The Narrative Fallacy of Probability

Toshihiko Miura

1 Fallacious Comments on the Boy or Girl Paradox

Our intuition has often been diagnosed as bad at probability judgment. The most popular illustration is the Boy or Girl paradox, which is almost as well known as the Monty Hall problem. We will inspect how plausible this diagnosis is, although only for the Boy or Girl paradox. We will indicate the fallacies involved in presenting an applied version of the paradox, and make a modest protest against the prevalent diagnosis for human intuition. Let us begin with the following puzzle:

> Gary Foshee, a collector and designer of puzzles from Issaquah near Seattle walked to the lectern to present his talk. It consisted of the following three sentences: "I have two children. One is a boy born on a Tuesday. What is the probability I have two boys?" (Bellos, 2010)

The first thing we must recognize is the identity or the essence of the puzzle we are thinking about. The puzzle keeps its identity even if the day of the week involved in the sentences is changed. This is a matter of course because potential questioners will not always have children born on Tuesdays. The same applies to the gender.

Whether or not the information about being born on a Tuesday is relevant to the right answer is Foshee's purpose of the problem presentation. So first of all we had better consider the puzzle in its classic form, namely without the phrase 'born on a Tuesday.'

This puzzle could also have been "I have two children. One of them is a girl. What is the probability I have two girls?" Both of the versions are qualified for the same purpose. Depending on the questioner's family make-up, we are ready for either version. Because boy and girl are symmetrical, the gender difference does not change the logical structure of this puzzle. This same problem is sometimes presented as the girl-version such as "…one of them is a girl. What is the probability that there are two girls?" (Mlodinow, 2008; Marks & Smith, 2011).

So, we must obtain the answer valid for the sentences: "I have two children. One is a x. What is the probability I have two xs?" where x refers to either boy or girl. In this respect, "the right answer" given by Gary Foshee and others (including many mathematicians and psychologists) is wrong.

As the puzzle is potentially composed of the two versions, when we are informed of which version is actually told by the parent, who in this case is the same person as the questioner, we obtain information useful for revising the probability through the questioner's version choice. Initially, we had the probability 1/4 for two boys; after hearing the boy-version, we obtain the posterior probability as follows:

The questioner's choice was the boy-version	#B
The questioner's children are boys	BB
The questioner's children are an elder boy and a younger girl	BG
The questioner's children are an elder girl and a younger boy	GB
The questioner's children are girls	GG

$$P(\#B) = 1/2, \quad P(BB) = P(BG) = P(GB) = P(GG) = 1/4$$

$$P(BB \mid \#B) = P(\#B \mid BB)\, P(BB) / P(\#B) = (1 \times 1/4) / 1/2 = 1/2$$

Gary Foshee and the commentators, such as Bellos (2010), Devlin (2010), Briggs (2010) and Rutherford (2010), regarded $P(BB \mid \#B) = 1/3$. They are mistaken. Their calculation is probably like this: $P(BB \mid \#B) = P(\#B \mid BB)\, P(BB) / P(\#B) = (1 \times 1/4) / 3/4 = 1/3$. In short, they thought that $P(\#B) = 3/4$. Instead, we must distinguish between the questioner having at least one boy and the questioner selecting the boy-version as the question this time. #B is the latter, not the former.

In the first person presentation by Foshee and Bellos, the parent is identified in the beginning, and he chooses gender according to his knowledge of his children. Then, $P(\#B)$ is not 3/4, but 1/2, because it is untrue that 3/4 of all parents with two children would freely choose the boy-version with probability one.

To be sure, let us calculate $P(\#B)$.

$$P(\#B) = P(BB)\, P(\#B \mid BB) + P(BG)\, P(\#B \mid BG) + P(GB)\, P(\#B \mid GB) + P(GG)\, P(\#B \mid GG)$$
$$= (1/4 \times 1) + (1/4 \times 1/2) + (1/4 \times 1/2) + (1/4 \times 0) = 1/2$$

Foshee's presumption was $P(\#B \mid BG) = P(\#B \mid GB) = 1$, not 1/2, but the presumption is boy-biased. In the situation BG or GB, the parent as the questioner would choose the boy-version or the girl-version with equal probability.

