The Monty Hall Problem

Richard D. Gill

Mathematical Institute, University of Leiden, Netherlands

http://www.math.leidenuniv.nl/~gill

17 March, 2011

Abstract

A range of solutions to the Monty Hall problem is developed, with the aim of connecting popular (informal) solutions and the formal mathematical solutions of introductory text-books. Under Riemann's slogan of replacing calculations with ideas, we discover bridges between popular solutions and rigorous (formal) mathematical solutions by using various combinations of symmetry, symmetrization, independence, and Bayes' rule, as well as insights from game theory. The result is a collection of intuitive and informal logical arguments which can be converted step by step into formal mathematics.

The Monty Hall problem can be used simultaneously to develop probabilistic intuition and to give a deeper understanding of the paradox, not just to provide a routine exercise in computation of conditional probabilities. Simple insights from game theory can be used to prove probability inequalities.

Acknowledgement: this text is the result of an evolution through texts written for three internet encyclopedias (wikipedia.org, citizendium.org, statprob.com) and each time has benefitted from the contributions of many other editors. I thank them all, for all the original ideas in this paper. That leaves only the mistakes, for which I bear responsibility.

1 Introduction

Imagine you are a guest in a TV game show. The host, a certain Mr. Monty Hall, shows you three large doors and tells you that behind one of the doors there is a car while behind the other two there are goats. You want to win the car. He asks you to choose a door. After you have made your choice, he opens another door, revealing a goat. He then asks you whether you want to stay with your initial choice, or switch to the remaining closed door. Would you switch or stay?

The host, naturally, knows in advance which of the three doors hides the car. This means that whatever door you initially choose, he can indeed open a different door and reveal a goat. Stronger still: not only *can* he do this; you also know he certainly *will* do this.

The instinctive, but incorrect, answer of almost all newcomers to the problem is: "*Stay*, since it is equally likely that the car is behind either of the two closed doors".

However, under very natural assumptions, the good answer is "*Switch*, since this doubles my chance of winning the car: it goes from one third to two thirds".

Because of this conflict the *Monty Hall problem* is often called the *Monty Hall paradox*.

The key to accepting and understanding the paradox is to realize that the (subjective) probabilities relevant for the decision are not determined by the situation (two doors closed) alone, but also by what is known about the development that led to this situation. In statistical terminology, the data is not an unordered set of two closed doors, but an ordered set, where the ordering corresponds to the roles of the closed doors: the one chosen by you, and the one left un-chosen by the host. We have to model the data-generating mechanism as well as the data.

There are several intuitive arguments why switching is a good stategy. One is the following. The chances are 1 in 3 that the door initially chosen hides the car. When that happens *staying* is good, it gives the car. Both of the other two doors hide goats; one is revealed by the host, but *switching* to the other door just gives the other goat. Complementarily to this, the chances are 2 in 3 that the door initially chosen hides a goat. When that happens, *staying* is not good: it gives a goat. On the other hand, switching certainly does give the car: the host is forced to open the other door hiding a goat, and the remaining closed door is the door hiding the car.

In many repetitions, one third of the times the stayer will win and the switcher will lose; two thirds of the time the stayer will lose and the switcher will win.

The (wrong) intuitive answer "50–50" is often supported by saying that the host has not provided any new information by opening a door and revealing a goat since the contestant knows in advance that at least one of the other two doors hides a goat, and that the host will open that this door or one of those doors as the case may be. The contestant merely gets to know the identity of one of those two. How can this "non-information" change the fact the remaining doors are equally likely to hide the car?

However, precisely the same reasoning can be used *against* this answer: if indeed the host's action does not give away information about what is behind the closed doors, how can his action increase the winning chances for the door first chosen from 1 in 3 to 1 in 2? The paradox is that while initially doors 1 and 2 were equally likely to hide the car, after the player has chosen door 1 and the host has opened door 3, door 2 is twice as likely as door 1 to hide the car. The paradox (apparent, but not actual, contradiction) holds because it is equally true that initially door 1 had chance 1/3 to hide the car, while after the player has chosen door 1 and the host has opened door 3, door 1 *still* has chance 1/3 to hide the car.

