CALCULATED FICTION

MATH WORKSHOP 4

October 19, 2004

Since it's probability week, we'll use a random process to put you into groups this time. As always, work at the same speed as the rest of your group. Remember to keep your work on your own paper and turn it in at the end of the workshop.

Each group should get a deck of cards and two dice. You'll also need a coin for flipping.

ASTRALAGI: IT'S SHEEP FOR "HEY, I WAS USING THAT"

Let's start out by gathering some empirical data. In the first few activities, we'll conduct repeated *Bernoulli trials*, experiments in which there are two outcomes, one which occurs with probability p and one which occurs with probability $1 - p$. For example, for one flip of a fair coin, $p = .5$.

ACTIVITY

1. Flip a coin 8 times, keeping track of how often it comes up heads. Let's call one sequence of 8 flips a trial. As a group, conduct at least 60 trials, and combine your data into one histogram. (Some division of labor will be useful here! Everybody can conduct trials on their own and then share their results.) Be sure you make a copy of the histogram on your own paper.

   Now look at the drawing of the bell curve on p. 286 of TLOM. How does your histogram compare with the bell curve?

2. On a graph, plot the probabilities corresponding to row 8 of Pascal's Triangle, as discussed on TLOM pp. 275-6. (You may want to use your calculator to turn them into decimals.) How does your graph compare to your histogram from activity 1? How does it compare with the bell curve?

3. This time, let's say a trial is 5 rolls of a single die; keep track of how often the number that comes up is bigger than 4 (i.e., is 5 or 6). What's the probability of success in a single roll?

   As a group, conduct at least 40 trials, and combine your data into one histogram. How does it compare with the bell curve? How do you explain the difference between this histogram and the one from activity 1?

The moral of this story is that repeated Bernoulli trials yield a histogram that looks like some version of the normal curve. Hence the Greeks were wrong in thinking that there was no order to be found in chance events (TLOM, p. 272); this pattern shows up consistently everywhere you look. That's right, *everywhere*. Where are you looking right now?  See, there it is.

ACTIVITY

4. Now let a trial be a roll of two dice; keep track of the sum of the two numbers that show up. Conduct at least 50 trials and collect your results into a histogram.

Now make the histogram that would result if you conducted 36 trials and each possible outcome came up exactly once. Does this look like the normal curve? Should it? Are we still conducting Bernoulli trials in this case?

ANOTHER REASON WHY MATHEMATICIANS ARE USEFUL

The experimental approach we've been taking so far is nice for approximating certain probabilities and, sometimes, for finding confirmation of our computations.

The Monte Carlo method, for example, is the method of using a computer to provide an approximate solution to a mathematical problem. The computer performs random sampling experiments and uses the relative rate of successes to estimate the answer. The method even applies to problems with no obvious relationship to probability. For example, it can be used to estimate the value of π as follows:

Draw a 2 inch by 2 inch square, and draw a circle with radius 1 inside of it. Now, the area of the circle is $\pi r^2 = \pi\ (1 \text{ in})^2 = \pi \text{ in}^2$, and the area of the square is 4 in^2. Hence the ratio of the circle's area to the square's area is $(\pi \text{ in}^2)/(4 \text{ in}^2) = \pi/4$. To estimate this ratio, we simply have the computer choose a large number of points at random inside the square, and we keep track of how many of those points fall inside the circle. The ratio of points inside the circle to points chosen is then a reasonable estimate of the value of $\pi/4$. For example, if we have the computer choose 40,000 points and 31,407 of them fall inside the circle, then our estimate of $\pi/4$ is 31407/40000 = .785175. (The actual value is close to .785398.)

The downside of the experimental approach is that it doesn't give us the exact answer; all it can do is approximate the correct answer. It's quite important, for example, that we have a better way to find the value of π, since a lot depends on us getting it right.

Another problem is that it's often not practical to estimate the probability of rare events by experimentation.

ACTIVITY

5. A poker hand consists of five cards dealt from a standard deck of 52 (4 suits, 13 ranks in each suit; ask if you need more explanation). A flush is a hand in which all five cards are in the same suit. A straight is a hand in which the ranks are consecutive (for example, 8-9-10-J-Q, from any suits). A straight flush is both a straight and a flush. Deal out 50 poker hands; how many of them are flushes? Are you comfortable using these data to estimate the probability of being dealt a straight flush? What if you dealt 100 hands? 1000 hands?

This suggests that it would be useful to have some theoretical tools for computing exact probabilities. (In case you're interested, the actual probability is being dealt a straight flush is about .0000153908, so about 1 in 64,974 hands is a straight flush.)

JUST THE FACTS, MA'AM

Forthwith, some basic facts about and tools for working with probability.

