ECON 455, Discussion Section 4

TA: Shane Auerbach (sauerbach@wisc.edu) ; Date: 02/20/15

Office: SS 6470. OH: Wed 8:00-9:30am; Thu 4:15-5:45pm; or by appt.

1. (Experiment) The point of these instructions is in the solutions so as not to bias the experiment.

   (a) Write a random sequence of twenty H's and T's, i.e. outcomes from twenty sequential fair coin tosses.

   (b) What is the most (and least) likely total number of T's and H's in the random sequence?

   Solution: Obviously the most likely is ten of each, because there are more possible sequences that involve ten of each than any other particular ration.

   (c) What is the most (and least) likely random sequence of T's and H's?

   Solution: Trick question: All are equally likely!

   (d) Other (more interesting) questions in the solutions, so as not to bias the experiment.

   Solution: I am interested in the occurrence of *streaks* in sequential coin tosses. If the outcome of n tosses in a row is heads (H), that constitutes a streak of n. Similarly if we have n consecutive tails (T) outcomes. In this experiment, I will compare the longest streaks in your sequence with the longest streaks when we actually draw randomly (or pseudo-randomly, actually) in simulations.

   We can calculate the average longest streak when we're tossing the coin n times using the following (terribly inelegant) Matlab code which tosses a coin n times ten thousand times. We then average over the longest streaks in each of those ten thousand trials to get an expectation for the longest streak in the given trial.

```
clear
clc
n=20;
samplesize=10000;
longstreak=zeros(1,samplesize);
for k=1:samplesize
v = randi([0 1],n,1)';

if or(isequal(v,ones(1,n)),isequal(v,zeros(1,n)))==1
    longstreak(k)=n;
else
    diffsforones = diff([0 v 0]);
    one_starts = find(diffsforones == 1);
    one_ends = find(diffsforones == -1);
    one_streak = max(one_ends-one_starts);

    diffsforzeros = diff([1 v 1]);
    zero_starts = find(diffsforzeros == -1);
    zero_ends = find(diffsforzeros == 1);
    zero_streak = max(zero_ends-zero_starts);

    longstreak(k)=max(zero_streak,one_streak);
end
end
```

```
meanlongstreak = mean(longstreak);
```

Running this gave me a mean longest streak of 4.67 (it'll give a different, but hopefully similar, number each time). The point of interest here is that average longest streak from people in the class was about 3.6, which is significantly shorter than 4.67.

Does this perhaps illustrate any of the biases we've discussed in class? You could say that it illustrates the Gambler's Fallacy in a far more nuanced setting. You're all smart enough to know that after getting two heads in a row, the probability of getting another heads on the next coin toss is still a half. However, when we ask you to create a random sequence (and there is no wrong answer to that exercise because all such sequences are equally likely, but when we aggregate your answers they may seem inconsistent with randomness), you tend to have alternation between H and T more often than a random sequence would in expectation. In a sense, there's a little voice in your head saying:

I've just written down 3 H's in a row. That's fine, but it has to go back to T at some point. And 4 H's in a row is pretty unlikely, so in some sense it looks more random if I go back to T now.

When a random number generator gets three H's in a row, it doesn't care whether the fourth H makes the sequence look *less random*, it just draws it with probability 1/2. The point is that even if you're very aware of the gambler's fallacy, you probably still exhibit it at some subconscious level that we can get at through experiments like this.

And now, for fun (and candy), let's guess the average longest streak if we're tossing a coin 100 times. What about 10000 times? For 100 times, the average longest streak is about 7. For 10000 times, the average longest streak is about 13.7. Your first reaction to this might be that you're surprised at how low it is. But keep in mind that the probability of getting 14 heads in a row is $(0.5)^{14} = \frac{1}{16000}$. With 10000 tosses, you have a lot of opportunities to get streaks that long, sure, but they're still very unlikely (each exponentially less likely than the last). Interestingly, while thinking that 13.7 is surprisingly low, I bet if I asked you to write out 10000 H's and T's, few people would have streaks longer than 13 (or even close to that).

Final note: You might ask why we simulated an answer instead of just solving for the expectation analytically. The answer is that doing so is very tricky. Here's a very related Stack Exchange question which should demonstrate the complexity of this problem:

http://math.stackexchange.com/questions/59738/probability-for-the-length-of-the-longest-run-in-n-bernoulli-trials)

Final final note: Aren't computers amazing? My crappy computer did an experiment involving flipping a coin ten thousand times. And it did this experiment ten thousand times (for a total of one hundred million coin tosses!) in eleven seconds. It could probably do it in a small fraction of that had I coded it more efficiently too. Life lesson: get good at programming. OK - technically the computer wasn't actually flipping coins ten million times, rather it was converting a sequence of ten million random numbers it already had into zeros and ones then searching for the longest streaks. Still pretty cool.