Strictly speaking, this symmetric consideration applies only to a completely spontaneous gender-mention by the parent. Thus, in Devlin (2010), Keith Devlin refers to Gary Foshee's questioning. So his writing about himself: "I tell you I have two children and that (at least) one of them is a boy, and ask you what you think is the probability that I have two boys." may not be his free choice. He may be copying Foshee's expression because of the contingent truth of Devlin's own family make-up. If this is the case, the puzzle Devlin presented is actually the same as the other version: "I have two children. A stranger asked me if I have a boy. I answered yes. What is the probability I have two boys?"

In this version, there is no parental free choice of the version. Gender was first determined independently. If Devlin had at least one boy, he would have expressed the sentence equivalent to the one in #B. That is, $P(\#B \mid BG) = P(\#B \mid GB) = 1$. This is an excessively benevolent interpretation of Devlin's actual questioning though.

We cannot exactly conclude the ambiguity is due to the first person narrative mode. The third person narrative in the broad sense can be ambiguous in this respect. Let us suppose the puzzle were presented by a person in the story, different from the parent, such as “Bob told me ‘Alice has two children. One is a boy.’ What is the probability she has two boys?” Here Bob probably would have had some reason in advance to mention a boy, and Alice was selected as a qualified parent for the condition. Then, P (#B | BG) = P (#B | GB) = 1. This could also be false; Bob might have identified Alice first. As long as the person in the story selects the gender, the suspicion often remains that the parent was selected first. The issue is not always in the first person but in the choice of gender dependent on the knowledge about the children concerned.

The ambiguity-free presentation of this puzzle is the one where it is explicit that the choice of gender comes first, according to which a parent is randomly chosen, as we saw above as a benevolent interpretation in favor of Devlin. This is accomplished in two ways. First, a person in the story who does not know the gender should chose the gender in advance. Second, the real questioner outside the story, as in ordinary math problems, should directly choose the gender. In the second case, which has been called ‘the third person narrative’ in the narrow sense, the situation is described from the impersonal, transcending point of view, or no viewpoint, such as “Alice has two children. One is a boy. What is the probability she has two boys?” In this case, Alice is whoever satisfies the definition of the problem, or even a mere stipulated character. Here the phrases ‘x has two children. One is a boy’ refer to the necessary and sufficient condition of any parent employed for the puzzle. Alice can be understood as a random sample from the reference class. In that case, P (#B) = 3/4; that is 3/4 of all parents with two children could have been employed for the puzzle with the same probability.

About the Boy or Girl paradox, it is well-known that in general (under the most natural assumptions) the correct answer is 1/3 if a parent or a family is chosen at random and 1/2 if one child is encountered at random (Bar-Hillel & Falk, 1982). However, it does not seem to be recognized clearly (with a few exceptions such as Marks & Smith, 2011, Kaos, 2011) that in general the correct answer is 1/3 if the gender mentioned in the narrative is chosen at random and 1/2 if the gender mentioned in the narrative is determined causally by the children’s make-up. We will call this kind of insensibility the narrative fallacy of probability.

2 Applied Versions of the Paradox

We confirmed that the right answer to Foshee’s question “I have two children. One of them is a boy. What is the probability I have two boys?” is 1/2, not 1/3. Then, Foshee’s puzzle has the second part with a specific day of the week, which composes the essential aim of his challenge. In this essential part too though, Foshee and Devlin committed a narrative fallacy. Let us move to Foshee’s question with the phrase ‘on a Tuesday’ that we saw at the beginning, and his explanation from Bellos (2010).

