

Poisson

18.600 Problem Set 5, due October 20

Welcome to your fifth 18.600 problem set! We'll be thinking more about Poisson random variables and the corresponding processes. Food for thought: should we replace the expression "Bad things happen in threes" with "When one bad thing is expected, the chance of 3 is $1/6e$ "?

A. FROM TEXTBOOK CHAPTER FOUR:

1. Theoretical Exercise 16: Let X be a Poisson random variable with parameter λ . Show that $P\{X = i\}$ increases monotonically and then decreases monotonically as i increases, reaching its maximum when i is the largest integer not exceeding λ . *Hint:* Consider $P\{X = i\}/P\{X = i - 1\}$.
2. Theoretical Exercise 25: Suppose that the number of events that occur in a specified time is a Poisson random variable with parameter λ . If each event is "counted" with probability p , independently of every other event, show that the number of events that are counted is a Poisson random variable with parameter λp . Also, give an intuitive argument as to why this should be so. As an application of the preceding result, suppose that the number of distinct uranium deposits in a given area is a Poisson random variable with parameter $\lambda = 10$. If, in a fixed period of time, each deposit is discovered independently with probability $\frac{1}{50}$, find the probability that (a) exactly 1, (b) at least 1, and (c) at most 1 deposit is discovered during that time.

B. ANSWER THE FOLLOWING:

1. Compute the expectation of $(-\log X)^n$, where n is a positive integer and X is a uniform random variable on the interval $[0, 1]$. Do this in two ways. First, write $Y = (-\log X)$, find the cumulative density function F_Y , work out f_Y and compute $E[Y^n]$. Second compute $\int_{-\infty}^{\infty} f_X(x)(-\log x)^n dx$ directly using something like wolframalpha. Do your answers agree?
2. Let X be uniform on $[0, 1]$ and compute the expectation and variance of X^n when n is a positive integer.

C. In Regular Bus City, there is a shuttle bus that goes between Stop A and Stop B, with no stops in between. The bus is perfectly punctual and arrives at Stop A at precise five minute intervals (6:00, 6:05, 6:10, 6:15, etc.) day and night, at which point it immediately picks up all passengers waiting. Citizens of Regular Bus City arrive at Stop A at Poisson random times, with an average of 5 passengers arriving every minute, and board the next bus that arrives.

1. Suppose that you visit this city and that you arrive at Stop A at a time chosen uniformly at random from the times in a day. How long do you expect to have to wait until the next bus?
2. How many citizens of Regular Bus City do you expect to be on the bus that you take?

In Poisson Bus City, there is a shuttle bus that goes between Stop A and Stop B, with no stops in between. The times at which the bus arrives at Stop A are a Poisson point process with one bus arriving every five minutes on average, day and night, at which point it immediately picks up all

passengers waiting. Citizens of Poisson Bus City (like those of Regular Bus City) arrive at Stop A at Poisson random times, with an average of 5 passengers arriving every minute, and board the next bus that arrives. [Hint: You can use the fact that in a Poisson point process with parameter λ , the waiting time until the first event is an exponential random variable with parameter λ , and the expectation of such a random variable is λ^{-1} .]

3. Suppose that you visit this city and that you arrive at Stop A at a time chosen uniformly at random from the times in a day. How long do you expect to have to wait until the next bus?
4. How many citizens of Poisson Bus City do you expect to be on the bus that you take?
5. Are the following two statements true or false? If they are both true, explain in words the apparent discrepancy:
 - (i) When you visit, buses in Poisson Bus City seem on average to come twice as slowly and to be twice as crowded as those in Regular Bus City
 - (ii) In both cities, buses come on average every five minutes and people come on average five times per minutes, so that over the long haul there are 25 people per bus on average—so buses are on average equally crowded in the two cities.

Remark: Poisson Bus City is not the worst case scenario. Suppose that buses come in pairs (one right behind the other) with the pairs arriving as a Poisson point process with one pair every 10 minutes on average. And suppose that whenever this happens, everybody gets in the first bus and leaves the second bus empty. Now if you arrive at a random time, you can expect your bus to take four times as long to come and be four times as crowded as in Regular Bus City (assuming that like others you get on the first bus in a pair). On a real life bus route with many stops, the closer a bus is to the bus ahead of it, the faster it can go (since it is picking up fewer passengers) which can lead to this kind of clumping.

D. Each day (independently of all other days) Jill has a $1/5000$ chance of hearing a particular fact: let's say the fact that Henry Mancini composed "The Pink Panther Theme." Jill stores something in long term memory after hearing it 3 times. Use Poisson approximations to (approximately) answer the following:

1. What is the probability that, by the time Jill is 10,000 days old, she knows that Henry Mancini wrote "The Pink Panther Theme"?

