Cauchy, beta, gamma and the infinite expectation paradox

18.600 Problem Set 7, due April 19

Welcome to your seventh 18.600 problem set! This problem set features problems about beta, Gamma, and Cauchy random variables. These random variables are not quite as ubiquitous as others we have discussed (exponential, uniform, normal, Poisson, binomial) but they are fun and do come up. The problems should help you internalize the definitions and some of the standard interpretations. In particular we will explore the idea of beta distributions as Bayesian posteriors for the p associated to a biased coin, where your prior was uniform but then you saw a few tosses.

Many of you are familiar with *Pascal's wager*. The general idea is that if choosing A over B comes with a finite cost but a positive probability (however small) of an infinite payoff, then one should always choose A. Pascal's conclusion was that if living a virtuous life leads (with even a tiny probability) to an eternal reward, then it is a worthwhile sacrifice to make. A common criticism is that this kind of thinking can lead to violence (killing heretics who *might* lead souls astray, or dissidents who *might* obstruct an endless Marxist utopia) as well as virtue. A more mathematical concern is that in principle there may be many choices, each of which we expect to do an infinite amount of good (and perhaps also an infinite amount of harm) and that there is no obvious mathematical way to compare the competing infinities.

The comparison difficulties associated with infinite expectations can arise even when the payoffs themselves are finite with probability one (e.g., if the utility payout is a Cauchy random variable). This problem set illustrates this point with a particularly vexing form of a famous envelope switching paradox. Interestingly, in this paradox, the conditional expectations used for decision making are all finite; but a certain *a priori* expectation is infinite, and that is the root of the paradox. I hope that you enjoy thinking about the story, and that it causes you at most a finite amount of existential angst.

Please stop by my weekly office hours (2-249, Wednesday 3 to 5) for discussion.

A. FROM TEXTBOOK CHAPTER FIVE: Theoretical Exercise 26: If X is a beta random variable with parameters a and b show that

$$E[X] = \frac{a}{a+b},$$

$$\mathrm{Var}(X) = \frac{ab}{(a+b)^2(a+b+1)}.$$

B. Let p be the fraction of MIT students who love Ariana Grande — or, more precisely, the fraction who will *say* they love Ariana Grande when you ask (making it clear that you *absolutely* require a simple yes or no answer). Let's make believe that your initial Bayesian prior for p is uniform on $[0,1]$. Now ask three of your fellow students (actually do this!) one at a time whether they love Ariana Grande, and write the pair (# yes answers so far, # no answers so far) before you start and after each time you ask a question. For example, you will write the pairs

$$(0,0),(1,0),(2,0),(3,0)$$

if everyone you ask loves Ariana Grande. Pretend that you have chosen your people uniformly at random from the large MIT population, so that each answer is yes with probabilty p and no with probability $(1-p)$ independently of the other answers. Then write down each of the four number pairs, and beside each one draw a rough picture of the graph of the revised probability density function for p that you would have at that point in time, along with its algebraic expression, which should be a polynomial whose integral from 0 to 1 is 1. You can use graphing software if you want. Beside each graph write down the corresponding conditional expectation for p (using the results from part A) given what you know at that time.

Remark: Last year this question was asked about Taylor Swift. If you have a recommendation for next year's question, let me know. :) The ideal is somebody who is well known but whose popularity within MIT is something most of us would be *a priori* unsure of, so that the uniform prior doesn't seem *too* unreasonable. (I have seen enough opinion polls to suspect that very few national *politicians* are loved by more than 70 percent of a population...)

C. The following is one formulation of a famous "two envelope" paradox. Jill is a money-loving individual who, given two options, invariably chooses the one that gives her the most money in expectation. One day Harry, a trusted (and capable of delivering) individual, offers her the following deal as a gift. He will secretely toss a fair coin until the first time that it comes up tails. If there are n heads before the first tails, he will place 10^n dollars in one envelope and 10^{n+1} dollars in the second envelope. (Thus, the probability that one envelope has 10^n dollars and the other has 10^{n+1} dollars is 2^{-n-1} for $n \geq 0$.) Harry will then hand Jill the pair of envelopes (randomly ordered, indistinguishable from the outside) and invite her to choose one. After Jill chooses an envelope she will be allowed to open it. Once she does, she will be allowed to either keep the money in the first envelope or switch to the second envelope and keep whatever is in the second envelope. However, if she decides to switch, she has to pay a one dollar "switching fee."

1. If Jill finds 100 dollars in the first envelope she opens, what is the conditional probability that the other envelope contains 1000 dollars? What is the conditional probability that the other envelope contains 10 dollars?