The Monty Hall problem (MHP) became internationally famous after its publication vos Savant (1990) in a popular weekly magazine led to a huge controversy in the media. It has been causing endless disputes and arguments since then.

2 The origins of MHP

Also known as the as the *Monty Hall paradox*, the *three doors problem*, the *quizmaster problem*, and the problem of the *car and the goats*, the problem was introduced by biostatistician Steve Selvin (1975a) in a letter to the journal *The American Statistician*. Depending on what assumptions are made, it can be seen as mathematically identical to the Three Prisoners Problem of Martin Gardner (1959a,b). It was named by Selvin after the stage-name of the actual quizmaster, Monty Halperin (or Halparin) of the long-running 1960's TV show "Let's make a Deal". Selvin's letter provoked a number of people to write to the author, and he published a second letter in response, Selvin (1975b). One of his correspondents was Monty Hall himself, who pointed out that the formulation of the Monty Hall problem did not correspond with reality: in actual fact, Monty only occasionally offered a player the option to switch to another door, and he did this depending on whether or not the player had made a good or bad initial choice.

The problem, true to reality or not, became world famous in 1990 with its presentation in the popular weekly column "Ask Marilyn" in *Parade* magazine. The author Marilyn Vos Savant, was, according to the *Guiness Book of Records* at the time, the person with the highest IQ in the world. Rewriting in her own words a problem posed to her by a correspondent, Craig Whitaker, vos Savant asked the following:

> "*Suppose you're on a game show, and you're given the choice of three doors: behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice?*".

Vos Savant proceeded to give a number of simple arguments for the good answer: switch, it doubles your chance of winning the car. One of them was the previously mentioned argument that a stayer wins if and only if a switcher loses. A stayer only wins one third of the time. Hence a switcher only loses one third of the time, and wins two thirds of the time.

Another intuitive reasoning is the following: one could say that when the contestant initially chooses door 1, the host is offering the contestant a choice between his initial choice door 1, or doors 2 and 3 together, and kindly opens one of doors 2 and 3 in advance.

By changing one aspect of the problem, this way of understanding why the contestant indeed should switch may become even more compelling to the reader. Consider the 100-door problem: 99 goats and one car. The player chooses one of the 100 doors. Let's say that he chooses door number 1. The host, who knows the location of the car, one by one opens all 99 of the *other* doors but one – let's say that he skips door number 38. Would you switch?

The simple solutions often implicitly used a frequentist picture of probability: probability refers to relative frequency in many repetitions. They also do not address the issue of whether the specific door opened by the host is relevant: if the player has initially chosen door 1, could it be that the decision to switch should depend on whether the host opens door 2 or door 3? Intuition says no, but we already saw that naive intuition can be misleading.

These topics are taken up in the next section.

3 A more refined analysis

Marilyn vos Savant was taken to task by Morgan, Chaganty, Dahiya and Doviak (1991a), in another paper in *The American Statistician*, for not computing the conditional probability that switching will give the car, given the choice of the player and which door was opened by the host. Possibly with a frequentist view of probability in mind, Morgan et al. took it for granted that the car is hidden uniformly at random behind one of the three doors, but did not assume that the probability that the host would open door 3, rather than door 2, given that player has chosen door 1 and the car is behind door 1, is one half. Vos Savant (1991) responded angrily in a letter to the editor. Further comments were given by Seymann (1991), Bell (1992) and Bhaskara Rao (1992), then (after a long pause) Hogbin and Nijdam (2010), followed by a final response Morgan et al. (2010). This last note incidentally finally revealed Craig Whitaker's original wording (spelling and grammatical errors corrected by RDG) of the problem in his letter to vos Savant:

> "*I've worked out two different situations based on whether or not Monty knows what's behind the doors. In one situation it is to your advantage to switch, in the other there is no advantage to switch. What do you think?*"

The solution to be given in this section takes explicit account of which door was opened by the host: door 2 or door 3 – and does so in order to argue that this does not change the answer. This is probably the reason why lay persons on hearing the simple solutions of the previous section, do not see any need whatsoever for further analysis. Having seen that the strategy of "always switching" gives a success rate of 2/3, while "always staying" gives a success rate of 1/3, there seems little point in pondering whether or not the success rate of 2/3 could be improved.