2. Don't forget the likelihood-ratio formulation of Bayes' Rules:

$$\frac{P(A|n)}{P(B|n)} = \frac{P(A)}{P(B)} \frac{P(n|A)}{P(n|B)}$$

In words, we can say that the posterior equals the prior multiplied by the signal. Recall the model introduced in the lectures. A professor is trying to decide if a student is good or bad. The prior belief that the student is good is 1/2. Then the professor receives signals (tests), where g represents a good test performance and b represents a bad test performance. Once a professor believes it more than 50% likely that a student is good (bad), she will misinterpret a bad (good) test performance as good (bad) with probability q. Also good students do good on tests and bad students do bad on tests with $P(g|G) = P(b|B) = p$.

(a) Suppose $q = 1/2$ and $p = 3/4$. Given a naive professor observes $\tilde{g}\tilde{g}\tilde{g}$, what is her posterior probability that the student is good?

Solution:

$$\frac{P(G|\tilde{g}\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g}\tilde{g})} = \frac{P(G)}{P(B)} \frac{P(\tilde{g}\tilde{g}\tilde{g}|G)}{P(\tilde{g}\tilde{g}\tilde{g}|B)}$$

Let's take the items one-by-one:

- $P(G) = P(B) = \frac{1}{2}$
- $P(\tilde{g}\tilde{g}\tilde{g}|G) = \left(\frac{3}{4}\right)^3 = \frac{27}{64}$
- $P(\tilde{g}\tilde{g}\tilde{g}|B) = \left(\frac{1}{4}\right)^3 = \frac{1}{64}$

Therefore

$$\frac{P(G|\tilde{g}\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g}\tilde{g})} = 27$$

The last step is to calculate the probability from the ratio. Since

$$P(G|\tilde{g}\tilde{g}\tilde{g}) + P(B|\tilde{g}\tilde{g}\tilde{g}) = 1,$$

it follows that $P(G|\tilde{g}\tilde{g}\tilde{g}) = 27/28$

(b) Now do the same exercise for a sophisticated professor.

Solution:

$$\frac{P(G|\tilde{g}\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g}\tilde{g})} = \frac{P(G)}{P(B)} \frac{P(\tilde{g}\tilde{g}\tilde{g}|G)}{P(\tilde{g}\tilde{g}\tilde{g}|B)}$$

Let's take the items one-by-one:

- $P(G) = P(B) = \frac{1}{2}$
- $P(\tilde{g}\tilde{g}\tilde{g}|G)$.
  - Given your prior was 1/2 and you perceived the first test as good, it must have been a good test performance. However, either of the second and third tests could have been bad and misread. Therefore

$$\begin{aligned} P(\tilde{g}\tilde{g}\tilde{g}|G) &= P(ggg|G) + P(gbg|G)P(\tilde{g}|b) + P(ggb|G)P(\tilde{g}|b) + P(gbb|G)(P(\tilde{g}|b))^2 \\ &= \frac{27}{64} + 2 \cdot \frac{9}{64} \cdot \frac{1}{2} + \frac{1}{64} \cdot \left(\frac{1}{2}\right)^2 = \frac{145}{256} \end{aligned}$$

- $P(\tilde{g}\tilde{g}\tilde{g}|B)$.

– Given your prior was 1/2 and you perceived the first test as good, it must have been a good test performance. However, either of the second and third tests could have been bad and misread. Therefore

$$\begin{aligned}P(\tilde{g}\tilde{g}\tilde{g}|B) &= P(ggg|B) + P(gbg|B)P(\tilde{g}|b) + P(ggb|B)P(\tilde{g}|b) + P(gbb|B)(P(\tilde{g}|b))^2 \\ &= \frac{1}{64} + 2\cdot\frac{3}{64}\cdot\frac{1}{2} + \frac{9}{64}\cdot\left(\frac{1}{2}\right)^2 = \frac{25}{256}\end{aligned}$$

Therefore

$$\frac{P(G|\tilde{g}\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g}\tilde{g})} = \frac{145}{25} = \frac{29}{5}$$

And we can solve that $P(G|\tilde{g}\tilde{g}\tilde{g}) = \frac{29}{34}$. While that's still pretty high, it's significantly lower than $\frac{27}{28}$ as the sophisticated professor accounts for her bias.