2. If Jill finds 100 dollars in the first envelope she opens, how much money does Jill expect to win from the game if she does not switch envelopes? (Answer: 100 dollars.) How much does she expect to win (net, after the switching fee) if she *does* switch envelopes?

3. Generalize the answers above to the case that the first envelope contains 10^n dollars (for $n \geq 0$) instead of 100.

4. Jill concludes from the above that, no matter what she finds in the first envelope, she will expect to earn more money if she switches envelopes and pays the one dollar switching fee. This strikes Jill as a bit odd. If she knows she will always switch envelopes, why doesn't she just take the second envelope first and avoid the envelope switching fee? How can she be maximizing her expected wealth if she spends an unnecessary "switching fee" dollar no matter what? How does one resolve this apparent paradox? (Use the hints in the Lecture 24 slides as needed. Even with hints, it may take time to make peace with this problem.)

D. Alice and Bob are interested in having a child and, after difficulty conceiving, decide to undergo a medical procedure called IVF. In their universe, each couple has a random quantity p, uniformly distributed on $[0, 1]$, which indicates the probability that they will conceive a child after a cycle of IVF treatment. (The value p depends on permanent biological characteristics of Alice and Bob, but its value is unknown to them, so we model it as a random variable.) If Alice and Bob attempt multiple cycles, each one succeeds with the same probability p, independently of what happens on previous cycles.

(a) Explain intuitively why (in this universe) the probability that Alice and Bob conceive after one cycle should be .5 (i.e., the expected value of p).

(b) *Given* that Alice and Bob did not conceive during the first $(k-1)$ cycles, what is the updated Bayesian probability density for the random variable p?

(c) Use the answer in (b) to explicitly compute the expected value for p, given that the couple did not conceive during the first $(k-1)$ cycles. The answer is the conditional probability that the couple conceives during the kth cycle, given that they did not conceive during the first $(k-1)$ cycles. (One can prove in general that if one *first* chooses r in some random fashion and *then* tosses a coin that is heads with probability r, the overall probability of heads is the expectation $E[r]$.)

(d) Compute the conditional probability describe in (c) in a different way: imagine that $X_0, X_1, X_2, \ldots, X_k$ are uniformly and independently distributed on $[0, 1]$. Write $p = X_0$ and declare that the jth cycle succeeds if and only if $X_j < X_0$. Show that this model is equivalent to the one initially described, and then explain why the probability that $X_k < X_0$, *given* that X_0 is the smallest of the set $\{X_0, X_1, \ldots, X_{k-1}\}$, should be $1/(k+1)$. [Hint: use symmetry to argue that *a priori* the rank ordering of $X_0, X_1, \ldots, X_k$ is equally likely to be given by each of the $(k+1)!$ possible permutations.]

(e) Suppose that instead of being uniform the random variable p is *a priori* distributed on $[0, 1]$ according to the density function $f(x) = 2 - 2x$. (This might be more realistic, see remark below.) Under this assumption, compute the probability of success on the kth cycle given that the first $(k-1)$ cycles failed. [Hint: recognize $f(x)$ as itself a beta random variable and reduce to the previous case.]

Remark: This problem was inspired by a NY Times article called *With in vitro fertilization persistence pays off* (look it up) which reports on a large study:

> The rate of live births for participants after the first cycle in the new study was 29.5 percent, compared with 20.5 percent ater the fourth cycle, 17.4 percent after the sixth cycle, and 15.7 percent after the ninth cycle.

The numbers start a bit below our answer in (e) (since $.295 < 1/3$) and end up larger (since $.157 > 1/11$). This may suggest that p values are not distributed according the f that we guessed (somewhat arbitrarily) in (e). On the other hand, maybe different people have different Bayesian priors for p (based on age, known physical issues, etc.) and those whose p values are expected *a priori* to be small tend to discontinue IVF after fewer cycles; if so, this could explain the higher reported success rates for later cycles.

E. Suppose $X_1, X_2, \ldots, X_8$ are independent Cauchy random variables. Compute the probability that $X_1 + X_2 + X_3 > X_4 + X_5 + X_6 + X_7 + X_8 - 8$. (Hint: try combining the spinning flashlight story with left-right symmetry and the fact that the average of independent Cauchy random variables is itself a Cauchy random variable.)

F. On Planet A a site called rottentomatoes.com analyzes movie reviews. Each review is classified "fresh" if it seems on balance positive, "rotten" otherwise. Each movie has an *a priori* quality parameter $p \in [0, 1]$. After it is released, professional reviewers arrive one at a time and write reviews, each of which is fresh with probability p (independently of the others). The Tomatometer Score is the overall percentage of reviews that were fresh, expressed as a number between 0 and 100. One can show — using the *strong law of large numbers*, which will appear later in this course — that no matter what p is, the Tomatometer Score will (with probability one) converge to $100p$ in the limit as the number of reviews tends to infinity; so e.g. if $p = 3/5$ then the Tomatometer score will converge to 60 in the long run.