This is mathematically a true fact, given the only probabilistic assumption which we have made (and used) so far: initially all doors are equally likely to hide the car. However, a really short rigorous *and* intuitive proof of this fact does not seem to exist.

Let's make a more careful analysis, in which a further (natural) assumption will indeed be used. In this section, *probability* is used its daily-life Bayesian or subjective sense: that is to say, probability statements are supposed to reflect the state of knowledge of one person. That person will be a contestant on the show who initially knows no more than the following: he'll choose a door; the quizmaster (who knows where the car is hidden) will thereupon open a different door revealing a goat and make the offer that the contestant switches to the remaining closed door. The argument will be kept intuitive or informal, however the student of probability theory will be able to convert every step into formal mathematics, if the need is felt to do so.

For our *tabula rasa* contestant, initially all doors are equally likely to hide the car. Moreover, if he chooses any particular door, and if the car happens to be behind that particular door, then as far as this contestant is concerned, the host is equally likely to open either of the other two doors.

The contestant initially chooses door number 1. Initially, his odds that the car is behind this door are 2 to 1 against: it is two times as likely for him that his choice is wrong as that it is right.

The host opens one of the other two doors, revealing a goat. Let's suppose that for the moment, the contestant doesn't take any notice of *which* door was opened. Since the host is certain to open a door revealing a goat whether or not the car is behind door 1, the information that an unspecified door is opened revealing a goat cannot change the contestant's odds that the car is indeed behind door 1; they are still 2 to 1 against.

Now here comes the further detail which we will take account of in this solution: the contestant also gets informed which specific door was opened by the host – let's say it was door 3. Does this piece of information influence his odds that the car is behind door 1? No: from the contestant's point of view, the chance that the car is behind door 1 obviously can't depend on whether the host opens door 2 or door 3 – the door numbers are arbitrary, exchangeable.

Therefore, also knowing that the host opened specifically door 3 to reveal a goat, the contestant's odds on the car being behind his initially chosen door 1 still remain 2 to 1 against. He had better switch to door 2.

4 Explicit computations

Students of probability theory might feel uneasy about the informality (the intuitive nature) of the last argument. Ordinary people's intuition about probability is well known to be often wrong – after all, it is ordinary intuition which makes most people believe there is no point in switching doors! To feel more secure, students of probability theory might consider the mathematical concept of symmetry and use the law of total probability to show how symmetry leads to statistical independence between the events "Car is behind door 1" and "Host opens door 3" when it is given that the contestant chose door 1. Alternatively, they might like to explicitly use Bayes' theorem, in the form known as Bayes' rule: posterior odds equals prior odds times likelihood ratio (*aka* Bayes factor). They just have to check that under the two competing hypotheses (whether or not the car is behind the door chosen by the contestant, door 1), the fact that it is door 3 (rather than door 2) which gets opened by the host has the same probability 1/2.

Either of these routes can be used to convert the last step of the argument in the previous section into a formal mathematical proof, see the Appendix to this paper.

An alternative approach is to use symmetry *in advance* to dispose of the door-numbers. Suppose without loss of generality (since later we will condition on its value anyway) that the contestant's initial choice of door number X is uniformly distributed over the three door numbers $\{1, 2, 3\}$. Independently of this, the car is hidden behind door C, also uniformly at random. Given X and C, the host opens door number H uniformly at random from the door numbers different from both X and C (in number, there are either one or two of them). Let Y be the remaining closed door, so (X, H, Y) is a random permutation of $(1, 2, 3)$. By symmetry it is uniformly distributed over the set of six permutations. We know that either $C = X$ or $C = Y$ with probabilities 1/3 and 2/3 respectively. By symmetry, the conditional probability that $C = X$ given the value of (X, H, Y) – one of the six permutations of $(1, 2, 3)$ – cannot depend on that value and hence the event $\{C = X\}$ is statistically independent of (X, H, Y).

The actual numbers of door chosen and door opened are irrelevant to deciding whether to switch or stay.

