" are different (2/3 vs. 1/3). But then—when Monty Hall opens the empty door—player A learns that the other door = door 2. As a rational person (and we are certainly talking about rational expectations here) A will update his probabilities, given the new information. The probability that the other door wins remains 4/7 but since he now knows that the other door = door 2, he has to change the probability that door 2 will win to 4/7. How could he consistently believe that the other door = door 2, that the probability that the other door will win is 4/7 while the probability that door 2 will win is 3/7 or, at least, not 4/7?

However, if all this were true, then one would run into problems if one were to assign probabilities in single cases. Since A knows that "his" other door = door 2, he would, as a rational person, have to assign the same probability (4/7) to "My other door will win" and to "Door 2 will win." For analogous reasons, B would have to assign a probability of 3/7 both to "My original door will win" and to "Door 2 will win." Again: How could probabilities differ here for A and B? Only if they do not share all the relevant information. Which information could that be? The only candidate is this one: A (B) but not B (A) knows that he has initially picked door 1 (2). However, whether on has initially picked this or that door does not matter at all when it comes to one's probabilistic reasoning about switching and sticking. Hence, both A and B share the same relevant information. And according to the principle of non-arbitrariness (N) their probabilities should match. Hence, one ought to give up probabilistic arguments for or against switching or sticking in isolated individual cases. They only make sense in the long run.

All this explains why there is no inconsistency when one says that the chances of winning are 4/7 if one switches but that there is no reason to switch to door 2 (given that door 2 = the other door). In one case one is using a (non-rigidly referring) description, "the other door," whereas in the other case one is using a (rigidly referring) proper name, "door 2." Using the proper name, one is talking about an individual case whereas one can talk about samples of many cases by using a description. With respect to samples of cases there is an argument for switching but with respect to an isolated individual case there is no probabilistic argument at all, neither for switching or for sticking.

III. One Player and his Counterfactual Counterpart

There is another way to make the argument. In the original Monty Hall scenario there is only one player, A. For the sake of the reductio, it is okay to assume here that one can apply probabilistic considerations to single cases. If A has originally picked door 1 and Monty Hall has just opened door 3, then the probability that the originally unchosen door ("the other door") will win is, of course, 2/3. Since A knows that the other door = door 2, he will update the probability that door 2 will win to 2/3; accordingly, the probability that door 1 will win is 1/3 now. A also knows the following thing: Had everything been the same except that he had picked door 2 originally, then the probability that door 1 will win would be 2/3 and the probability that door 2 will win would be 1/3. In other words, he knows that the probability that door 2 (or door 1 for that matter) will win is either 2/3 or 1/3; what it is depends only on his initial pick. If door 2 has a 2/3 chance of winning, that is because A had (given the circumstances of this case) initially picked door 1; had he initially picked door 2, then the probabilities would be the reverse ones. But all that just seems absurd. It violates (N) and its relatives:[6] Probabilities do not vary with irrelevant factors. Now, whether one has originally picked one or the other door, cannot make a difference as to one's (epistemic) probabilities. There is no relevant difference between the two doors (see above). Hence, there is something wrong with the conclusion that if A has originally picked door 1, then door 2 will have a 2/3 winning chance whereas if A has originally picked door 2, door 1 will have a 2/3 winning chance. The actual case in which A originally picks door 1 doesn't differ from the counterfactual case in which A originally picks door 2 in such a way that the distribution of probabilities in both cases should differ.

Suppose the player is thinking about the game before it starts. He doesn't know yet what his initial choice will be but he knows he is a switcher. He thus already knows that whichever door he will choose initially he will switch to another door. Suppose he decides to make a random initial choice. He then knows that his final choice will be determined by his initial random choice. This makes his final choice—whatever it turns out to be—appear rather arbitrary.

To be sure, as long as othe relevant probability statements exclusively use definite descriptions (like "the other door") everything is fine. One can just say that switching to the other door has a winning chance of 2/3 and there is no problem with that because one can always (and does always) refer to a plurality of cases. However, as soon as one talks about individual cases, the person can identify the other door as one particular door and thus make probability statements using (rigid) proper names. And then one runs into the above absurdity that the probability that door 2 wins is 2/3 because one has originally picked door 1 whereas it would have been 1/3 had one originally picked door 2.

IV. Moorean Insecurities

How would the player see her own probability assignments? It seems that A, given the circumstances just mentioned, would have to see them as arbitrary. A would have to say something Moore-paradoxical:

> Switching to door 2 gives me a 2/3 chance of winning but this belief is completely arbitrary.

This only adds to the serious problems one runs into if one makes probability statements about isolated individual Monty Hall situations.

One can make the same general point with respect to two player-situations. Again, the assumption is that Monty Hall has just opened an empty door; A and B initial picks are door 1 and 2 respectively. Suppose A thinks that switching to door 2 will give him a 4/7-chance of winning. He doesn't know what B has originally picked but he knows that if B's original choice was door 2, then B's subjective probability of winning by switching to door 1 will be 4/7. A doesn't know whether he is in this situation but he knows that if he is then B's reasoning in favor of door 1 would have to be accepted as much as his own reasoning in favor of door 2. But if one can imagine and cannot exclude circumstances in which one's own actual reasoning is as convincing as an alternative and incompatible way of reasoning, then one's own actual reasoning loses its convincing force. One ends up in a Moore-paradoxical situation.