Recall from lecture yesterday that all probabilities are between 0 and 1, inclusive. A probability of 0 means the event happens with 0% likelihood and is therefore impossible; a probability of 1 means the event happens with 100% likelihood and is therefore certain. The numbers in between correspond to percent likelihoods in the same way. For example, a probability of .375 means the event happens with 37.5% likelihood – that is, in the long run, if the experiment were repeated a large number of times, we could expect the number of successes to be 37.5% of the number of trials. Note that this means *no probability can be less than 0 or greater than 1*; you can't have a probability of, say, 2, because you can't run an experiment 100 times and get 200 successes.

Our basic formula for computing theoretical probabilities is this one:

$$P(event) = \frac{\#\{successes\}}{\#\{outcomes\}}$$

We require that *all outcomes are equally likely.* Recall the dice example from yesterday's lecture.

To see how this formula works, let's look at some examples.

Flipping a coin. What's the probability of getting heads on a flip of a single fair coin? Here there are two outcomes, which we can call H and T (for heads and tails). If we want P(heads), there is only one success – H – and two outcomes, so P(heads) = ½ = .5.

Flipping a coin twice. What's the probability of getting one H and one T if we flip a fair coin twice? (I'll stop saying "fair"; assume unless I say otherwise that all the objects we use are fair.) To begin, let's decide what the outcomes are. There are two candidates for the outcome sets: {HH, HT, TH, TT} (where the order tells us the order in which they came up on the coin) and {2 H's, 2 T's, 1H & 1T}. These have different sizes, and give us *different answers*. On the first view, there are two successes (HT and TH), so the probability is #{successes}/#{outcomes} = 2/4 = 1/2. On the second view, there is only one success and there are three outcomes, so the probability is 1/3. Which is correct? (If you can't tell, do a few experiments with a coin to see.) 1/2 is correct, because the outcomes in that case are equally likely. 1/3 is incorrect because the outcomes on that view are *not* equally likely. (And if you don't believe me, I have a dice game I'd like to play with you...)

Rolling a die. What's the probability of getting a 2 on the roll of one die? Here there are six outcomes: 1,2,3,4,5,6. There's only one success: 2. Thus the probability is 1/6.

What if we instead ask about the probability of getting an odd number on the roll of one die? There are still six outcomes, but now there are three successes: 1,3,5. Thus the probability in this case is 3/6 = 1/2.

Winning the lottery. Suppose the lottery works as follows: six winning numbers are chosen from a group of 40 numbers. You win if your ticket has all six winning numbers. What is the probability of winning this lottery if you have one ticket? Here it seems clear that there's only one success – having the one correct combination of winning numbers. What's an outcome, though? Well, an outcome is a choice of six numbers from the 40 available. For each such choice, there's a different possible outcome (and it makes a big difference to whoever has the winning ticket if you change to a different choice!). How many ways are there to choose six numbers from among 40?

[Pause.]

C(40,6)! No, just C(40,6). (C(40,6)! is a verrrrry large number. See, this is a joke I have made. End transmission) Remember combination numbers? They and permutation numbers are quite useful for calculating probabilities. Here we see that the probability of winning the lottery is #{successes}/#{outcomes} = 1/C(40,6) = 1/3838380 ≈ 2.605×10^{-7} = .0000002605. Go buy your tickets now, kids! There's a reason they call the lottery a tax on the stupid. (Zing!)

ACTIVITY

6. If an experiment consists of flipping a fair coin 10 times in a row, how many outcomes are there? Make sure your outcomes are equally likely.

7. If an experiment consists of dealing a poker hand, how many outcomes are there? There are only four ways to get a royal flush in poker. What's the probability of being dealt a royal flush?

What about the probability of two events? Let's think about he probability that A happens or B happens or both. Here's the rule:

$$P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$$

Why isn't it just P(A) + P(B)? We can see that that's wrong by thinking about an example.

ACTIVITY

8. Say an experiment is two coin flips. Let event A be "anything but two T's", and let event B be "anything but two H's". Determine P(A) and P(B). Why can't P(A or B) = P(A) + P(B)? Devise an experiment to confirm the formula above for these two events, and carry it out.

One last rule: What's the probability that event E *doesn't* happen? Well, with probability 1 it either happens or it doesn't happen, so 1 = P(E) + P(not E), so (with a little rearranging)

$$P(\text{not } E) = 1 - P(E).$$

We can also say P(failure) = 1 - P(success).

ACTIVITY

9. What's the probability of not being dealt a royal flush in one hand of poker?

PROBABILITY AS A MEASURE OF BELIEF

We've been interpreting the probability of an event as measuring how frequently the event will occur in the long term. But what about statements like "There's a 75% chance that Mom will be mad when she sees what we did to the family room." How do we interpret such a claim? After all, either she will be mad or she won't. How could we ever know if the claim is true or not?

Essentially, we can't know for sure. In this case we can't repeat the experiment enough times to get empirical data (if we did, Mom would be likely to take preventive action, like locking us in the storm cellar every time she leaves the house). Hence such intuitive estimates represent something different from the ones we've been considering. Here the 75% is a measure of my belief; it is a *subjective* probability.