Note the use of the trick of symmetrization – randomization over the door initially chosen – in order to simplify the mathematical analysis.

Almost all introductory statistical texts solve the Monty Hall problem by computing the conditional probability that switching will give the car, from first principles. Arguments for the chosen assumptions, and for the chosen approach to solution, are usually lacking. Gill (2011) argues that Monty Hall can be seen as an exercise in the art of statistical model building, and actually allows many different solutions: as one makes more assumptions, the conclusions are stronger but the scope of application becomes smaller; moreover, the meaning and the meaningfulness of the assumptions and of the result are tied to the user's interpretation of probability. The task of the statistician is to present a menu of solutions; the user is the one who should choose according to his resources and wishes. Rosenthal (2005, 2008) is one of the few who at least uses Bayes' rule to make the solution more insightful.

5 Variations

Many, many variations of the Monty Hall problem have been studied in the enormous literature which has grown up about the problem. The book Rosenhouse (2010) is a good resource, as are also the wikipedia pages on the topic. We just consider two variations here.

5.1 The biased host

Morgan et al. (1981)'s main innovation was to allow the host to have a bias to one door or the other. Suppose that when he has a choice between doors 2 and 3, he opens door 3 with probability q. The Bayes factor for the hypotheses that the car is or is not behind door 1 therefore becomes $q : 1/2$. The prior odds were $1 : 2$ so the posterior odds become $q : 1$. This can be anything between $0 : 1$ and $1 : 1$, but whatever it is, it is not unfavourable to switching. A frequentist player who knows that the car has been hidden by a true uniform randomization, but does not know anything about the probabilistic nature of Monty's brain processes with regards to choosing a door to open, should switch anyway. He does not actually know the conditional probability that switching gives him the car, but he does know the unconditional probability is 2/3.

The appendix gives an alternative and elementary proof, without using Bayes, of the fact that all the conditional probabilities of winning by switching are at least 1/2, hence you might as well switch, and hence there is no strategy giving a better overall win chance than 2/3

5.2 Game theory

In the literature of game theory and mathematical economics, starting with Nalebuff (1987), the Monty Hall problem is treated as a finite two stage two person zero sum game. The car is hidden by the host (in advance), the contestant independently chooses a door. The host opens a door revealing a goat. The contestant is allowed to choose again. The contestant wants to win the car, the host wants to keep it. If we allow the two "game-players" (host, contestant) randomized strategies, then according to von Neumann's minimax theorem, they both have a minimax stategy, and the game has a value say p, such that if the contestant uses his minimax strategy, then whatever strategy is used by the host, the contestant will go home with the car with probability *at least* p; while on the other hand, if the host uses his minimax strategy, then whatever strategy is used by the contestant, the contestant will go home with the car with probability *at most* p.

It is not difficult to show, and symmetry is one way to establish this, that the minimax strategy of the host is: hide the car uniformly at random, and open either door with equal chance when there is a choice. The minimax strategy of the contestant is: choose a door uniformly at random and thereafter switch, regardless of the which door is opened by the host.

With his minimax strategy the contestant wins the car with probability 2/3 exactly, whatever strategy is used by the player. With the host's minimax strategy, the contestant can't do better than 2/3 (random initial choice and thereafter switch).

A wise player would be recommended to choose a door number in advance, at home, by a fair randomization, and later switch. He'll get the car with probability 2/3, he cannot do better, and his ego won't be damaged when his initial choice turned out to have been right and yet he switched and lost the car.

References

Bell, W. (1992), Comment on Morgan et al. (1991a), *Am. Statist.* **46** 241.

Bhaskara Rao, M. (1992), Comment on Morgan et al. (1991a), *Am. Statist.* **46** 241–242.

Gardner, M. (1959a), "Mathematical Games" column, *Scientific American*, October 1959, 180–182. Reprinted in *The Second Scientific American Book of Mathematical Puzzles and Diversions.*

Gardner, M. (1959b), "Mathematical Games" column, *Scientific American*, November 1959, p. 188.

Gill, R.D. (2011), The Monty Hall Problem is not a Probability Puzzle – it's a challenge in mathematical modelling, *Statistica Neerlandica* **65** 58–71.