Alice reads more than Jill and has a better memory for trivia. Each day (independently of all others) Alice has a $1/1000$ chance of learning that Henry Mancini wrote "The Pink Panther Theme," and she stores information in long term memory after hearing it twice.

2. What is the probability that, by the time Alice is 10,000 days old, she knows that Henry Mancini wrote "The Pink Panther Theme"?

E. This problem addresses the Gompertz model for the duration of human life. But it starts out as another story about buses. Let X_1, X_2, X_3 be a Poisson point process of parameter 1 on $[0, \infty)$. Recall that this implies that X_1 and $X_2 - X_1$ and $X_3 - X_2$, etc., are i.i.d. exponential random variables each with parameter 1. Now for each integer $i \geq 1$, let $Y_i = \log X_i$. In Poisson Bus City, you might imagine that a bus line starts operating at time zero, and thereafter bus arrivals correspond to the times X_i .

On Accelerating Frequency Planet (AFP) the bus arrival times are Y_1, Y_2, \dots . That is, each arrival time is the *natural logarithm* of a point in the Poisson point process. Time is measured from $-\infty$ to ∞ on AFP, so it is possible that some bus arrival times are negative.

1. Show that, on AFP, given any constants $a < b$, the number of buses that arrive between times a and b is a Poisson random variable with parameter $\int_a^b e^x dx$.
2. Explain (with a sentence each) why the following things are true: the number of buses that arrive during the time interval $(-\infty, 0]$ is Poisson with parameter 1, while with probability one *infinitely* many buses arrive after time 0. If $\alpha = \ln 2 \approx .7$, then for each k the expected number of buses that arrive during $[k\alpha, (k+1)\alpha]$ is twice as large as the the expected number that arrive during $[(k-1)\alpha, k\alpha]$. (In other words, the doubling time for the bus arrival-frequency rate is α .) Moreover, the *first* bus's arrival time is a random variable whose median is $\ln(\ln(2)) \approx -.37$.

Note you can approximate an ordinary Poisson point process with parameter λ by partitioning time into disjoint intervals of the form $[t, t + \epsilon)$ for small ϵ and asserting that each interval independently contains a bus with probability $\lambda\epsilon$. Things are similar on AFP, except that the probability is approximately $e^t\epsilon$ when ϵ is small; in some sense, this is like saying that the Poisson parameter λ is “time dependent” (and exponentially increasing) with $\lambda(t) = e^t$.

More to the story: at time -7 on AFP, an adorable but immobile sloth is born at the bus stop, where it lives until it is killed by the first bus that arrives. Since it is unlikely the first bus comes before time -7 , (b) suggests that the sloth's life span is a random variable with median about $7 - .37 \approx 6.63$. The standard unit of time on AFP is the duodecennium (i.e., twelve years), so that α “units” means $12\alpha \approx 8.32$ years and the sloth's median life span is $6.63 * 12 \approx 80$ years.

3. Suppose that on AFP, half of the buses have fat tires and half have thin tires (bus type decided by independent coin toss for each bus), and female sloths are only killed by fat tired buses, while males are killed by all buses. Argue that if the sloth is female, its life expectancy is about 8.32 years (i.e., α duodecennia) longer than if it had been male. (Hint: use problem A.2 and argue that the probability density function for the lifespan of a female born at time -7 agrees with that of a male born at time $-7 - \alpha$. Then note that having a bus between time $-7 - \alpha$ and -7 is very unlikely.)
4. Let p_k be the probability that the sloth dies during its k th year of life, *given* that it has survived for $(k-1)$ years. Argue that p_k is approximately the expected number of buses that arrive during that k th year when p_k is small — say, less than .1. (This is related to arguing that the probability that a λ Poisson random variable equals 1 is approximately λ

when λ is reasonably small.) Thus p_k grows (roughly) exponentially in k for the first 80 or so years of life.

5. Look up <https://www.ssa.gov/oact/STATS/table4c6.html> and https://en.wikipedia.org/wiki/Gompertz%E2%80%93Makeham_law_of_mortality and (after looking them over for five minutes) write a sentence or two about what you noticed — and in particular about how closely the p_k corresponding to humans match those of the sloths on AFP. (Three obvious differences: humans are much more likely than AFP sloths to die during the first year or so of life. Humans in late teens and twenties — especially males — die at a rate that is higher than the Gompertz law would predict, hence the “bump” in the otherwise nearly straight line on the Wikipedia chart. This may be in part be due to risky behaviors, which are either more pronounced among people that age or simply make a larger difference on the log scale because other causes of death are low. Finally, outside the “bump” period the gap between death rates for male and female humans is large but not as large as for AFP sloths.) You may find it helpful to consult <https://gravityandlevity.wordpress.com/2009/07/08/your-body-wasnt-built-to-last-a-lesson-from-human-mortality-rates/> for a somewhat breezier account of the Gompertz law story. (Notice also that this problem has another part after the next few remarks.)