Although subjective probabilities are horses of a different color, they're still horses, so we'll forge ahead with the assumption that subjective probabilities satisfy all of our rules for probabilities.

ACTIVITY

10. After she calmed down, I asked Mom what she thought the chances were of
    (a) rain today,
    (b) rain tomorrow,
    (c) rain both today and tomorrow,
    (d) rain today or tomorrow (or both).

    After a minute, her answers were 30%, 40%, 20%, and 60%. Are these consistent with our rules for probability? If not, why not, and do you think I should tell her?

PROBABILITEMS! WAIT... PROBLEMILITY! NO... PROBLABEMITY?

Here are some probability problems to get your gourd steaming.

ACTIVITY

11. Following the computation from class, find the probability that two people in a group of 40 have the same birthday. (As in class, ignore February 29.) Can you find any shortcuts for the computation?

12. There are 38 numbers on a roulette wheel: 0, 00, and 1 through 36. 0 and 00 are green, while the other numbers are evenly split between red and black (so there are 18 of each color).

    On a red/black bet at a roulette wheel, most casinos pay 2 to 1. That is, if you bet $1 that the next number will be a red one and you're right, you get $2 from the casino. If you're wrong, you lose your $1.

    On average, what's the probability that you win on a red/black bet? How much would the casino pay if the game were exactly fair?

13. *The Monty Hall Problem*. On a game show, you are asked to select one of three doors to open; there's a large prize behind one of the doors, and there's nothing behind either of the other two. Once you select a door, the host, who knows where the prize is, opens one of the other two doors *that he knows is a losing door*. (If both are losing doors, he selects at random.) Then he asks you if you would like to switch doors, from your original choice to the other unopened door.

    Should you switch? What is your probability of winning if you switch, and what is it if you don't? Devise an experiment to verify your results, and carry it out.

    Hint: It might be helpful to think about a slightly different version of the problem. What if there are 100 doors, and after you choose one, the host opens all but one of the others? So he opens 98 doors (all losers, natch). In this case, do you think the probability is different if you switch than if you don't?

14. We always talk about fair coins, but in practice we can't verify that even a single coin

is fair in the real world. So what should we do when it's really important to have a fair coin? Let's figure out how to simulate a fair coin using a possibly unfair one.

To begin, suppose we have two identical coins, both possibly unfair. How can we use these two coins to simulate a fair coin? That is, how can we use these two coins to conduct a Bernoulli trial that gives success with probability .5 and failure with probability .5?

Now use your two-coin answer to say how you could do it with only one coin.

15. An evil probabilist shows you two envelopes and tells you that one of them contains two times as much as the other one, but he doesn't tell you which one is which. He lets you choose one of the envelopes and open it. It turns out to have $10 inside. Now he gives you the opportunity to take the other envelope instead of the current one (and give back the $10 you already have), because the second box could contain twice as much (i.e., $20).

    To maximize the amount of money you end up with on average, should you choose the second envelope, or should you stick to your first choice?

    Hint: If you have $10, and you could double this with a chance of 1/2, or half it with a chance of 1/2, one would expect an average of 1/2 * $20 + 1/2 * $5 = $12.5 (so you would expect to gain $2.5)!

    Second Hint: The first hint is wrong. Your job is to figure out why.

REFERENCES

Some of the materials for this workshop were taken/adapted from these sources:

- *Discrete Mathematics and Its Applications*, 5th edition, by Kenneth H. Rosen
- *A First Course in Probability*, 5th edition, by Sheldon Ross

PROJECT

For this project, you may work with whomever you like.

1. Finish anything from the workshop that you didn't get to in class. (Give some thought to the problems at the end, but don't feel like you have to come up with a solid solution to each one; there may be too much here to do.)

2. My good friends Yelena and Zarya each have two children, none of them twins. Yelena says (truthfully) that at least one of her children is a boy. Zarya says (truthfully) that her older child is a boy.

Given the above information, which of these two friends is more likely to have TWO boys? For each of them, what is the probability of having two boys? (Assume that a child chosen at random is just as likely to be a boy as a girl.) Devise a strategy using cards to verify your results experimentally, and carry out your experiments.

3. You and I decide to play the following fun game: I write down 100 numbers on slips of paper and put them into a hat. You don't know what numbers I write down, or how big they are; you just know that there are 100 of them. You take slips of paper out of the hat one at a time and look at the numbers written on them.

Your task is to identify the largest number in the hat. The catch is that you don't necessarily get to look at all of the numbers before you make a claim about which one is largest. At any point, you can take a number out of the hat, look at it, and say, "This is the largest number." You only get to say that once, and you have to say it about the number you just drew out of the hat. If you're correct, you win and I pay you \$$x$. If you're incorrect (that is, if the number you say it about isn't the largest one), you pay me \$1.

Some questions:
(a) What's a good strategy to use when playing this game?
(b) What does x have to be in order to make the game fair? (This is a hard question; rather than trying for an exactly correct answer, make a general guess based on your answer from (a).)