Hogbin, M. and Nijdam, W. (2010), Letter to the editor, *Am. Statist.* **64** 193.

Morgan, J. P., Chaganty, N. R., Dahiya, R. C., and Doviak, M. J. (1991a), Let's make a deal: The player's dilemma, *Am. Statist.* **45** 284–287.

Morgan, J. P., Chaganty, N. R., Dahiya, R. C., and Doviak, M. J. (1991b), Rejoinder to Seymann's comment on "Let's make a deal: the player's dilemma", *Am. Statist.* **45** 289.

Morgan, J. P., Chaganty, N. R., Dahiya, R. C., and Doviak, M. J. (2010), Response to Hogbin and Nijdam's letter, *Am. Statist.* **64** 193–194.

Nalebuff, B. (1987), Puzzles: Choose a curtain, duel-ity, two point conversions, and more, *J. Econ. Perspectives* **1** (2) 157–163.

Rosenhouse, J. (2009), *The Monty Hall Problem*, Oxford University Press.

Rosenthal, J. S. (2005), *Struck by Lightning: The Curious World of Probabilities*, Harper Collins, Canada.

Rosenthal, J. S. (2008), Monty Hall, Monty Fall, Monty Crawl, *Math Horizons* **16** 5–7. `http://probability.ca/jeff/writing/montyfall.pdf`

Selvin, S. (1975a), A problem in probability (letter to the editor), *Am. Statist.* **29** 67.

Selvin, S. (1975b), On the Monty Hall problem (letter to the editor), *Am. Statist.* **29** 134.

Seyman, R. G. (1991), Comment on "Let's make a deal: the player's dilemma", *Am. Statist.* **64** 287–288.

vos Savant, Marilyn (1990), "Ask Marilyn" column, *Parade Magazine*, 9 September 1990, p.16.

vos Savant, Marilyn (1991), Letter to the editor "Marilyn vos Savant's reply", *Am.*

Statist. **45** 347. With response by Morgan, Chaganty, Dahiya and Doviak, pp. 347-348.

Appendix

A.1 Solution by arithmetic

For some non-mathematicians, numbers speak louder than words or mathematical formalism. Suppose the player chooses door 1. The probability that the car is behind door 1 is 1/3, and in that case, the host might open either door 2 or door 3, with equal chances, while if the car is behind door 2 or door 3 the host's choice is forced. Altogether there are just four possibilities:

1. The car is behind door 1 and the host opens door 2 with probability 1/6
2. The car is behind door 1 and the host opens door 3 with probability 1/6
3. The car is behind door 2 and the host opens door 3 with probability 1/3
4. The car is behind door 3 and the host opens door 2 with probability 1/3

We observe that the player who switches wins the car with probability 2/3 (cases 3 and 4). We also see that door 3 is opened by the host with probability $1/2 = 1/6+1/3$ (cases 2 and 3), as must also be the case by the symmetry of the problem with regard to the door numbers – either door 2 or door 3 must be opened, and the chance of each must be the same, on half, by symmetry.

Winning by switching in combination with door 3 being opened occurs with probability 1/3 (case 3). The *conditional probability* of winning by switching, given door 3 is opened, is therefore $(1/3)/(1/2) = 2/3$. Since this is the same as the overall chance 2/3 of winning by switching, we see that knowing the identity of the opened door doesn't change the chance of winning by switching. Not only does the switcher win with probability 2/3 of the time, he also wins with conditional probabillity 2/3 given that door 3 is opened by the host, and with conditional probability 2/3 given that door 2 is opened by the host.

In other words, the combined chance of winning by switching *and* door 3 (rather than door 2) being opened, 1/3, equals the product of the separate chances of "the car being behind the other door", 2/3, and "host opens door 3", 1/2. Whether or not the car is behind the door not opened by the host is *statistically independent* of whether the host opens door 2 or door 3.

This last fact could have predicted in advance, by the symmetry of the probabilistic ingredients to the problem. The contestant may simply ignore the door numbers: they do not change his chances of winning by staying or by switching.

