Chris Piech
CS109

Winter 2025
Feb 11th

CS109 Midterm Exam

This is a closed calculator/computer/phone/smart-watch/smart-toothbrush exam. You are, however, allowed to use notes in the exam. You have 2 hours (120 minutes) to take the exam. The exam is 120 points, meant to roughly correspond to one point per minute of the exam. You may want to use the point allocation for each problem as an indicator for pacing yourself on the exam.

In the event of an incorrect answer, any explanation you provide of how you obtained your answer can potentially allow us to give you partial credit for a problem. For example, describe the distributions and parameter values you used, where appropriate. It is fine for your answers to include summations, products, factorials, exponentials, and combinations. You can leave your answer in terms of Φ (the CDF of the standard normal) or Φ^{-1}. For example $\Phi(3/4)$ is an acceptable final answer.

1891

I acknowledge and accept the letter and spirit of the honor code. I pledge to write more neatly than I have in my entire life:

Signature: ____

Family Name (print): ____

Given Name (print): ____

Stanford Email (@stanford.edu): ________________________________

1. What a Time to Be Alive [26 points]

```
def main():
  mystery = poisson_sample(lambda = 2.5)
  if mystery < 2:
     return good_times()
  else:
     return what_a_time_to_be_alive()

def good_times():
  treasure = bernoulli(0.2)
  kindness = bernoulli(0.7)
  health = bernoulli(0.9)
  return (treasure == 1 or kindness == 1) and health == 1

def what_a_time_to_be_alive():
  k = 0
  for i in range(10000):            # repeat 10,000 times
    k += bernoulli(0.4)
  return k < 3998

def bernoulli(p):
  # returns 1 with probability p
  # returns 0 with probability (1-p)
  if random() < p: return 1
  else: return 0

def poisson_sample(lambda):
  # return a value sampled from the poisson distribution
  return scipy.stats.poisson.rvs(lambda)
```

a. (6 points) You call `good_times()`. What is the probability it returns `True`?

$$\begin{aligned}
P(\text{good\_times returns True}) &= P([\text{treasure} = 1 \text{ OR kindness} = 1] \text{ AND health} = 1) \\
&= P(\text{treasure} = 1 \text{ OR kindness} = 1) \cdot P(\text{health} = 1) \\
&= [P(\text{treasure} = 1) + P(\text{kindness} = 1) - P(\text{treasure} = 1 \text{ AND kindness} = 1)] \cdot P(\text{health} = 1) \\
&= [0.2 + 0.7 - 0.2 \cdot 0.7] \cdot 0.9
\end{aligned}$$

We break down the events above first by recognizing that the events for each Bernoulli are independent of each other (which lets us re-write the ”and” as a product), and then by applying Inclusion-Exclusion to break down probability of ”or”.

We can also use De Morgan’s Laws to convert the ”or” into an ”and”, with complements:

$$\begin{aligned}
P(\text{good\_times returns True}) &= P(\text{treasure} = 1 \text{ OR kindness} = 1) \cdot P(\text{health} = 1) \\
&= [1 - P(\text{treasure} = 0 \text{ AND kindness} = 0)] \cdot P(\text{health} = 1) \\
&= [1 - (1 - 0.2)(1 - 0.7)] \cdot 0.9
\end{aligned}$$

b. (8 points) You call `what_a_time_to_be_alive()`. What is the probability it returns `True`? Use an approximation so that your answer does not have a sum in it.

This function is calculating a Binomial sample, as the Binomial is a sum of Bernoullis. $k \sim \text{Bin}(n = 10000, p = 0.4)$. We want to find $P(k < 3998)$.

Since n is large, we need to approximate. Let X be the approximating RV for k. Since p is moderate, we should use a Normal distribution, setting the mean and variance equal to that of the original Binomial: $X \sim N(\mu = np = 4000, \sigma^2 = np(1-p) = 2400)$. Then, to approximate $P(k < 3998)$, we need to continuity correct:

$$P(k < 3998) \approx P(X < 3997.5) = \Phi\left(\frac{3997.5 - 4000}{\sqrt{2400}}\right)$$

c. (6 points) You call `main()`. What is the probability that `mystery` < 2?