Let’s list the equally likely possibilities of children, together with the days of the week they are born on. Let’s call a boy born on a Tuesday a BTu. Our possible situations are:

•When the first child is a BTu and the second is a girl born on any day of the week, there are seven different possibilities.

•When the first child is a girl born on any day of the week and the second is a BTu, again, there are

seven different possibilities.
•When the first child is a BTu and the second is a boy born on any day of the week, again there are seven different possibilities.
•Finally, there is the situation in which the first child is a boy born on any day of the week and the second child is a BTu – and this is where it gets interesting. There are seven different possibilities here too, but one of them – when both boys are born on a Tuesday – has already been counted when we considered the first to be a BTu and the second to have been born on any day of the week. So, since we are counting equally likely possibilities, we can only find six extra possibilities here.

Summing up the totals, there are 7 + 7 + 7 + 6 = 27 different equally likely combinations of children with the specified gender and day of birth, and 13 of these combinations are two boys. So the answer is 13/27, which is very different from 1/3. (Bellos, 2010)

Foshee and Bellos presuppose the answer 1/3 for the preliminary question without 'on a Tuesday.' If we understand the situation literally according to Bellos's presentation, the preliminary answer must not be 1/3 but 1/2. So it is gratuitously difficult to see why their logic for the final answer 13/27 is incorrect. To be clear on the puzzle, we should now distinguish the four possible variations including the one for which the final answer 13/27 is correct.

The common background of the following four situations is that we know the questioner has two children. The a priori stage is conditioned only by the common background. The first stage is further conditioned by the information that one is a boy. The second stage is further conditioned by the information that one is a boy born on a Tuesday. The ways of getting the information are different from one version to another.

Version 1 The questioner told us spontaneously: "One is a boy born on a Tuesday." (The literal interpretation of Foshee, Bellos and other's presentation.)
Version 2 The questioner told us spontaneously: "One is a boy." Someone ignorant of the situation then asked: "Do you have a boy born on a Tuesday?" The questioner answered yes.
Version 3 The questioner was asked by someone ignorant of the situation: "Do you have a boy?" The questioner answered yes, and then spontaneously added: "One is a boy born on a Tuesday."
Version 4 The questioner was asked by someone ignorant of the situation: "Do you have a boy born on a Tuesday?" The questioner answered: "Yes, one is a boy born on a Tuesday."

Finally the questioner asks us: "What is the probability I have two boys?" Now, what are the correct answers?

S…… report spontaneously
A …… asked and affirm

	The first stage "One is a boy"	The second stage "One is a boy born on a Tuesday."
Ver.1	S	S
Ver.2	S	A
Ver.3	A	S
Ver.4	A	A

Following Devlin's way of enumerating the cases, we obtain 14 possibilities for each child:

B-Mo, B-Tu, B-We, B-Th, B-Fr, B-Sa, B-Su
G-Mo, G-Tu, G-We, G-Th, G-Fr, G-Sa, G-Su

When the questioner tells you that one of her children is a boy born on a Tuesday, she reduces the possibilities, leaving the following:

(first child, second child)＝(B-Tu, B-Mo), (B-Tu, B-Tu), (B-Tu, B- We), (B-Tu, B-Th), (B-Tu, B-Fr), (B-Tu, B-Sa), (B-Tu, B-Su), (B-Tu, G-Mo), (B-Tu, G-Tu), (B-Tu, G- We), (B-Tu, G-Th), (B-Tu, G-Fr), (B-Tu, G-Sa), (B-Tu, G-Su), (B- Mo, B-Tu), (B-We, B-Tu), (B-Th, B-Tu), (B- Fr, B-Tu), (B-Sa, B-Tu), (B-Su, B-Tu), (G- Mo, B-Tu), (G-Tu, B-Tu), (G-We, B-Tu), (G-Th, B-Tu), (G- Fr, B-Tu), (G-Sa, B-Tu), (G-Su, B-Tu)

