Intermediate Math Circles
March 21, 2012
Probability

This week and next week we'll be talking about probability and random events. We'll start with some warm-ups, and discuss how there is a unique limiting shape if you repeat the same event enough times. We'll work on some problems about *conditional probability*: how likely is some event to happen *given that* some other event happened? Next week, we'll discuss expected values; along the way we'll see random walks, game theory, and gambling problems.
You can contact Dave with questions at dagpritc@uwaterloo.ca. Solutions will go on the website.

Warm-up

If you roll a die, what is the probability that the number (on top) is a prime number?
The number on top is equally likely to be 1, 2, 3, 4, 5, or 6. Three of those (2, 3, 5) are prime. So the probability is 3/6, or 1/2 in simpler terms.
What is the general rule for computing probabilities? What is the minimum and maximum possible probability?
If there are T equally likely outcomes, and G of them are good, then the probability of a good outcome is G/T. Every probability is between 0 and 1, and one way to argue this is to notice that the minimum value of G is 0, and the maximum value of G is T.
If you flip two coins, what is the probability that one of them comes up heads and one of them is tails?
There are two approaches, but which one is correct? You could argue that there are three possible outcomes (both heads, both tails, mixed) or four possible outcomes (HH, HT, TH, TT), giving an answer either of 1/3 or 2/4. The second answer is correct. The first answer is not the correct one because the "three possible outcomes" listed in the first group are not all equally likely. Why? Imagine that one of the coins is a penny and one of the coins is a nickel. Then there are clearly 4 possible outcomes: (penny H & nickel H, penny H & nickel T, penny T & nickel H, penny T & nickel T), and they are equally likely. Or you can imagine flipping a "first" coin and a "second" coin.
If two events A and B are completely unrelated to each other (e.g., A is "the penny comes up heads," B is "the nickel comes up heads") what is the relation between the probability that A happens (written $\mathsf{Pr}[A]$), the probability that B happens, and the probability that both of them happen (written $\mathsf{Pr}[A \text{ and } B]$)?
The formula is $\mathsf{Pr}[A \text{ and } B] = \mathsf{Pr}[A] \times \mathsf{Pr}[B]$. We say that A and B are *independent*. An example of non-independent events are $A =$ "the penny comes up heads" and $B =$ "the penny comes up tails;" then $\mathsf{Pr}[A] = \mathsf{Pr}[B] = 1/2$ but $\mathsf{Pr}[A \text{ and } B] = 0$.

We roll two dice and add up the numbers that show on top. What is the probability that the sum is 6?

We need to think again about which "equally likely events" we should use as a basis for our calculations. Similar to the coin problem, the right way to think about it is using ordered pairs of numbers from 1 to 6. (You could roll one die first and another second, or colour one die fir green and another sepia.) In class I'll show a grid of these pairs and we can count that out of the $36 = 6 \times 6$ outcomes, 5 of them add up to 6. So the answer is $5/36$.

Draw a diagram showing how likely each sum is.

If we do the same counting with all the other possible sums, we see that the lowest possible sum is 2 which has probability $1/36$, $\Pr[\text{sum is } 3] = 2/36$, $\Pr[\text{sum is } 4] = 3/36$, up to a maximum of $\Pr[\text{sum is } 7] = 6/36$, then going back down to $\Pr[\text{sum is } 12] = 1/36$.

What is the probability that the sum is even? (Hint: there is a long solution but also a short solution.)

The probability of getting an even sum can be calculated in two ways. One of them is $1/36 + 3/36 + 5/36 + 5/36 + 3/36 + 1/36 = 18/36 = 1/2$. Another way is much simpler. Notice that when you roll a single die, it is even $1/2$ the time and odd $1/2$ the time. So roll the first die. If it was even, the probability that you get an even number on the second roll and the sum is even, is $1/2$. If the first die was odd, the probability that you get an odd number on the second roll and the sum is even is $1/2$. In other words, no matter what happened on the first roll, the second die has probability $1/2$ of making the sum even, which proves the final answer is $1/2$. We'll come back to this a little later.

Aside: The Normal Distribution

The previous diagram was pretty stylish. What does it look like with 3 dice, 10 dice, or 100 dice? We'll check it out on the computer.

You can access the code we used at `http://sagenb.org/home/pub/4598/` and if you can change it yourself too. Click on "log in to edit a copy" and you can create an account, then make a copy and edit it.