A.2 Solution by formal probability calculus

For the following analysis we will assume all natural (subjectivist) uniformity assumptions: the car is initally equally likely behind any of the three doors, and if the host has a choice of door to open, he is equally likely to open either.

Let C be the number of the door hiding the car, uniformly distributed on $\{1, 2, 3\}$. Let X be the number of the door initially chosen by the player. We assume that X and C are independent. Sometimes we will take X as identically equal to 1, sometimes as uniformly distributed on $\{1, 2, 3\}$. In the latter case, conditioning on $X = 1$ brings us back to the former case. Conversely, even if we want to give results for the situation where the initial choice of the player is fixed, and specifically it is door 1, as a mathematical device one can "pretend" that X is uniform random, *conditioning* on $X = 1$ at the end in order to read off the wanted results.

We'll denote the door opened by the host by H and the remaining closed door by Y. Thus, the triple (X, H, Y) represents the door numbers of the three doors listed according to the *manifest* or *observed* roles: door chosen by player, door opened by host, (other) door left closed by host. It is a random permutation of the triple $(1, 2, 3)$.

We already defined C to denote the door hiding the car. With H still being the goat-door opened by the host, let G be the number of the other goat-door. Again, (C, H, G) is a random permutation of $(1, 2, 3)$. This triple represents the door numbers of the three doors listed according to their *hidden* roles (largely hidden to the player, that is): door hiding the car, door opened by the host, other door hiding a goat.

By our assumptions so far, $\mathrm{P}(X = C) = 1/3$: the initially chosen door has probability $1/3$ to hide the car. This means that the two random permutations (X, H, Y) and (C, H, G) are either equal, with probability $1/3$, or unequal, with probability $2/3$. In the latter case, $(X, H, Y) = (G, H, C)$, switching gives the car, $Y = C$.

This result corresponds to the simple solutions with which we started – solutions which take no account of whether the host opened door 2 or 3.

A.2.1 From simple to conditional, by symmetry

Suppose the player initially chooses door 1. The odds on this door hiding the car are 2:1 against. The host now certainly opens a door revealing a goat, whether or not the initial choice was correct. The odds on door 1 hiding the car, given the host has opened an (as yet unidentified) door are still 2:1 against. Now we are informed that the door opened happened to be door number 3. We want to compute $\mathrm{P}(C = 1|X = 1, H = 3)$ and we know $\mathrm{P}(C = 1|X = 1) = 1/3$. By the law of total probability, $\mathrm{P}(C = 1|X = 1) = \mathrm{P}(C = 1|X = 1, H = 2)\mathrm{P}(H = 2|X = 1) + \mathrm{P}(C = 1|X = 1, H = 3)\mathrm{P}(H = 3|X = 1)$. By symmetry, $\mathrm{P}(C = 1|X = 1, H = 2) = \mathrm{P}(C = 1|X = 1, H = 3)$. Since $\mathrm{P}(H = 2|X = 1) + \mathrm{P}(H = 3|X = 1) = 1$ we find $1/3 = \mathrm{P}(C = 1|X = 1) = \mathrm{P}(C = 1|X = 1, H = 3)$.

In words, by symmetry the door opened by the host is conditionally independent of whether or not the car is behind door 1, given the player chose door 1. Hence, with $X = 1$ fixed, the conditional probability of winning by switching (given the number of the door opened by the host) equals the unconditional probability of winning by switching, 2/3.

A.2.2 From simple to conditional, by Bayes

Bayes' rule says: conditional odds equals prior odds times likelihood ratio. Let's apply this to the odds on whether or not door 1 hides the car, let's keep the initial choice of the player fixed as door 1 too, and let's take the piece of data whose likelihood ratio we are going to use to update the odds, as the identity of the door opened by the host: let's take it specifically as door 3.

Initially the odds are 2:1 against the car being behind door 1, the door chosen by the player.

If the car is indeed behind this door, the host is equally likely to open either of the other doors. Thus the probability that he opens door 3 is 1/2.

If on the other hand the car is not behind this door, it is equally likely behind doors 2 or 3. The host is forced to open the other door, to the door hiding the car. The chance that he opens door 3 is therefore also 1/2.