From the code, mystery $\sim \text{Poi}(\lambda = 2.5)$. We can solve for $P(\text{mystery} < 2)$ using the Poisson PMF:

$$\begin{aligned} P(\text{mystery} < 2) &= P(\text{mystery} = 0) + P(\text{mystery} = 1) \\ &= \frac{2.5^0 e^{-2.5}}{0!} + \frac{2.5^1 e^{-2.5}}{1!} \\ &= e^{-2.5} + 2.5e^{-2.5} = 3.5e^{-2.5} \end{aligned}$$

d. (6 points) You call `main()`. What is the probability that `main` returns `True`? Let p_a, p_b, p_c be your answers to parts (a), (b), and (c) respectively.

Main() could return True if one of two mutually exclusive outcomes occurs: either mystery ¡ 2 and good_times returns True, or mystery $\geq$ 2 and what_a_time_to_be_alive() returns true. Thus, we can use the Law of Total Probability:

$$\begin{aligned} &P(\text{main() returns True}) \\ &= P(\text{mystery} < 2 \text{ AND good\_times() returns True}) + P(\text{mystery} \geq 2 \text{ AND what\_a\_time() returns True}) \\ &= p_c \cdot p_a + (1 - p_c)p_b \end{aligned}$$

2. Vibrant Variables [21 points]

For each of the following scenarios:
Step 1: Declare a single random variable of a type introduced in class that best models the scenario and specify its parameters e.g., $X \sim \text{Bin}(n = 10, p = 0.3)$.
Step 2: State the question in terms of your random variable e.g., $P(X = 2)$.
Step 3: Solve for the probability statement from Step 2.

a. (7 points) A doctor on call receives an average of 6 calls a day. These calls come in independently at all hours. What is the probability the doctor receives more than 8 calls today?

Let the random variable X represent the number of calls the doctor receives in a day. We are given an average rate of calls per day, which leads us to using a Poisson as follows:

$$X \sim \text{Poi}(\lambda = 6)$$

$$P(X > 8) = \sum_{i=9}^{\infty} \frac{6^i \cdot e^6}{i!}$$

It is equivalent to solve for the complement:

$$1 - P(X \leq 8) = \sum_{i=0}^{8} \frac{6^i \cdot e^6}{i!}$$

b. (7 points) A bioluminescent firefly emits a glow at an average rate of once every minute. A researcher is watching a single firefly. What is the probability that a researcher waits more than 5 minutes before the first glow?

Solution 1:
Let the random variable X represent the time until one glow. We can analyze this question as the probability that the time until we see one firefly glow is greater than 5 minutes, i.e., an Exponential distribution. Since the firefly glows once per minute, i.e. the time until one glow is 1 minute on average, $\lambda = 1$. From there, we have:

$$X \sim \text{Exp}(\lambda = 1)$$
$$P(X > 5) = 1 - (1 - e^{-1 \cdot 5}) = e^{-5}$$

Solution 2:
We can also solve this problem using a Poisson similar to section 3 as follows:

$$X \sim \text{Poi}(\lambda = 5)$$

$$P(X = 0) = \frac{5^0 \cdot e^{-5}}{0!} = e^{-5}$$

Here, λ is in units of "glow emissions per 5 minutes". So, in order for a researcher to wait more than 5 minutes before the first glow, there must be no glows in the first 5 minute interval as calculated above.

c. (7 points) Earthquakes are equally likely to occur at any point in a 24-hour day. Given an earthquake occurs, what is the probability that it happens in the 16-hour span between 6 am to 10 pm?

Let the random variable X represent the time when the given earthquake occurs. We can represent this time with a Uniform since earthquakes are equally likely to occur at any point in the day.

$$\begin{aligned} X &\sim \text{Uni}(0, 24) \\ P(6 \leq X \leq 22) &= P(X \leq 22) - P(X \leq 6) \\ &= \frac{22-0}{24-0} - \frac{6-0}{24-0} \\ &= \frac{22-6}{24-0} = \frac{2}{3} \end{aligned}$$

It's also possible to define the Uniform with a different range end, such as $\text{Uni}(0, 1)$, if the event of interest is also transformed to match this range, such as $P(6/24 < X < 22/24)$.

3. Board Game Analysis [15 points]

You are playing a board game where each turn you roll two fair six-sided dice.
If the two dice sum to 5, you get two cards.
If the two dice sum to 11, you get one card.
Otherwise, you get zero cards.

The grid on the right shows the sum of two dice for all 36 outcomes.

Value on first dice roll

Value on second roll

	1	**2**	**3**	**4**	**5**	**6**
1	2	3	4	5	6	7
2	3	4	5	6	7	8
3	4	5	6	7	8	9
4	5	6	7	8	9	10
5	6	7	8	9	10	11
6	7	8	9	10	11	12

Sum of the two dice

a. (5 points) What is the probability of getting at least one card in a turn?