Here's a diagram showing 100 dice. Here's the command — you can edit 100 to look at a different number of dice.

$$\texttt{makePlot(repeat(die, 100), min = 300, max = 400)}$$

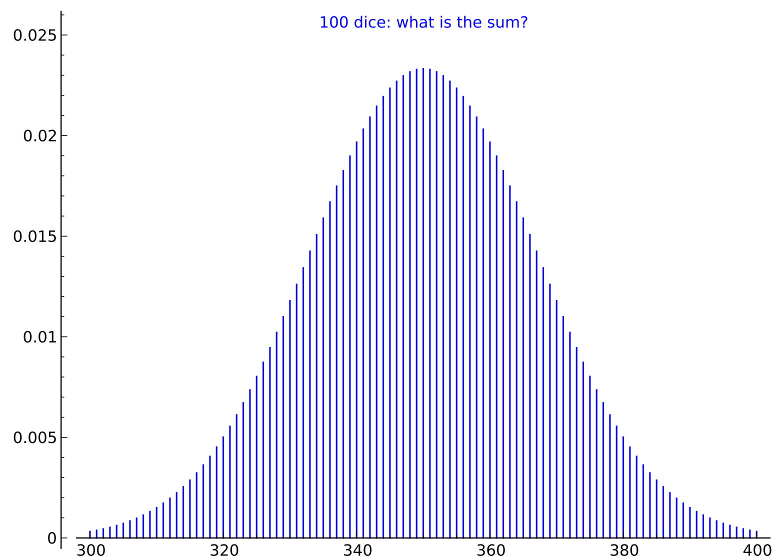

What about coins? Coins are pretty flat compared to dice so maybe a very different shape would emerge. Let's take 5 coins as an example. We'll fill in the left diagram using data collected in class. Then we'll compute the exact probabilities and fill in the diagram on the right.

We need to recall: the number of ways to select k out of n items, written $\binom{n}{k}$ and called "n choose k," equals

$$\frac{n}{k}\frac{n-1}{k-1}\cdots\frac{n-k+1}{1}=\frac{n!}{(n-k)!k!}.$$

The probability of getting k heads is $\binom{5}{k}/2^5$. This is because the possible outcomes are all possible 5-sequences of H and T, like (H, H, H, T, H). There are 2^5 of these sequences; the number of them with k Hs is $\binom{5}{k}$ by the definition of "choose." So the probabilities of 0, 1, 2, 3, 4, 5 heads are respectively 1/32, 5/32, 10/32, 10/32, 5/32, 1/32. In the experiment, the more flips you do, the closer you'll get to these ideal probabilities.

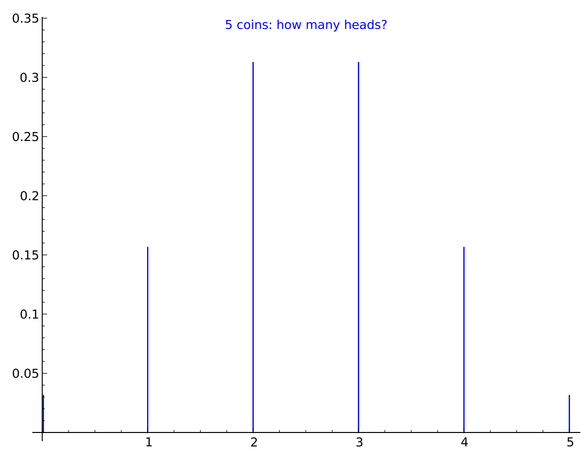

On the computer we'll try to experiment again with 100 coins instead of 100 dice. The shape is pretty similar, except that it has been moved and squished a little bit.
To see 100 coins, we run

$$\texttt{makePlot(repeat(coin, 100))}$$

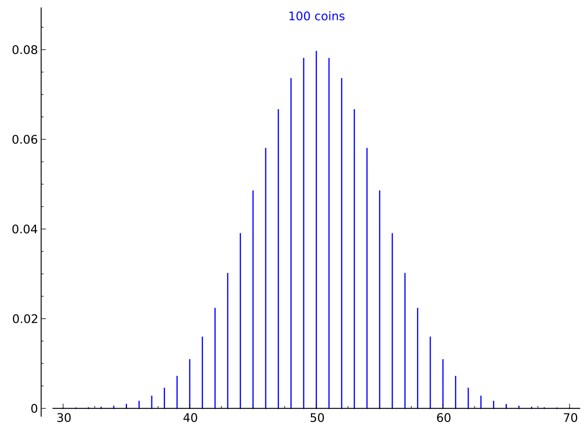

What about if we use a biased coin, or an unfair die?
Surprisingly, we get the same shape! Here are 100 coins, each biased 2/3 towards heads:

