Definition 1. A random variable X has a *discrete uniform distribution* if it is equally likely to assume any one of a finite set of possible values.

Examples. Roll a single die. Choose a number in a lottery.

Definition 2. A random variable X has a *Bernoulli distribution* with parameter θ (with $0 < \theta < 1$) if its probability distribution is $f(x;\theta) = \begin{cases} 1-\theta & \text{if } x = 0 \\ \theta & \text{if } x = 1 \end{cases}$. The outcome 1 is often referred to as "success" while 0 is "failure" and the random variable is often called a Bernoulli trial.

Examples. Flip a coin and count the number of heads. Ask one person a yes or no question.

1. Express the mean and variance of a Bernoulli random variable as functions of θ.

Definition 3. The total number of successes in n independent Bernoulli trials is a random variable with a *Binomial distribution.* More precisely, let $X_1, X_2, \ldots, X_n$ be n independent, identically distributed (iid) Bernoulli random variables, all with probability of success θ. The total number of successes is $X = X_1 + X_2 + \cdots + X_n$. The random variable X has a binomial distribution with parameters n and θ and its probability distribution function is

$$b(x; n, \theta) = \qquad\qquad\qquad\qquad \text{for } x = 0, 1, \ldots, n.$$

Examples. Flip n identical coins and count the number of heads. Ask n people a yes or no question on a survey.

2. a) Fill in the probability distribution function for the binomial random variable X above.

b) Use theorem 4.14 to find the mean and variance of X (as functions of n and θ).

Definition 4. Let $X_1, X_2, \ldots$ be a sequence of independent, identically distributed (iid) Bernoulli trials, all with probability of success θ. Let N be the trial on which the first success occurs. The random variable N is said to have a *geometric distribution* with parameter θ and its probability distribution function is

$$g(n : \theta) = \qquad\qquad\qquad\qquad \text{for } n = 1, 2, 3, \ldots$$

Examples. Flip a coin until the first heads appears. Roll a pair of dice until you first get a pair of sixes. Ask people a yes or no question until you first get a yes.

3. Fill in the probability distribution function of N.

4. The goal now is to find the mean and variance of this distribution. The following questions outline one method, but other methods exist (the moment-generating function or using the mean and variance of the negative binomial distribution below, for example). Use whatever approach you like.

a) Calculate the expected value of N. Hint: differentiate both sides of $\frac{1}{1-x} = \sum_{n=0}^{\infty} x^n$.

b) Calculate $E\left[N(N+1)\right]$. Hint: take the second derivative of both sides of $\frac{1}{1-x} = \sum_{n=0}^{\infty} x^n$.

c) Use $E(N)$ and $E\left[N(N+1)\right]$ to calculate $V(N)$.

Definition 5. Let $X_1, X_2, \ldots$ be a sequence of independent, identically distributed (iid) Bernoulli random variables, all with probability of success θ. Now let N be the trial on which the k^{th} success occurs (so the possible values for N are $k, k+1, k+2, \ldots$). Such a random variable is said to have a *negative binomial (or binomial waiting-time or Pascal) distribution* with parameters k and θ and its probability distribution function is

$$b^*(n; k, \theta) = \qquad\qquad\qquad\qquad \text{for } n = k, k+1, k+2, \ldots$$

5. Fill in the probability distribution function of the negative binomial random variable N.

Challenge. Find the mean and variance of the negative binomial distribution.

Definition 6. A random variable with the probability distribution function

$$p(x; \lambda) = \frac{\lambda^x e^{-\lambda}}{x!} \text{ for } x = 0, 1, 2, \ldots$$

is said to have a *Poisson distribution* with parameter $\lambda > 0$.

6. Find the moment-generating function of a Poisson random variable and use it to calculate the mean and variance of the distribution.