The questioner's final information was certainly "One is a boy born on a Tuesday." However, the datum to be used to reach the posterior probability of (B, B) is not the bare fact that one is a boy born on a Tuesday, but the fact that the questioner gave the information: "One is a boy born on a Tuesday." Let us call this meta-information #B-Tu. The value of P ((B, B)｜#B-Tu) depends on how we get #B-Tu. The processes of our obtaining #B-Tu are as follows. In Version 1, the questioner intentionally presented the conditional information: "One is a boy born on a Tuesday." In Version 4, by chance, the questioner has satisfied the given condition: "One is a boy born on a Tuesday." In Versions 2 and 3, the questioner intentionally performed half of the process. Especially in the second stage, only in Versions 1 and 3 is the information about a birthday intentionally selected by the parent as a true report on whatever his children's birthdays are.

Therefore,

P (#B-Tu | (B-Tu, B-Tu)) = 1
When X ≠ B-Tu,
P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu))
= 1/2 in Versions 1 and 3
= 1 in Versions 2 and 4

Notice that it is only in Version 4 where the parent is a random sample of the parents who have a boy born on Tuesday. In Version 1, the parent is a random sample of the parents who would spontaneously say "One is a boy born on Tuesday" at the specific occasion. As we saw earlier, Foshee, Bellos and Devlin fell into the narrative fallacy about this last distinction. They missed the distinction because of their confusing #B-Tu with the necessary and sufficient condition that "One is a boy born on a Tuesday." They regarded P (#B-Tu | (B-Tu, X)) and P (#B-Tu | (X, B-Tu) as 1. In Versions 1 and 3 though, the questioner might have used X, but not B-Tu, as a component of her question sentences with the chance of 1/2.

3 Calculations

Let us calculate each version one by one. The conclusion in each version will be the same whether or not it is calculated with P ((B, B)) at the a priori stage or at the first stage. So, for convenience, we will perform the a priori stage calculation for the continuous SS and AA patterns (Versions 1 and 4) and the first stage calculation for the articulated SA and AS patterns (Versions 2 and 3).

●Version 1 (SS pattern, so conveniently based on the a priori P ((B, B)))
P ((B, B)) = 1/4
P (#B-Tu) = 1/14
P (#B-Tu | (B-Tu, B-Tu)) = 1
P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu)) = 1/2 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (1/7 × 1/4) / (1/14) = 1/2

P (#B-Tu | (B, B)) is obviously 1/7 considering the symmetry of the 7 days, but a detailed analysis is as follows:

The category {(B, B)} involves 49 equiprobable cases concerning a day of the week. Among them, under the category {(B, B)with B-Tu} 13 cases are involved: (B-Tu, B-Mo), (B-Tu, B-Tu), (B-Tu, B- We), (B-Tu, B-Th), (B-Tu, B-Fr), (B-Tu, B-Sa), (B-Tu, B-Su), (B- Mo, B-Tu), (B-We, B-Tu), (B-Th, B-Tu), (B- Fr, B-Tu), (B-Sa, B-Tu), (B-Su, B-Tu). Among these only P (#B-Tu | (B-Tu, B-Tu))is 1, twice as much as any other P (#B-Tu | (B, B)), so P (#B-Tu | (B, B)) is calculated as (1 ×

1 + 12 × 1/2) / 49 = 1/7.

●Version 2 (SA pattern, so conveniently based on the first stage P ((B, B)))

P ((B, B)) = 1/2

P (#B-Tu) = 13/49 × 1/2 + 1/7 × 1/4 + 1/7 × 1/4 = 10/49

P (#B-Tu | (B-Tu, B-Tu)) = 1

P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu)) = 1 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (13/49 × 1/2) / (10/49) = 13/20

P (#B-Tu | (B, B)) is analyzed as follows.