$$\texttt{makePlot(repeat([1/3, 2/3], 100))}$$

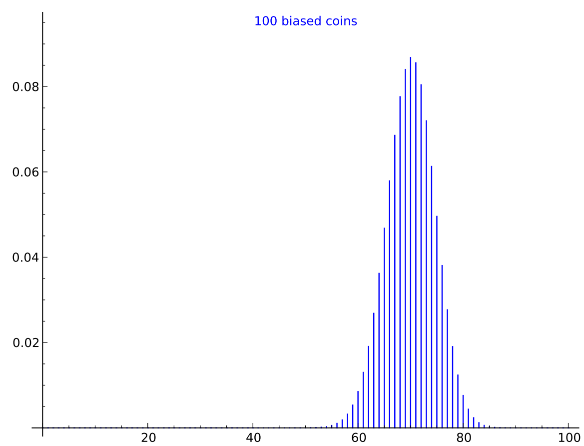

This is an example of the *Central Limit Theorem*: for any random variable, take N independent copies of it and add them up; as N tends to infinity, the resulting sum has a *normal distribution* with a particular mean (horizontal center) and standard deviation (squishing factor). *It doesn't matter what the original distribution's "shape" was at all!*
Next week we'll talk some more about the precise meaning of the mean and standard deviation. For *normal distributions* there is a shortcut: the mean is where the maximum/center point is located, and one standard distribution away is where the *inflection points* lie. An inflection point has maximum slope: it is where biking uphill would be hardest. Draw the mean and inflection points on the diagram below of a normal distribution.

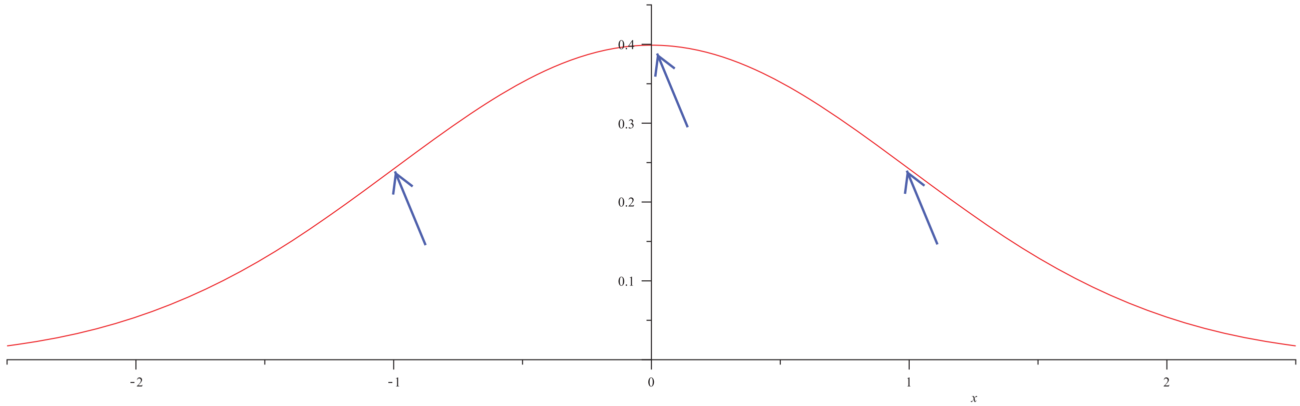

The central arrow points to the maximum point $(0, 0.39)$ — the mean is 0. The other two arrows point to the inflection points at $(\pm 1, 0.24)$, and the standard deviation is 1.

Law and Politics, Part I

The people in Probabilia are all either blue and green. The population is 90% green and 10% blue. There was a robbery last night, and Officer Ophelia saw the robber running away. She thinks that the robber was blue. However, at night time she has a 20% chance of making a mistake in trying to figure out the colour of anybody she sees.

Bob, a blue person found near the scene of the crime, is being held in custody. Ophelia argues that the robber was 80% likely to be blue, since her vision is only wrong 20% of the time. Bob's lawyer argues that Bob should be released because *the robber was more likely to be green than blue*! Why?