1. Suppose one movie has quality parameter .5 and another .6. Use normal approximations to estimate the probability the former gets a higher Tomatometer score than the latter after each movie has 143 reviews. (Hint: remember Harper and Heloise.)

2. Repeat the above with one movie having parameter .8 and the other .9, and with 100 reviews for each movie. (In both this problem and the previous one, the higher quality movie *probably* scores higher, but in neither case is it a sure thing.)

3. Imagine a studio makes a movie but has no idea in advance how well it will be received. The movie has a quality parameter p, but the studio does not know what it is and *a priori* considers p to be a *uniformly random variable* on $[0, 1]$. But then reviewers arrive one at a time to make reviews, each rating the movie fresh with probability p and rotten otherwise. Using beta random variables, give the *conditional* probability density for p given that one has seen f fresh and r rotten reviews so far.

4. Argue that if the studio does not know p, and knows only the number of f and r reviews seen so far, then it considers the probability of the *next* review being fresh to be $\frac{(f+1)}{(f+1)+(r+1)}$. (You can use the things derived in the IVF problem.) Using this compute the probability that the first four reviews are fresh, rotten, fresh, fresh in that order.

On Planet B, each released movie is initially given one fresh and one rotten review (to get the ball rolling). After that reviewers arrive one at a time to write and post reviews. But these reviewers do not form opinions independently; instead, each reviewer selects, uniformly at random, one of the *previously posted* reviews and writes a review of the same type (fresh or rotten). Let F_n be the fraction of the first n reviewers who rated a movie fresh. (We know

$F_2 = 1/2$, but F_3 could be 1/3 or 2/3, and F_4 could be 1/4, 2/4 or 3/4.) Ultimately an infinite number of reviewers arrives, and the Tomatometer score is the limit $\lim_{n\to\infty} 100F_n$.

5. What is the probability on this planet that the first four reviews (after the "get the ball rolling" two) are fresh, rotten, fresh, fresh in that order? Does your answer agree with the answer computed above for the same sequence on Planet A? Would this still be be true if we replaced "fresh, rotten, fresh, fresh" by *any* finite length sequence?

Remark: It seems oddly coincidental that each sequence has the same probability on Planet A as on Planet B, even though the mechanism for generating the sequence is *completely* different.

6. Use comparison to Planet A to argue that on Planet B the limiting Tomatometer score is a uniformly random variable on $[0, 100]$. (Google *Pólya's urn* for more on this model.)

Remark: On Planet A, you can imagine that a sufficiently skilled movie expert could figure out p after seeing an advance screening of the movie. This expert would then know *exactly* what the Tomatometer score would converge to in the $n \to \infty$ limit. But on Planet B, it is impossible to know anything at all about the limiting score just from seeing the movie.

Remark: Are the mechanisms of *our* world is closer to A (where reviewers see same movie but otherwise work independently) or B (where reviewers influence each other, and final consensus is unrelated to quality)? What explains why *Mona Lisa* and *Starry Night* are such iconic art works and *Baby Shark* has 2 billion views? I have no answer, but I include a story below.

Quoted Remark (from Cass R. Sunstein's book *On Rumors*): The Princeton sociologist Matthew Salganik and his coauthors 14 created an artificial music market among 14,341 participants who were visitors to a website that was popular among young people. The participants were given a list of previously unknown songs from unknown bands. They were asked to listen to selections of any of the songs that interested them, to decide which songs (if any) to download, and to assign a rating to the songs they chose. About half of the participants made their decisions based on the names of the bands and the songs and their own independent judgment about the quality of the music. This was the control group. The participants outside of the control group were randomly assigned to one of eight possible subgroups. Within these subgroups, participants could see how many times each song had been downloaded. Each of these subgroups evolved on its own; participants in any particular world could see only the downloads in their own subgroups.....

It turned out that people were dramatically influenced by the choices of their predecessors. In every one of the eight subgroups, people were far more likely to download songs that had been previously downloaded in significant numbers—and far less likely to download songs that had not been so popular. Most strikingly, the success of songs was highly unpredictable. The songs that did well or poorly in the control group, where people did not see other people's judgments, could perform very differently in the "social influence subgroups." In those worlds, most songs could become very popular or very unpopular, with everything depending on the choices of the first participants to download them.