V. Conclusion

The general morals of all this is the following one. In the modified version of Monty Hall discussed here, it doesn't make sense to apply statistical probabilities to a singular case. If one does, one faces arbitrariness and Moore-paradoxality. This is the main result of the *reductio* proposed here. Since the original version of the Monty Hall-scenario is relevantly similar, it also doesn't make sense to attribute probabilities to single cases in the original version. If the best argument so far for switching in an isolated individual case (not in a series of cases) fails, then one might wonder whether probabilistic arguments say anything at all about isolated individual cases.[7]

Aberdeen University

NOTES

1. For other variations of the same problem and for related problems, cf. Gardner 1959: 180–182, 188; Copi 1968: 433; Lindley 1971: 43–44. Cf. also Falk 1993; Ellin 1993–1995; Bradley and Fittelson 2003. The Monty Hall problem was initially brought to the attention of a broader public by Marilyn vos Savant: cf. vos Savant 1992: 199–209.

2. One can also give a more frequentist exposition of this argument. Suppose the door with the prize behind is called "door 1" ("A" standing for the player, "W" for winning and "L" for losing). Then there will be three cases that are of equal probability, mutually exclusive and jointly exhaustive:

A chooses	M.H. opens door no.	switch	stick
1	2 or 3	L	W
2	3	W	L
3	2	W	L
		2/3 W	1/3 W

3. On the question what is and what isn't irrelevant when it comes to informational states, cf. Freund 1965, Shafer 1983, and Sobel 1992. For similar questions regarding a different puzzle, cf. Elga 2000 and Lewis 2001.

4. What about other information the players acquire during the game: that Monty opens door x, and that nothing/the prize is behind x? Furthermore, the players can infer from this information either what the other player has chosen (if Monty opens a full door) or which door the other player has not chosen (if Monty opens an empty door). All this is uncontroversial. However, these kinds of information are as "irrelevant" here as the information discussed in the text above (about which door oneself has initially picked). The reasons why these kinds of information are irrelevant here are similar or even analogous to the ones mentioned in the case above. There is no need to go into the details here.

5. The probabilities here are *de dicto* probabilities: From the perspective of A, one is talking about the probability that door so-and-so will win. One is not talking about *de re* probabilities ("There is a door x such that so-and-so is true of it and the probability that x will win = p"). See also Horgan 2000: 585–586.

6. (N) unproblematically applies here if one regards A and his counterfactual counterpart as two persons.

7. The author's thanks go to Joseph Ellin, Arthur Falk, Terence Horgan, Paul Moser, and Jordan Howard Sobel.

REFERENCES

Bradley, Darren, and Branden Fittelson. 2003. "Monty Hall, Doomsday and Confirmation." *Analysis*, vol. 63, pp. 23–31.

Copi, Irving M. 1968. *Introduction to Logic*, 3rd ed. (New York: Macmillan).

Elga, Adam. 2000. "Self-Locating Belief and the Sleeping Beauty Problem." *Analysis*, vol. 60, pp. 143–147.

Ellin, Joseph. 1993–95. "Monty Hall No Newcomb Problem." *Proceedings of the Heraclitean Society*, vol. 17, pp. 149–152.

Falk, Arthur. 1993. "Summer 1991: The Monty Hall Problem; Fall 1993: The Two Envelopes Puzzle; and Now: Doomsday." *Proceedings of the Heraclitean Society*, vol. 17, pp. 64–78.

Freund, John E. 1965. "Puzzle or Paradox." *The American Statistician*, vol. 19, no. 4, pp. 29–44.

Gardner, Martin. 1959a. "Mathematical Games." *Scientific American*, vol. 201, no. 4 (October), pp. 180–182.

_______ . 1959b. "Mathematical Games." *Scientific American*, vol. 201, no. 5 (November), p. 188.

Gillman, Leonard. 1992. "The Car and the Goats." *The American Mathematical Monthly*, vol. 99, no. 1, pp. 3–7.

Horgan, Terence. 1995. "Let's Make a Deal." *Philosophical Papers*, vol. 24, pp. 209–222.

_______ . 2000. "The Two-Envelope Paradox, Nonstandard Expected Utility, and the Intensionality of Probability." *Noûs*, vol. 34, pp. 578–603.

Lewis, David. 2001. "Sleeping Beauty: Reply to Elga." *Analysis*, vol. 61, pp. 171–176.

Lindley, Dennis V. 1971. *Making Decisions* (London: John Wiley).

Morgan, John P., N. Rao Chaganty, Ram C. Dahiya, and Michael J. Doviak. 1991. "Let's Make a Deal: The Player's Dilemma." *The American Statistician*, vol. 45, no. 4, pp. 284–287.

Moser, Paul K., and D. Hudson Mulder. 1994. "Probability in Rational Decision-Making." *Philosophical Papers*, vol. 23, pp. 109–128.

Shafer, Glenn. 1985. "Conditional Probability." *International Statistical Review*, vol. 53, pp. 261–277.

Sobel, Jordan Howard. 1992. "Kings and Prisoners (and Aces)." *Proceedings of the Biennial Meetings of the Philosophy of Science Association*, vol. 1, pp. 203–216.

vos Savant, Marilyn. 1992. *Ask Marylin* (New York: St. Martin's).