Conditional Probability

Conditional probability means "the probability that something happens, **if** something else happens." This helps to keep track of the calculations used in the blue/green question, and has many other uses.

Here are two equivalent ways to describe conditional probability. First, assume we are keeping track of a set of "equally likely outcomes," like in the original definition of probability. Then *the conditional probability of A, conditioned on B*, written $\mathsf{Pr}[A \mid B]$, is defined as

$$\mathsf{Pr}[A \mid B] = \frac{\text{the number of outcomes where } A \text{ and } B \text{ both happened}}{\text{the number of outcomes where } B \text{ happened.}}.$$

It means, *if a random experiment causes B to be true, what fraction of the time will A also be true?* For example, let X be the number on the top of a die. Then

$$\mathsf{Pr}[X \text{ is prime} \mid X \text{ is a non-square number}] = \frac{3}{4}$$

because X can be a non-square number in 4 ways (2, 3, 5, 6), and a non-square prime number in 3 ways (2, 3, 5). What is

$$\mathsf{Pr}[X \text{ is a non-square number} \mid X \text{ is prime}]?$$

The value of $\Pr[X \text{ is a non-square number} \mid X \text{ is prime}]$ is 1. This is because there are 3 ways for X to be prime, and 3 ways for X to be both prime, and a non-square.
An equivalent way to define conditional probability is,

$$\Pr[A \mid B] = \frac{\Pr[A \text{ and } B]}{\Pr[B]}.$$

Can you prove the two definitions are equivalent? (Hint: use the original definition of probability.)

Two classical problem involving conditional probabilities

Q. You flip three coins in the air. If all three are the same (all heads or all tails) you win \$1 from me, otherwise you lose \$1 to me. I argue that this is a fair game, for the following reason: we know two of the coins are guaranteed to be identical, so there is a 50% chance that the third coin is too, thus you're just as likely to win as to lose. Am I telling the truth?
I am lying, in order to steal your money! We already did some calculations with coins and we know that the probability they are all the same is $\binom{3}{3}/2^3 + \binom{3}{0}/2^3 = 1/8 + 1/8 = 1/4$. But what is the problem with my argument? It has to do with the guarantee that two out of the three coins will agree. Let M be the event that the first two coins were the same. It's certainly true that if the first two coins match (M), the third, unflipped coin has a 50% chance of matching them. So

$$\Pr[\text{you win} \mid M] = 1/2.$$

However, if the first two coins do not match, while it is true that the third coin will match one of them, we don't get to re-flip anything, and we already know that they won't all be the same. So you are guaranteed to lose: $\Pr[\text{you win} \mid \text{not } M] = 0$. Using the LOTP, we can confirm the actual probability of winning:

$$\Pr[\text{you win}] = \Pr[\text{you win} \mid M]\Pr[M] + \Pr[\text{you win} \mid \text{not } M]\Pr[\text{not M}] = \frac{1}{2} \times \frac{1}{2} + 0 \times \frac{1}{2} = 1/4.$$

When we discuss the solution, we'll write down a way to calculate $\Pr[A]$ from $\Pr[B], \Pr[A \mid B]$, and $\Pr[A \mid \text{not } B]$. It is called *the law of total probability (LOTP)*.

$$\Pr[A] = \Pr[A \mid B] \times \Pr[B] + \Pr[A \mid \text{not B}] \times \Pr[\text{not } B].$$

Do you see how to prove this? Start by arguing that

$$\Pr[A] = \Pr[A \text{ and } B] + \Pr[A \text{ and } (\text{not } B)].$$

Q. A cat had two kittens. If one of the kittens was female, what's the probability that both kittens were female?
Intuitively, the answer is 1/2, but this answer is too high. There are four equally likely outcomes for the kittens' genders: MM, MF, FM, FF where F=female and M=male. Let B denote the event that one of the kittens is female, and A denote the event that they are both female. Since (A and B) happens in only one outcome, and B happens in three outcomes, we have that $\Pr[A \mid B] = 1/3$.

Back to Probabilia

We will assume that the robber was a random person from the population, and therefore has 90% probability of being green, and 10% probability of being blue. What is the probability that the robber was blue and Officer Ophelia saw them as blue? What is the probability that the robber was green and Officer Ophelia saw them as blue? Finally, what is the probability that the robber was actually blue (like Bob), conditioned on the robber being seen as blue?