The category {(B, B)} involves 49 equiprobable cases concerning a day of the week. Among them, under the category {(B, B) with B-Tu} 13 cases are involved, among which every P (#B-Tu | (B, B)) is 1, so P (#B-Tu | (B, B)) is calculated as 13/49.

●Version 3 (AS pattern, so conveniently based on the first stage P ((B, B)))

P ((B, B)) = 1/3

P (#B-Tu) = 1/7 ∵ The existence of a boy is already informed at the first stage, so the information given at the second stage must be limited to boys' birthdays, in order to keep the opportunity of the questioning with the non-trivial answer other than probability 0.

P (#B-Tu | (B-Tu, B-Tu)) = 1

P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu)) = 1/2 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (1/7×1/3) / (1/7) = 1/3

P (#B-Tu | (B, B)) is analyzed in the same way as we saw in Version 1.

●Version 4 (AA pattern, so conveniently based on the a priori P ((B, B)))

P ((B, B)) = 1/4

P (#B-Tu) = 1 – (13/14 × 13/14) = 27/196 ∵ In every case other than (X, Y) for X, Y ≠ B-Tu, the questioner says yes, and starts giving the puzzle.

P (#B-Tu | (B-Tu, B-Tu)) = 1

P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu)) = 1 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (13/49 × 1/4) / (27/196) = 13/27

P (#B-Tu | (B, B)) is analyzed in the same way as we saw in Version 2.

For reference, let us examine the ordinary third person narrative (presentation from no point of view or the transcending viewpoint) of ordinary math problems, as the fifth version, where the information on the children is provided from outside the story.

●Version 5 (First, based on the a priori P ((B, B)))
P ((B, B)) = 1/4
P (#B-Tu) = 1 – (13/14 × 13/14) = 27/196 ∵ In every case other than (X, Y) for X, Y ≠ B-Tu, the questioner can truly describe the parent employed for the puzzle.
P (#B-Tu | (B-Tu, B-Tu)) = 1
P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu)) = 1 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (13/49 × 1/4) / (27/196) = 13/27

P (#B-Tu | (B, B)) is analyzed in the same way as we saw in Version 2.

(Next, based on the first stage P ((B, B)))
P ((B, B)) = 1/3
P (#B-Tu) = 27/147 ∵ The same reason as in Version 2.
P (#B-Tu | (B-Tu, B-Tu)) = 1
P (#B-Tu | (B-Tu, X)) = P (#B-Tu | (X, B-Tu) = 1 X ≠ B-Tu

P ((B, B) | #B-Tu) = P (#B-Tu | (B, B)) P ((B, B)) / P (#B-Tu) = (13/49 × 1/3) / (27/147) = 13/27

P((B, B))	a priori	first stage	second stage
Ver.1	1／4	1／2	1／2
Ver.2	1／4	1／2	13／20
Ver.3	1／4	1／3	1／3
Ver.4	1／4	1／3	13 ／27
Ver.5	1／4	1／3	13 ／27

Lynch (2011) stated "the question that Foshee actually answered was: 'Of all two-child families with at least one child being a boy born on a Tuesday, what proportion of those families have two boys?' The correct answer to the question he actually posed is P = 1/2" (p. 72). Lynch is right. Marks & Smith (2011) also correctly indicated the parallel error of Mlodinow (2008) on the pioneering applied Boy or Girl paradox featuring a girl named Florida.

Foshee, Bellos and Devlin's calculations are correct only for Versions 4 and 5 (and only structurally for Version 2). This has been intuitively obvious at the start, because in their actual version (Version 1) the questioner appears to have given no informative knowledge at the second stage. The impression of no informative value in Versions 1 and 3 came not only from our knowledge that birthday and gender are mutually independent, but also from our tacit recognition that a birthday was never mentioned independently of the original information source on gender. So the possibilities were never tried or tested by any random sampling.