Remark: Gompertz law (i.e., exponential mortality rate growth) appears to apply pretty well to both humans and animals (with a species dependent doubling rate α). Google Gompertz mortality and find out more. If you wanted a story to explain the exponential growth, one naive one would be that “glitches” in the body accumulate exponentially, and your death rate is proportional to the number of glitches. Another simplistic story is that if there is a genetic mutation that *causes* an organism to die at age X , then the rate at which natural selection eliminates that gene decreases with the proportion of organisms who reach age X . So mutations that cause functioning to break down in old people accumulate faster than those that affect young people. Note also that although Gompertz law holds pretty well in developed nations, it does not hold in settings where a large fraction of deaths are caused by predators, wars, infectious diseases, etc. that are as deadly to the young as the old.

Remark: The assumption of an exponential increase in “bus arrival frequency” suggests that an unhealthy habit that *doubles* your probability of dying within any given year should *subtract* about α units of life expectancy (one doubling period). Medical advances that eliminate half the number of fatal buses should *add* about α units of life expectancy. If medicine eliminated 7/8 of the buses, this should increase life expectancy by 3α units, about 25 years. Lifestyle choices that decrease the death rate by 29 percent (i.e., which multiply death rate during any given year by $1/\sqrt{2}$) should add about $\alpha/2$ units (about four years) to life expectancy (since your death rate each year is what it would be if you were 4 years younger).

Remark: *Life expectancy* is computed by estimating what fraction of people of each numerical age die each year, and then calculating how long newborns would expect to live if they died with the corresponding probability each year. Per discussion above, doubling the death rate each year should *roughly* correspond to subtracting eight years of life expectancy. Decreasing the death rate by 29 percent *roughly* corresponds to adding four years of life expectancy. Pre-pandemic US life

expectancy per https://en.wikipedia.org/wiki/List_of_countries_by_life_expectancy was about 79 years. Countries with life expectancy 8 years *lower* included Guatemala, Bangladesh, and Ukraine. Countries with life expectancy 4 years *higher* included Japan, Switzerland, and Singapore. Factors presumably include opioids, guns, cigarettes, obesity, alcoholism, pollution, diet, traffic, etc. Covid was tough on the US, over extra 1.3 million excess deaths https://www.cdc.gov/nchs/nvss/vsrr/covid19/excess_deaths.htm#dashboard and US life expectancy fell 2.5 years from 2019 to 2021 before rebounding a year in 2022. <https://www.medrxiv.org/content/medrxiv/early/2023/03/01/2023.02.26.23286363.full.pdf>.

Remark: In a world where Gompertz applies precisely with a doubling period of 8 years (and people's life spans are independent of each other), if you have $8 = 2^3$ students of age 20 and one professor of age $44 = 20 + 3 \times 8$, then the time until the first student dies should agree in law with the time until the professor dies. Essentially, being three doubling periods older means you expect eight times as many buses coming your way. But being eight people instead of one means that (as a group) you also have eight times as many buses coming your way. (If you want to extend the bus metaphor, imagine a different independent lane of buses for each person, with frequency rates depending on that person's age...) If you have a class of 128 students of age 20 and one older professor of age 76, then the time until the first student dies agrees in law with the time until the professor dies. If two parents are 24 years older than a child then, of the buses coming towards these three people, $1/17$ are coming toward the child and $8/17$ are coming toward each parent (since each parent has 8 times as many buses coming its way as the child). In fact, one can show that *given* the arrival times of the buses, we can assign each bus independently to one of the three people, with probabilities $8/17$ for each parent and $1/17$ for child. Looking at the first bus, we find that the probability that the child is the first of the three people to die is $1/17$.

6. Suppose that, in the world of the previous remark, there are four siblings of ages 20, 20, 28, and 36. What is the probability the oldest one dies first?