The probability of obtaining at least one card in a turn is the sum of the probabilities that the dice sum to 5 and that the dice sum to 11, because the outcomes are mutually exclusive. From the table, there are 4 ways for the dice to sum to 5 and 2 ways for the dice to sum to 11. Since there are a total of 36 possible outcomes when rolling two dice, the probability is computed as follows:

$$\frac{4}{36}+\frac{2}{36}=\frac{6}{36}=\frac{1}{6}.$$

b. (5 points) What is the expectation of the number of cards received in a turn?

Let X be the number of cards received in a turn. The expectation $E[X]$ is given by the weighted sum of the possible outcomes, where each outcome is multiplied by its probability:

$$E[X]=\sum_{k=0}^{2} kP(X=k).$$

From the problem setup and the same approach at in part a, we know:

- $P(X=0)=\frac{30}{36}$.
- $P(X=1)=\frac{2}{36}$.
- $P(X=2)=\frac{4}{36}$.

We compute:

$$E[X]=0\cdot P(X=0)+1\cdot P(X=1)+2\cdot P(X=2).$$

Substituting the given probabilities:

$$E[X]=1\cdot\frac{2}{36}+2\cdot\frac{4}{36}.$$

The expected number of cards received in a turn is $\frac{5}{18}$.

c. (5 points) What is the variance of the number of cards received in a turn?

The variance is given by $\text{Var}(X)=E[X^2]-E[X]^2$

To find the second term, we just need to square the expectation from part b:

$$E[X]^2 = \left(\frac{5}{18}\right)^2$$

To find the first term, we apply the Law of the Unconscious Statistician:

$$E[X^2] = 2^2 \cdot P(X = 2) + 1^2 \cdot P(X = 1)$$

Plugging in:

$$\mathrm{Var}(X) = 4 \cdot \frac{4}{36} + 1 \cdot \frac{2}{36} - \left(\frac{5}{18}\right)^2 = \frac{137}{324}$$

4. Rival Production [19 Points]

A rival is producing items. We would like to estimate the number of items, N, that they have produced. We notice that each item has a unique serial number and we assume that when we acquire (sample) items each serial number on the item is a positive integer equally likely to be any number from the set $\{1, 2, \ldots, N\}$.

For example, if you randomly acquired (sampled) 5 items produced at the factory, you might see the serial numbers $\{7, 12, 14, 24, 33\}$ which should give you a clue as to what N could be!

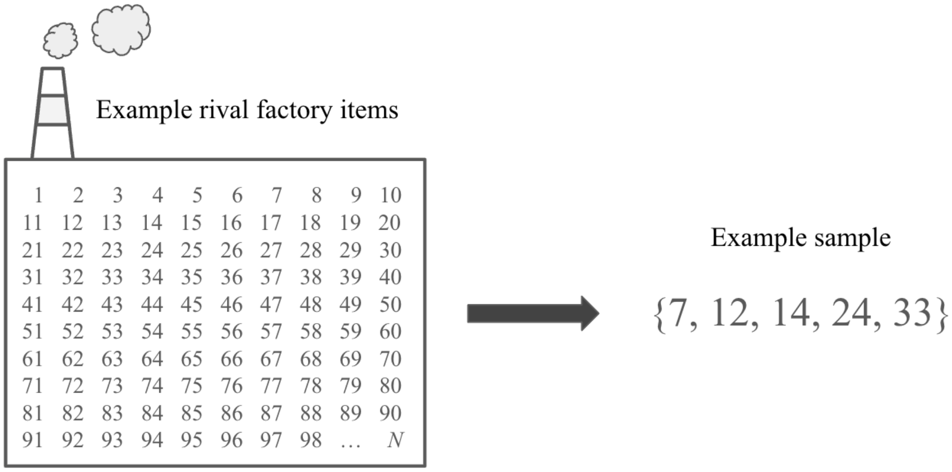

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	…	N

a. (7 points) For part (a) only, assume $N = 100$. We sample 5 items. What is the probability that the largest serial number in our sample is 33?

In this problem, we are sampling without replacement, since each unique combination of five distinct integers between 1 and N is a potential outcome. Since all items are equally likely to be sampled, each set of five samples is equally likely. We can thus compute probabilities using the event space and sample space method.