4 The Lesson Neglected

To make sure of our conclusion on Version 1, let us try a thought experiment. Let each parent all over the world, who has just two children, issue a puzzle about his or her children in the template form "One is a x born on a y. What is the probability I have two xs?" (x refers to gender; y refers to a day of the week).

Half of them would choose boy for x, whom we will call #B parents. #B parents are composed of 100% of parents in the category (B, B), 50% of parents in the category (B, G) and (G, B) and 0% of parents in the category (G, G). So P ((B, B) | #B) = 1 / (1 + 1/2 + 1/2 + 0) = 1/2, not 1/3. Here Foshee, Bellos and Devlin committed their first narrative fallacy.

One seventh of #B parents would choose Tuesday for y; we will call them #B-Tu parents. #B-Tu parents are composed of 100% of parents in the category (B-Tu, B-Tu), 50% of parents in the category (B-Tu, B-X), (B-X, B-Tu), where X ≠ Tu, (B-Tu, G) and (G, B-Tu), and 0% of parents in other categories. So P ((B, B) | #B-Tu) = (1 × 1 + 12 × 1/2) / (1 × 1 + 12 × 1/2 + 14 × 1/2 + 0) = 1/2. The final probability P ((B, B) | #B-Tu) is the same as P ((B, B) | #B). Information about the day is irrelevant, as our naïve intuition had told us. Here Foshee, Bellos and Devlin committed their second narrative fallacy.

Versions 1 and 4 correspond to the two versions of Monty Hall problem. One has the rule that Monty should always open an empty door; the other has the rule that Monty should open a door at random. In the former version, the apparent datum does not give any expected information, although in the Monty Hall problem, due to its specific situation, the player can obtain a hint about what to do. The latter version has a risk that the game will stop en route. In the Monty Hall problem, Monty might open the winning door. In Version 4, the parent might say "No, I do not have a boy born on a Tuesday. What is the probability I have two boys? " which is not the same puzzle as the version "Yes, I have a boy born on a Tuesday. What is the probability I have two boys? "

When I heard Version 1 for the first time, I though the information about the day was irrelevant and that Version 4 is where the information of the day is useful. My intuition, and probably many other people's naïve intuition too, seems much more reliable than Devlin and Briggs warn. Actually, both Devlin and Briggs confess their first impression that the correct answer is 1/2 upon their first reading. Of course they were right, but their mathematics spoiled their intuition! They strangely failed to utilize the lesson of the Monty Hall problem, namely, the importance not just of the contents but also the lesser-known importance of the process of getting the information. Still more strangely, even Rutherford, who pointed out consciously many misunderstandings on the original version of the Boy or Girl paradox, and mentioned the first person versions as his versions 3, 9 and 10 (Rutherford, 2010, pp. 168, 170), does not criticize the authors who suggest that the correct answer to Version 1 is 1/3.

After his unfortunately mistaken commentary, Devlin presented an applied quiz:

> Now that your intuition has been primed, let me leave you with this problem. I tell you I have two children, and (at least) one of them is a boy born on April 1. What probability should you assign to the event that I have two boys? If you think that is going to be too cumbersome, simply tell me whether the probability is close to 1/2 or to 1/3, or to some other simple fraction, and provide an estimate as to how close. (Once more, you should assume all birth possibilities are equally likely, ignoring in particular the well known seasonal variations in actual births.) (Devlin, 2010)

Here again, our naïve intuition (non-primed intuition) tells us that the information on the birthday is irrelevant, and the intuition is definitively right. Because this problem is presented in the Version 1 style, the answer Devlin assumes to be correct is doubly wrong. First, as we saw, the probability we should assign to the parent (whether Keith Devlin himself or not) having two boys after we hear his words: "I tell you I have two children, and (at least) one of them is a boy." is not 1/3 but 1/2. Devlin should not have presented the problem in the first person. To ensure 1/3 as the correct answer, the information that one of them is a boy needs to be given independently of the parent's and all other person's knowledge on his children's gender. Second, after we hear the additional information "one of them is a boy born on April 1," the probability we should assign to the event that the parent has two boys should not ever change. To ensure the revision of probability as Devlin hopes, the information that a boy was born on April 1 needs to be given independently of the parent's self-enumeration.

