Bag of Tricks

You have a bag with 50 playing cards in it. Some cards have blue backs and some have pink backs. The cards can either be from red suits (hearts and diamonds) or from black suits (clubs and spades). We know the following facts about the cards in the bag:

- There are 20 cards with pink backs: 3 from red suits (hearts and diamonds) and 17 from black suits (clubs and spades).
- There are 30 cards with blue backs: 6 from red suits and 24 from black suits.

You will draw cards from this bag (without looking).

1. When you draw a card from the bag, what's the chance that it will have a pink back? You may express the chance as a percentage (like 70%) or as a fraction (like 7/10).

2. What's the chance that it's from a red suit?

Now, we split up the cards with pink backs from the cards with blue backs, and put them in separate bags. You can pick which of these bags you will draw from.

3. If you draw a card from the bag with blue-backed cards, what's the chance that it's from a red suit?

4. If you really want a red-suited card, should you draw a card from the bag with pink-backed cards or from the bag with blue-backed cards?

5. What if, instead of splitting the cards by the color of their backs, we split them by the color of the suit? If you draw a card from the bag with all of the red-suited cards, what's the chance that it has a pink back?

Now we're going to change what's in the bag. We take all the pink-backed cards from 2 of the original bags and put them into a box with the blue-backed cards from only one of the original bags. Now there are 40 pink-backed cards and 30 blue-backed cards, but the portions of red-suited and black-suited cards with each type of back stay the same. So:

- There are 40 cards with pink backs: 6 from red suits (hearts and diamonds) and 34 from black suits (clubs and spades).
- There are 30 cards with blue backs: 6 from red suits and 24 from black suits.

6. If you draw a pink-backed card from this box, what's the chance that it's from a red suit? Is the chance higher, lower, or the same as it was with the initial bag?

7. If you draw a red-suited card from this box, what's the chance that it has a pink back? Is this higher, lower, or the same as it was with the initial bag?

8. What if we took the pink-backed cards from 10 of the original bags and put them together in a box with the blue-backed cards from just one of the original bags - what would be the chance that a red-suited card drawn from this box has a pink back?

Now let's make a game out of this (in fact, there used to be a game show that worked like this): You are shown a bag with 3 poker chips in it. Two of the chips are blue and one is white. The white chip is worth $100, but the blue chips are worth nothing.

In the game, you first draw one chip from the bag, but you're not allowed to look at it. Then, the game show host looks in the bag and always finds and takes out a blue chip. Finally, he gives you a choice between keeping the chip you already drew (but haven't looked at yet) or trading your chip for the one that's still in the bag.

9. Should you keep your chip or take the one from the bag? Does it matter? Why or why not?

Let's move away from bags and cards and chips and on to something more practical: Doctors have discovered a rare new disease that is easy to treat if it's caught early, but patients don't start showing symptoms until much later. They have also developed a test for the early stages of the disease, but this test isn't always accurate. The doctors performed the test on 10,000 people and then watched to see which ones eventually got sick. Of these 10,000 people, only 1% (100 people) turned out to have the disease. And the test gave the correct result (it said that the healthy people were healthy and the sick people were sick) 90% of the time.

10. How many people were sick and had the test say they were sick? How many were sick, but the test said they were healthy?

11. How many people were healthy and had the test say they were healthy? How many were healthy, but the test said they were sick?

12. If you took the test and it said you were sick, what is the chance that you actually are sick?

Studies have shown that doctors often fail to answer these types questions correctly, and wildly over-estimate the chance that a person really has a disease when a test says they have it.

Here's another situation: You are on the jury for a hit and run trial in which a taxi driver is accused of hitting a pedestrian. A witness claims he saw a green taxi hit the pedestrian, and the prosecutor's case relies heavily on this witness's testimony. But it was a rainy night, and, while there's no reason for the witness to lie, he could be wrong about what he saw.

You know a few other things about the situation: only 2 companies run taxis in the area - one company only has green taxi cars, and the other company only has blue taxi cars. The company with blue cars is much more successful, so you know that 85% of the taxis on the road that night were blue. In addition, the defense attorney had the witness take a vision test where he tried to identify the color of a bunch of cars in conditions similar to the conditions of the crime. In this test, the witness correctly identified the color of a car (as blue or green) 80% of the time.

13. Before doing any calculations, take a minute to guess - given that the witness claims he saw a green taxi, do you think it's more likely that the taxi was really green or that it was really blue?

14. Now calculate it: what is the chance that the taxi that hit the pedestrian was green? Was your guess right?

One more situation: These days, many sports require their athletes to take drug tests. But, as with all medical tests, drug tests are never 100% accurate. Let's say that Major League Baseball wants to test players for a drug that is very rarely used, but they have a very accurate test for it. Only 0.5% of players (1 out of 200) actually uses this drug. When the test is given to people who do use the drug, it comes back positive 99% of the time. When it's given to people who don't use the drug, it comes back negative 98% of the time.

Note: often tests have different percentages of "false negatives" - that is, the test comes back negative on someone who did take the drug, than they give of "false positives" - that is, the test comes back positive on someone who did not take the drug. The test in this example has a "false negative" rate of 1% (the test will come back negative on 1% of people who do take the drug), and a "false positive" rate of 2% (the test will come back positive on 2% of people who do not take the drug).

15. If a player takes this test and it comes back positive, what's the chance that he actually took the drug?

16. What would the false positive and false negative rates have to be in order to make the chance that a person who tested positive actually took the drug to be 80%? You can change either the false positive rate or the false negative rate.

See http://yudkowsky.net/bayes/bayes.html if you'd like to learn more about this!