(c) Suppose the sophisticated professor's prior probability that the student is good is $\frac{50}{100}$. What is her posterior given she observes $\tilde{g}\tilde{g}$?

Solution: Since the professor started off unbiased, the first test results must have been a g. The second could have been a b and been misinterpreted. Therefore

$$\begin{aligned}P(\tilde{g}\tilde{g}|G) &= \frac{9}{16} + \frac{3}{16}\frac{1}{2} = \frac{21}{32} \\ P(\tilde{g}\tilde{g}|B) &= \frac{1}{16} + \frac{3}{16}\frac{1}{2} = \frac{5}{32}\end{aligned}$$

Therefore $\frac{P(G|\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g})} = \frac{21}{5}$, so $P(G|\tilde{g}\tilde{g}) = 21/26$

(d) Suppose the sophisticated professor's prior probability that the student is good is $\frac{51}{100}$. What is her posterior given she observes $\tilde{g}\tilde{g}$?

Solution: Now the professor is biased even before seeing the first test result. Therefore, it could be that the student actually did poorly on either or both of the tests and it was misperceived.

$$\begin{aligned}P(\tilde{g}\tilde{g}|G) &= \frac{9}{16} + 2\frac{3}{16}\frac{1}{2} + \frac{1}{16}\left(\frac{1}{2}\right)^2 = \frac{49}{64} \\ P(\tilde{g}\tilde{g}|B) &= \frac{1}{16} + 2\frac{3}{16}\frac{1}{2} + \frac{9}{16}\left(\frac{1}{2}\right)^2 = \frac{25}{64}\end{aligned}$$

Therefore

$$\frac{P(G|\tilde{g}\tilde{g})}{P(B|\tilde{g}\tilde{g})} = \frac{P(G)}{P(B)}\frac{P(\tilde{g}\tilde{g}|G)}{P(\tilde{g}\tilde{g}|B)} = \frac{51/100}{49/100}\frac{49/64}{25/64} = \frac{51}{25}$$

It follows that her posterior is $P(G|\tilde{g}\tilde{g}) = \frac{51}{76}$.

(e) Compare your answers to (c) and (d). What is counter-intuitive here? Where is it coming from in the model?

Solution: It's counter-intuitive that the posterior in (d) is lower than the posterior in (c). The professor observed the same signals in both cases and had a higher prior in (d), yet her posterior in (d) is lower than in (c). What's generating this in the model is that when the sophisticated professor has a prior of 1/2, she trust that the first signal she receives is an accurate reflection of the test performance.

In (d) however, even though she initially believes the student is more likely to be good, that belief (and the bias it generates) causes her to doubt the accuracy of the first signal she receives.

(f) What is the professor's posterior if it is equally likely that good students and bad students do well on tests and she has no bias? Answer in words, appealing to the likelihood-ratio formulation of Bayes' Rule above.

Solution: Then $P(n|A) = P(n|B)$, so the posterior just equals the prior, i.e. the evidence is uninformative.

3. (Mandatory drug testing)(Angner 5.28) In July 2011, the state of Florida started testing all welfare recipients for the use of illegal drugs. Statistics suggest that some 8 percent of adult Floridians use illegal drugs; let us assume that this is true for welfare recipients as well. Imagine that the drug test is 90 percent accurate, meaning that it gives the correct response in nine cases out of ten.

(a) What is the probability that a randomly selected Floridian welfare recipient uses illegal drugs and has a positive test?

Solution: Let's let D mean on drugs, so $P(D) = 0.08$. Let $+$ denote a positive test. Then

$$P(D \cap +) = P(D)P(+|D) = 0.08 \cdot 0.9 = 0.072$$

(b) What is the probability that a randomly selected Floridian welfare recipient does not use illegal drugs but nevertheless has a positive test?

Solution:

$$P(\neg D \cap +) = P(\neg D)P(+|\neg D) = 0.92 \cdot 0.1 = 0.092$$

(c) What is the probability that a randomly selected Floridian welfare recipient has a positive test?

Solution:

$$P(+) = P(D \cap +) + P(\neg D \cap +) = 0.164$$

(d) Given that a randomly selected Floridian welfare recipient has a positive test, what is the probability that he or she uses illegal drugs?

Solution:

$$P(D|+) = \frac{P(D \cap +)}{P(+)} = \frac{0.072}{0.164} = 18/41 \approx 44\%$$

(e) If a Florida voter favors the law because he thinks the answer to (d) is in the neighborhood of 90 percent, what fallacy might he be committing?

Solution: Base rate neglect!