Devlin (and Foshee and others) should have presented the problem in the third person (exactly, from no viewpoint, such as Version 5) or in the second person (introduced as a response to an interrogative such as Version 4).

The correct answer to Devlin's last problem is just 1/2. Or, benevolently taken as Version 3, the correct answer is 1/3, and remains 1/3 all along. The information on the birthday is irrelevant. The answer Devlin intended to be correct, namely "very close to 1/2," has the same value as the answer we will get if the problem is taken as a Version 4 type. Unfortunately, the Version 4 calculation is beside the point if we follow the literal meaning of Devlin's presentation.

5 Mathematical Intuition and Narrative Skill

If taken as the Version 4 type, our naïve intuition does not protest against the claim that the information on the birthday is very influential. As long as we understand correctly the described situation and how it was described (ex. who told it with what knowledge and intention), our intuition about probability seems much more reliable than sometimes pointed out by mathematicians and psychologists. The problem is perhaps much less in our probability judgment and more in our ability to read and understand correctly what situation is explicitly described and how. This diagnosis of ours is basically homologous to the current criticisms of the wordings of the original puzzle (ex. Rutherford, 2010), and of the birthday version (ex. Lynch, 2011).

When we read a novel in which the characters' background circumstances are not explicitly described, we can usually make suitable inferences, namely, almost correct probabilistic inferences, without setting ourselves to mathematical calculations. Our intuition works well enough if we already adequately understand the narrated facts and the narration itself, including whether it is reported by the person in the story or by a person outside the story. When we are in the context of art or literature, we are not as vulnerable to the narrative fallacy of probability as we are in the context of mathematics. We are generally careful about the narrative mode--point of view, tense, tone, mood, degree of spontaneity, etc.--when appreciating artwork, but such carefulness is felt pointless in mathematics.

Perhaps Devlin and others were mistaken not in their diagnosis on human intuition but rather about the narration of the puzzle. They were especially indifferent to or underestimated the influence of the narrative point of view on the interpretation of the text. Maybe many of what we call tricky probability problems (Briggs, 2010) should actually be classified under literature or narrative theory, or even practical narrative skills, not under the category of mathematics. We tend to be wrong more often on the reading, writing and reporting than on the probability inferences.

References

Bar-Hillel, Maya & Falk, Ruma. 1982. "Some teasers concerning conditional probabilities." *Cognition*. 11 (2): pp. 109–122.

Bellos, Alex, 2010. "Magic numbers: A meeting of mathemagical tricksters." *New Scientist*, *Issue 2762,* pp. 44-49.

Briggs, William M. 2010. "Tricky Probability Problem: Chance of Two Sons, One Born on Tuesday?" June 29, 2010. http://wmbriggs.com/post/2553/

Devlin, Keith. 2010. "Probability Can Bite." April 2010. Devlin's Angle. http://www.maa.org/external_archive/devlin/devlin_04_10.html

Lynch, Peter. 2011. "The Two-Child Paradox: Dichotomy and Ambiguity." *Irish Mathematical Society Bulletin 67* , pp.67-73.

Marks, Stephen & Smith, Gary. 2011. "The Two-Child Paradox Reborn?" *Chance Vol. 24,* pp. 54-59.

Kaos, 2011 *PANDORA BOX* 2011-10-02 (in Japanese) http://d.hatena.ne.jp/nakaoshi_kaos/20111002

Mlodinow, L. 2008. *The Drunkard's Walk: How Randomness Rules Our Lives*. Pantheon Books.

Rutherford, Mel D. 2010. "On the use and misuse of the "Two Children" brainteaser." *Pragmatics & Cognition Vol. 18:1,* pp. 165–174.