If we consider the randomly sampled items as unordered, we have $|S| = \binom{100}{5}$, i.e. the number of unique ways to choose 5 objects from 100 objects. To count the event space (number of samples with 33 as the highest), use the following generative story:
(1) Pick a set of four distinct numbers out of the range from 1 to 32,
(2) Pick 33 as the fifth number.
Thus, $|E| = \binom{32}{4} \cdot 1$, and $P(\text{largest serial number is } 33 | N = 100) = \frac{\binom{32}{4}}{\binom{100}{5}}$.

You can similarly consider the randomly sampled items as ordered. $|S| = \frac{100!}{95!} = 100 \cdot 99 \cdot 98 \cdot 97 \cdot 96$, and $|E| = 5(1 \cdot 32 \cdot 31 \cdot 30 \cdot 29)$. (Note that we have to multiply by 5 to account for the fact that we could choose to sample 33 in any of the possible positions in our order.)

b. (10 points) Your prior belief is that every value of N between 33 and 100 (inclusive) is equally likely. What is your updated probability mass function for N, given that you sampled 5 items and the largest serial number was 33?

This is a Bayesian inference problem. We are given the random variable N, the (unknown) total number of values or objects. Let L be the largest value in a sample of 5. We observe that $L = 33$. Setting up Bayes' Theorem and using Law of Total Probability in the denominator:

$$\begin{aligned} P(N = n|L = 33) &= \frac{P(L = 33|N = n) \cdot P(N = n)}{P(L = 33)} \\ &= \frac{L = 33|N = n) \cdot P(N = n)}{\sum_{i=33}^{100} P(L = 33|N = i) \cdot P(N = i)} \end{aligned}$$

Since there are $100 - 33 + 1 = 68$ values in the range $[33, 100]$, and we are told each of those possibilities is equally likely, we can use the prior

$$P(N = n) = \frac{1}{68} \text{ for } 33 \leq n \leq 100.$$

From part (a), we obtain the likelihood

$$P(L = 33|N = n) = \frac{\binom{32}{4}}{\binom{n}{5}} \text{ for } n \geq 33.$$

(since in part a, we assumed $N = 100$). Plugging these in:

$$P(N = n|L = 33) = \frac{\frac{\binom{32}{4}}{\binom{n}{5}} \cdot \frac{1}{68}}{\sum_{i=33}^{100} \frac{\binom{32}{4}}{\binom{i}{5}} \cdot \frac{1}{68}}$$

We only sum from $i = 33$ to 100 since the prior is 0 outside of this range.

c. (3 points) Given that you sampled 5 items and the largest serial number was 33, what is the probability that $N < 50$?

Let $P(N = n|L = 33)$ be the PMF computed in part (b). Using the same prior from part (b), the outcomes that satisfy this event are those where N is between 33 and 49, inclusive, so we sum over the PMF for this range of values of N:

$$\begin{aligned} P(N < 50|L = 33) &= \sum_{n=33}^{49} P(N = n|L = 33) \\ &= \sum_{n=33}^{49} \frac{\frac{\binom{32}{4}}{\binom{n}{5}} \cdot \frac{1}{68}}{\sum_{i=33}^{100} \frac{\binom{32}{4}}{\binom{i}{5}} \cdot \frac{1}{68}} \end{aligned}$$

5. Accidental Caps-Lock Press [20 Points]

Did you know that on many computers the caps-lock key needs to be pressed longer than other keys for it to activate? For this problem, assume these distributions over how long a user presses the caps-lock key:

The caps-lock key press time in seconds, given the press was an **accident**, is distributed as an Exp($\lambda = 1$).
The caps-lock key press time in seconds, given the press was **intentional**, is distributed as a $N(\mu = 2, \sigma^2 = 4)$.
When someone presses the caps-lock key, the probability it was by accident is 0.2.

a. (6 points) The user **intentionally** hits the caps-lock. What is the probability that the key press is longer than half a second?

Let T be a random variable representing the key press time. Let A represent the event that a key press was an accident, and let A^C represent the event that a key press was intentional. When a user presses a key intentionally, the key press time is exponentially distributed $T \mid A^C \sim N(\mu = 2, \sigma^2 = 4)$.

Since we have observed a user press a key intentionally, we will solve for $P(T > 0.5|A^C)$.

$$\begin{aligned}
P(T > 0.5 \mid A^C) &= 1 - P(T < 0.5) \\
&= 1 - \Phi\left(\frac{0.5 - 2}{\sqrt{4}}\right) \\
&= 1 - \Phi\left(\frac{-1.5}{2}\right) \\
&= 1 - \Phi(-0.75) = \Phi(0.75)
\end{aligned}$$

b. (10 points) A user hits a key for precisely 0.25 seconds. What is the probability that it was an accident?