We'll use the Law Of Total Probability, broken down into simple steps. The probability that the robber was blue, and Ophelia saw them as blue, is

$$\mathsf{Pr}[\text{was blue and looked blue}] = \mathsf{Pr}[\text{was blue}]\mathsf{Pr}[\text{looked blue} \mid \text{was blue}] = 10\% \times 80\% = 8\%.$$

The probability that the robber was green, and Ophelia saw them as green, is

$$\mathsf{Pr}[\text{was green and looked blue}] = \mathsf{Pr}[\text{was green}]\mathsf{Pr}[\text{looked blue} \mid \text{was green}] = 90\% \times 20\% = 18\%.$$

So, the total probability that the robber looked blue is

$$\mathsf{Pr}[\text{looked blue}] = \mathsf{Pr}[\text{was green and looked blue}] + \mathsf{Pr}[\text{was blue and looked blue}] = 18\% + 8\% = 26\%$$

and finally, the probability of being blue, given that Ophelia saw someone blue, is

$$\mathsf{Pr}[\text{was blue} \mid \text{looked blue}] = \frac{\mathsf{Pr}[\text{was blue \& looked blue}]}{\mathsf{Pr}[\text{looked blue}]} = \frac{8\%}{26\%} = 4/13 \sim 0.307.$$

So the lawyer may argue the robber was 69.3% likely to be green, and only 30.7% likely to be blue!
(Note that this argument of the lawyer makes an assumption that it is a random person who committed the crime. You could try to make a more refined argument by coming up with an estimate of the demographics in the vicinity of the crime.)

More Problems

1. A cat had a black kitten and an orange kitten. The black kitten was male. What's the probability that the orange kitten was male?

2. If we both pick random numbers from 1 to 10 (each one with probability 1/10), what is the probability yours is higher than mine? That they are equal?

3. I have three red M&Ms and two blue ones. I eat one at random and give two random ones to you. If you get at least one red M&M, what is the probability I ate a blue one? Is it higher than 2/5?

4. If you draw two cards from a standard deck of 52 cards, what's $\mathsf{Pr}[\text{1st card is an ace}]$? $\mathsf{Pr}[\text{2nd card is an ace}]$? $\mathsf{Pr}[\text{2nd card is an ace} \mid \text{1st card is an ace}]$? $\mathsf{Pr}[\text{both cards are aces}]$?

5. **Let's make a deal.** You are on a game show where there are three shiny boxes. The host, Guy Smiley, shows you that one box contains a gold coin, and the two other boxes contain turnips. (Assume you do not want a turnip, even though they are very nutritious.)

The boxes are then closed and shuffled randomly. The rules say first that you pick one box, without opening it. Guy Smiley will then open a *different* box, and in particular he always opens a box containing a turnip, and eats it. That leaves two boxes: the one you picked, and another one you didn't pick. Finally, you have the choice to keep your originally chosen box, or switch to the other unopened box. If the one you end up with has the gold coin you keep it!

Is it better to keep your original box, or switch boxes? Hint: even though you have 2 choices and you don't know which one has the turnip, it is not true that both boxes are 50% likely to contain the gold coin.

6. Assume $\Pr[A]$ is not 0 or 1. Show that the following three statements are either all true, or all false: $\Pr[A \mid B] > \Pr[A], \Pr[B \mid A] > \Pr[B], \Pr[A \mid B] > \Pr[A \mid \text{not } B]$. We say that events A and B are *positively correlated* when any of the above inequalities are true. A major component of statistics looks at correlations (which is different from causation).

7. It turns out YOU were arrested for being a robber. But the prince has made you an offer. He gave you two bags, 10 skull coins, and 10 happy face coins. First, you get to arrange the coins in the bags however you like; each coin must be in a bag. Second, a random bag will be selected. Third, a random coin from that bag will be selected. If it is a skull you go to jail, but if it is a happy face you go free. How can you minimize your probability of going to jail? (Hint: it is much less than 1/2.)

8. Which of these three modified dice is the best, when they are compared head-to-head and the larger number wins? Die A has values $\{2, 2, 4, 4, 9, 9\}$ printed on its sides. Die B has sides $\{1, 1, 6, 6, 8, 8\}$. Die C has sides $\{3, 3, 5, 5, 7, 7\}$.

9. **Law and Politics, Part II.** The country of Probabilia has a population with more men than women. The president has declared that a new policy will be started in order to increase the number of female births:

   Policy : When a family has a son, they are not allowed to have any more children.