Remark: Per CDC, smoking roughly triples mortality https://www.cdc.gov/tobacco/data_statistics/fact_sheets/health_effects/tobacco_related_mortality/index.htm and reduces life expectancy by over a decade. If this were true exactly (on our purely Gompertz 8-year-doubling world) then a smoker's death rate would be the same as that of a non-smoker who was $8 \cdot \log_2(3) \approx 12.67$ years older. If there were two young people of the same age (one smoker, one not) the non-smoker could say "You should not smoke. You are shortening your life by 12 years," and the smoker could counter "Yes, but don't be too sanguine about your own life span. There is still a $1/4$ chance that I will outlive you." The point is that for both individuals the uncertainty is quite high. Thinking back to our original model, if X is a rate one exponential and $Y = 12(7 + \log X)$, then $SD(Y) = 12SD(\log X)$ which comes out to about 15 on wolframalpha. This is consistent with <https://www.nber.org/papers/w14093>, which states that the standard deviation of adult lifespan in the US is about 15 years. As the linked paper notes, this uncertainty makes it hard to plan for things like retirement and inheritance. C'est la vie.

Remark: To *drastically* increase human life span (so we live to 150, say) we'll have to change α — i.e., to slow the exponential growth of the mortality rate. How can we do this? Can we freeze blood and tissue from when we're younger and reintroduce it later? Can we reset age markers? Can we manually edit the genes that cause aging? Somebody in this class should figure this out. We have a lot of buses coming our way. (Google *naked mole rat Gompertz* for inspiration.)

F. US elections are complicated. Let's tell a simpler story using uniform random variables.

1. Suppose f is an affine function on an interval $[a, b]$ and X is a uniform random variable on $[a, b]$. Show that $E[f(X)] = f[(a + b)/2]$.
2. Suppose f is a function on $[0, n]$ whose restriction to $[k, k + 1]$ is an affine function when $k \in \{0, 1, \dots, n - 1\}$. (This f is a “piecewise linear” function.) Suppose X is uniformly distributed on $[0, n]$. Show that $E[f(X)] = \frac{1}{n} \left(f(1 - \frac{1}{2}) + f(2 - \frac{1}{2}) + \dots + f(n - \frac{1}{2}) \right)$.

In Simple Electoral Country there are 51 states (all of exactly the same size) and only two parties: Party A and Party B. Let α be a uniformly random variable in $[-3, 3]$ that indicates “how many more percentage points popular than typical” Party A is during a certain year. In the 51 states (ordered from least to most A-friendly) the percentages voting for Party A are $25 + \alpha$, $26 + \alpha$, $27 + \alpha$, \dots , $75 + \alpha$. Only the middle five states (those with vote percentages $48 + \alpha$, $49 + \alpha$, $50 + \alpha$, $51 + \alpha$, $52 + \alpha$) are “swing states” in the sense that the state election winner depends on α . The overall election winner is the party that wins 26 or more states.

In some ways, elections in Simple Electoral Country are even closer than α would seem to suggest. For example, if $\alpha = .5$ then A wins 26 states, but in one of those A wins only 50.5 percent of the vote. Swapping the votes of .5 percent in *just that single state* would be enough to tie up the race (and going infinitesimally further would change the outcome). The number of votes one “needs to flip” to change the outcome is a mere $.005/51 = 1/(10200)$ fraction of the total number of votes in the country. Party A wins the election by 1 percentage point (50.5 to 49.5 percent) but one can change the election outcome by flipping less than .01 percent of the total votes.

3. Show that if $\alpha \in [0, 3]$ then the minimal number of percentage points of the overall vote one would need to flip (change from A to B) in order to tie up the election is given by

$$f(\alpha) = \begin{cases} \alpha/51 & \alpha \in [0, 1] \\ 1/51 + 2(\alpha - 1)/51 & \alpha \in (1, 2] \\ 3/51 + 3(\alpha - 2)/51 & \alpha \in (2, 3] \end{cases}$$

4. Conclude that if α is chosen uniformly from $[-3, 3]$ then the *expected minimal number of percentage points of votes one has to flip to change the outcome* is $\frac{1}{3}(\frac{1}{2} + \frac{4}{2} + \frac{9}{2})/51 = 7/153$. This comes to about .046 percent (which is *much* less than 3 percent).

Remark: If each state has 3 million voters, then the expected margin of victory is 3 percent (which comes to $.03 * 153,000,000 = 4,590,000$ votes). However the *expected number of votes needed to flip up the outcome* is a mere 70,000 (plus a couple of tiebreaker votes).

Remark: In reality, if you plot vote percentages of individual states along a line, they look more like the bus arrival times of Poisson Bus City than those of Regular Bus City (i.e., they are not evenly spaced). Can you adapt the story above to a model with some kind of “Poisson spacing”?

Remark: Even with Poisson spacing, the US system is more complex than the above because states have variable sizes and electoral weights. Still, US presidential margins—like those in the model—tend to be *pretty small* (no presidential candidate has won more than 53 percent of the popular vote since the 1980's) while the number of vote flips needed to change the outcome is often *very small* (tens of thousands rather than millions).