We have observed that a user's key press time was precisely 0.25 seconds, and we are interested in the probability that the key press was accidental. So, we will find $P(A|T = 0.25)$ using Bayes Theorem.

$$\begin{aligned}
P(A|T = 0.25) &= \frac{P(T = 0.25|A) \cdot P(A)}{P(T = 0.25)} \\
&= \frac{P(T = 0.25|A) \cdot P(A)}{P(T = 0.25|A) \cdot P(A) + P(T = 0.25|A^C) \cdot P(A^C)}
\end{aligned}$$

Now, we can substitute $P(A) = 0.2$ for the prior, and $P(A^C) = 1 - 0.2 = 0.8$ in the denominator.

$$P(A|T = 0.25) = \frac{P(T = 0.25|A) \cdot 0.2}{P(T = 0.25|A) \cdot 0.2 + P(T = 0.25|A^C) \cdot 0.8}$$

Finally, we can use the PDFs of the Normal and Exponential distributions.

$$P(A|T = 0.25) = \frac{1 \cdot e^{-(1)(0.25)} \cdot 0.2}{1 \cdot e^{-(1)(0.25)} \cdot 0.2 + \frac{1}{2\sqrt{2\pi}} e^{\frac{-(0.25-2)^2}{2(2)^2}} \cdot 0.8}$$

c. (4 points) You want to choose a decision boundary time t. If the user presses the caps-lock key for fewer than t seconds, you won't register it. You would like to only register the press if you are over 95% sure that it was not an accident. Explain, as if teaching, how you could solve for t. Two sentences max. Equations are encouraged!

We would like to only register the key press if the posterior probability of a non-accidental key press is at least 95%, i.e.:

$$P(A^C \mid T = t) > 0.95.$$

Using similar logic to Part (b), we can use Bayes's Rule to expand this to

$$P(A^C \mid T = t) = \frac{f(T = t \mid A^C)P(A^C)}{f(T = t)} = \frac{f(T = t \mid A^C)P(A^C)}{f(T = t \mid A^C)P(A^C) + f(T = t \mid A)P(A)}.$$

From here, we can find the appropriate cutoff value of t by setting this expression equal to 0.95, substituting for all of the probabilities exactly as in Part (b), and solving for t.

We also accepted answers that explained how to compute the cutoff value of t via the equation $P(A \mid T = t) = 0.05$.

Additional Notes:

- Substituting gives the equations

$$\frac{\frac{1}{2\sqrt{2\pi}} e^{-\frac{(t-2)^2}{2 \cdot 2^2}} \cdot 0.8}{\frac{1}{2\sqrt{2\pi}} e^{-\frac{(t-2)^2}{2 \cdot 2^2}} \cdot 0.8 + e^{-t} \cdot 0.2} = 0.95$$

or

$$\frac{e^{-t} \cdot 0.2}{e^{-t} \cdot 0.2 + \frac{1}{2\sqrt{2\pi}} e^{-\frac{(t-2)^2}{2 \cdot 2^2}} \cdot 0.8} = 0.05.$$

- Many students used a similar approach as above but wrote that we should solve for the value of t that creates equality in the two equivalent inequalities $P(A^C \mid T \geq t) > 0.95$ or $P(A \mid T \geq t) < 0.05$. Technically, though, the information $T = t$ and the information $T \geq t$ are distinct events and give distinct posterior distributions for the event A.

6. The Large Language Model That Keeps on Trying [18 Points]

A large language model (LLM) is trying to solve a coding task. We observe that 16% of the time, it correctly solves the task on the first attempt. But how does that probability change if it gets to re-attempt the same problem (after being told its previous solutions didn't work)?

a. (6 points) We ask the LLM to attempt to solve the coding task 10 times. If each attempt has an independent 16% chance of solving the coding task, what is the probability that the LLM does **not** solve the task in ≤ 10 tries?

This problem has multiple valid answers.

Approach 1: Using core probability, we can recognize that the probability of the LLM not solving the task for 10 tries is the product of a fail on each try, multiplied 10 times:

$$P(\text{does not solve task in 10 tries}) = (1-0.16)^{10} = 0.84^{10}$$

You could frame this as a Binomial and use the PMF to get $\binom{10}{0}0.16^0 \cdot 0.84^{10}$, which simplifies to the same numeric answer.