   The President thinks this is a good idea because then you can have some families with several daughters, but every family will have at most one son. Will it actually help achieve the goal of having more daughters born overall than sons?

1. This time, the presence of colours means the answer is actually 1/2, not like the 1/3 we had with the earlier cat question. There are exactly two outcomes where the black kitten is male: (black male, orange male) and (black male, orange female).

2. You can compute the answer by listing all 100 possibilities, but here is a shorter answer. No matter what number you pick, note that I have a 1/10 chance of picking the same one as you. So overall the probability that we pick the same number is 1/10. The remaining $1 - 1/10 = 9/10$ of the time we pick different numbers, and by symmetry, my number is larger just as often as yours is larger:

$$\Pr[\text{mine is bigger} \mid \text{our numbers are not equal}] = 1/2.$$

   So yours is higher $1/2 \times 9/10 = 19/20$, or 45% of the time.

3. Let EB be the event that the one I eat is blue, and GR be the event that you get at least one red M&M. We want to find $\Pr[EB \mid GR]$. We know that $\Pr[EB] = 2/5$. Conditioned on EB, we have $\Pr[GR \mid EB] = 1$; and we have $\Pr[GR \mid \text{not } EB] = 5/6$. Overall, by LOTP, $\Pr[GR] = 1 \times \frac{2}{5} + \frac{5}{6}\frac{3}{5} = \frac{9}{10}$. The conditional probability of $\Pr[EB \mid GR]$ is then $\Pr[EB \text{ and } GR]/\Pr[GR] = \frac{2}{5}/\frac{9}{10} = 4/9$. This is slightly higher than the probability of EB if you knew nothing at all about GR.

4. The first card is an ace with probability $4/52 = 1/13$. The second card is the same (since there's nothing different about the first card and the second card). However, conditioned on that the first card is an ace, there are only 3 aces left out of the remaining 51 cards. So the conditional probability $\Pr[\text{2nd card is an ace} \mid \text{1st card is an ace}]$ is $3/51 = 1/17$. Finally, the probability that both of the first two cards are aces is $\frac{1}{13}\frac{1}{17} = \frac{1}{221}$.

5. This is a famous problem! Here is a useful morsel of intuition: at the start, you had a $1/3$ chance of picking the gold coin. The turnip did not affect this. Either keeping your original choice, or switching, always wins. So switching, you win with probability $2/3$.
   Here's a more detailed explanation. Let G be the event that you initially selected the box with the gold coin; note $\Pr[G] = 1/3$. Conditioned on G, switching boxes will make you lose for sure. But conditioned on not-G, switching boxes will make you win for sure. So, $\Pr[\text{switching wins}]$ equals

$$\Pr[\text{switching wins} \mid G]\Pr[G] + \Pr[\text{switching wins} \mid \text{not } G]\Pr[\text{not } G] = 0 \times 1/3 + 1 \times 2/3$$

   which is $2/3$, and you can argue diectly or by a similar computation that not-switching wins with probability $1/3$.

6. First,

$$\Pr[B \mid A] > \Pr[B] \Leftrightarrow \frac{\Pr[A \text{ and } B]}{\Pr[A]} > \Pr[B] \Leftrightarrow \Pr[A \text{ and } B] > \Pr[A]\Pr[B]$$

   which, by symmetry, also is equivalent to $\Pr[A \mid B] > \Pr[A]$.
   For the second equivalence, we give an informal explanation but it is not hard to make it into a formal one. By LOTP and the fact that probabilities are non-negative, $\Pr[A]$ is a mixture (strictly convex combination) of $\Pr[A \mid B]$ and $\Pr[A \mid \text{not } B]$. Thus, either all three or equal, or they are distinct and $\Pr[A]$ is in the middle of $\Pr[A \mid B]$ and $\Pr[A \mid \text{not } B]$.

7. You should distribute the coins as follows: one bag just gets one happy face coin; all other coins go in the other bag. Then you go to jail only with probability $1/2 \times 0 + 1/2 \times 10/19 = 5/19$.
   Extra tricky solution: as above, but put the bag with one coin inside the other bag too. (So there is a coin inside a bag inside a bag, and 19 other coins in the outer bag.) Then you go to jail only with probability $1/4$.

8. This is a trick! No die is best. A beats B $5/9$ of the time, B beats C $5/9$ of the time, and C beats A $5/9$ of the time.

9. This policy does not have any effect on the ratio of boys and girls being born! We'll talk about it next week.