The Bayes' factor of likelihood ratio for the new piece of information is therefore 0.5/0.5 or 1. The new information is non-informative for the question at hand (of course, it *is* informative about whether or not the car is behind door 3, and also about whether or not it is behind door 2).

The posterior odds therefore remain 2:1 against. The player should switch.

A.3 The Holy Grail of MHP studies

The Holy Grail of MHP studies was for me, for a long time, to find a stupendously elementary proof of the fact that, in the case of a possibly biased host, there is no strategy for the player with a better overall success chance than 2/3. Hence the player might just as well ignore all specific door numbers and just switch.

Well, here is one, which I learnt from a wikipedia editor. I use rather a lot of words below: first to introduce the problem, and then to solve it. It could all be said in much fewer, but I hope this way there can be no misunderstanding.

A.3.1 The problem

Let's suppose the car is hidden behind one of the three doors by a fair randomization. The contestant chooses Door 1. Monty Hall, for reasons best known to himself, opens Door 3 revealing a goat. It can be shown using Bayes' theorem (or better still, Bayes' rule) that whatever probability mechanism is used by Monty for this purpose, the conditional probability that switching will give the car is at least 1/2. We know that the unconditional probability (i.e. not conditioning on the door chosen by the contestant, nor the door opened by Monty) is 2/3. Using the law of total probabiltity and the fact that all conditional probabilities of winning by switching are at least 1/2 proves that 2/3 overall win-chance can't be beaten.

Do we need Bayes to come to this conclusion? Surely, nobody in their right mind could imagine that there could exist some mixed strategy (sometimes staying, sometimes switching, perhaps with the help of some randomization device, and all depending on which doors were chosen and opened) which would give you a better overall (i.e. unconditional) chance than 2/3 of getting the car.

Is there an elementary proof? A short proof using words and ideas, no computations?

Yes there is, and I learnt it from a wikipedia editor. Here it goes.

A.3.2 The solution

Obviously the player only needs to consider deterministic strategies for himself. Now suppose Monty Hall makes his choice of door to open, when he does it at all, by tossing a possibly biased coin (a possibly different coin for each door). He might just as well toss his three coins in advance and just "look up" the action which is needed, if and when an action is needed. Now suppose the player also gets to see the results of the three coin tosses in advance. He now knows even more, so he cannot do worse (provided he uses all the available information as best as he can).

But now we are effectively in the so-called "Monty crawl" situation: this is the problem where the door which Monty would open in each of the three situations where he does have a choice (because the player is standing at the door hiding the car) is actually fixed in advance and known to the player.

We want to show that for the Monty crawl problem, there still is no strategy with an overall win-chance of more than 2/3.

There are just two cases to consider now.

Suppose the coin says that Monty would open door 3, if he had a choice between 2 and 3. Then whether the car is behind door 1 or door 2, Monty is certain to open door 3. His action tells us nothing so we may as well switch.

Suppose the coin says that Monty would open door 2, if he had a choice between 2 and 3. Then the fact he opens door 3 shows us that the car must be behind door 2, so we must switch.

Either way, we might as well switch. If we switch anyway, our overall win chance is 2/3. So this is the best overall win chance which is available for us for the Monty crawl problem. Hence the best in general.

A.3.3 Conclusion

To sum up: you can't do better than 2/3 overall because you can't do better than 2/3 in the situation that would be most favourable to you, Monty crawl. And therefore, because you can't do better than 2/3 overall, the chance of winning by switching must be at least 1/2 in each separate situation which you can distinguish (reductio ad absurdam and law of total probability).

We are in fact using Bayes in the Bayes' rule form, but only for the situation when the evidence we are given is certain under both hypotheses, and for the situation when it is certain under one, impossible under the other.

We are solving Monty Hall by use of the more simple problem Monty crawl. It reduces the problem just to an enumeration of two possible cases.

We are using the insight of all game theorists that one can always reduce everything to the extreme (deterministic) case. We are actually using game theory to prove an inequality about conditional probabilities!

Since 2/3 overall is the best you can do, and you can achieve that by always switching, it's a waste of time to look at the specific door numbers and a waste of time to figure out conditional probabilities with Bayes' theorem or whatever.