Approach 2: Let A be the number of attempts until a success. We can define $A \sim Geo(0.16)$.

$$\begin{aligned}
P(A > 10) &= 1 - P(A \leq 10) \\
&= 1 - \sum_{i=1}^{10}(1-0.16)^{i-1}0.16 \\
&= 1 - \sum_{i=1}^{10}(0.84)^{i-1}0.16
\end{aligned}$$

b. (6 points) Instead of our assumption in part (a), assume the probability that an LLM correctly solves the task on attempt k, given that the LLM failed on the previous $k-1$ attempts, is $1/(10 \cdot k)$. What is the probability it does **not** solve the task in ≤ 10 tries?

Let T be the number of tries that the LLM takes to solve the task. We are given that $P(T = k \mid T \neq k-1) = \frac{1}{10k}$ and $P(T = 1) = 0.16$. We want to find $P(T > 10)$. Note that this distribution is discrete as T can only take on integer values.

$$\begin{aligned}
P(T > 10) &= P(T \neq 10 \mid T > 9)P(T > 9) \\
&= (1 - P(T = 10 \mid T > 9))P(T > 9) \\
&= (1 - P(T = 10 \mid T \neq 9))P(T \neq 9 \mid T > 8)P(T > 8) \\
&= \ldots \\
&= (1-0.16) \cdot \Pi_{k=2}^{10}(1 - P(T = k \mid T \neq k-1)) \\
&= (1-0.16) \cdot \Pi_{k=2}^{10}(1 - \frac{1}{10k})
\end{aligned}$$

We also accepted $\Pi_{k=1}^{10}(1 - \frac{1}{10k})$, since it was ambiguous whether the 0.16 information from the preamble applied to this part of the problem.

c. (5 points) Instead of our assumption in part (b) we would like our probabilities to match empirical observations. A team at Stanford has observed that the probability that an LLM can solve a particular problem within k attempts is:

$$P(T \leq k) = e^{(-k^{-1/10})}$$

where T is the number of attempts needed to solve the problem. Derive an expression for the probability of success on trial k (given $k-1$ fails). Your expression should be in terms of just k and e. You do not need to simplify.

Hint: Use the fact that the probability of "no successes by trial k" is equal to the probability of "no successes by trial $(k-1)$" and "no success on trial k (given $k-1$ fails)".

There are two ways of exploring this problem. The first assumes the LLM stops after 1 success. The other assumes it can succeed multiple times (up to n), so we must condition on the face that there are $k-1$ failures before the success on the kth attempt.

Approach 1:

$$\begin{aligned} P(T = k) &= P(T \leq k) - P(T \leq k-1) \\ &= e^{-k^{-1/10}} - e^{-(k-1)^{-1/10}} \end{aligned}$$

Approach 2:

$$\begin{aligned} P(T = k | T > k-1) &= 1 - P(T \neq k | T > k-1) \\ &= 1 - \frac{P(T \neq k, T > k-1)}{P(T > k-1)} \\ &= 1 - \frac{1 - P(T \leq k)}{1 - P(T \leq k-1)} \\ &= 1 - \frac{1 - e^{-k^{-1/10}}}{1 - e^{-(k-1)^{-1/10}}} \end{aligned}$$

That's all folks! We hope you had fun. Here are some optional notes for further curiosity.

i. **Board Game Analysis** is based off a popular game called Settlers of Catan.

ii. **Rival Production**: During World War 2, the Allies needed to know how many tanks Nazi Germany was producing. First they sent spies to Germany who estimated Germany produced 1,400 tanks per month. Separately, they noticed that the serial numbers on gear boxes on German tanks were unique and sequential. Using this observation, and a sample of gear box serial codes, the mathematicians used the math you derived to estimate the amount of tanks produced by Nazi Germany. They estimated production was 270 tanks per month. After the war, German records confirmed an actual production rate of 276 tanks per month—the probabilistic method was incredibly accurate!

iii. **Accidental Caps-Lock Press**: On a Mac, you need to press and hold the Caps Lock key for about half a second for it to activate. This delay is intentional to prevent accidental activation. This is not the case for other keys!

iv. **The Large Language Model That Keeps on Trying**: Large Language Monkeys was a very popular paper written by Stanford professor Azalia Mirhoseini in 2024. In the paper Azalia and team showed the empirical result that $P(T \leq k) = e^{(-k^{-1/10})}$. While writing this midterm, the teaching team came up with a proof that the simple equation used in part (b) is the correct asymptotic approximation of the result from part (c). Neat!