1. **Coin Flips.** Imagine flipping an unbiased coin N times. Let $N_{\rm H}$ be the number of heads results, and $f = N_{\rm H}/N$ be the fraction of such results.

   (i) What is the probability of observing a particular sequence of heads (H) and tails (T) results, e.g., H T T T T H H T T H H T H... ?

   (ii) How many possible flip sequences yield exactly $N_{\rm H}$ heads results? Your answer should involve the factorial function, $M! \equiv M \times (M-1) \times (M-2) \times \ldots \times 3 \times 2 \times 1$.

   (iii) Write an exact equation for the probability $P(N_{\rm H})$ of observing $N_{\rm H}$ heads results when the coin is flipped N times.

   (iv) Stirling's approximation,

   $$\ln M! \approx M \ln M - M \quad \text{for large } M,$$

   allows you to simplify your result in part (iii) assuming N is very large. First, we consider a hand-wavy way to "derive" Stirling's approximation. We know that the integral of a function $g(x)$ can be approximated by a Riemann sum:

   $$\int_a^b g(x) \approx \sum_{i=0}^{(b-a)/\Delta x} g(a + i\Delta x)\Delta x$$

   when Δx is sufficiently small. If $b - a \gg 1$, $\Delta x = 1$ can be small enough for a good approximation of the integral. Follow this line of argument to show Stirling's approximation. (Hint: you will want to consider $g(x) = \ln x$ and an appropriate choice of a and b).

   (v) Armed with Stirling's approximation, show that $P(N_{\rm H})$ can be written in the large deviation form

   $$P(N_{\rm H}) \approx e^{-NI(f)}$$

   when N is sufficiently large to justify Stirling's approximation. Identify and plot $I(f)$ as a function of f.

   (vi) Reflect on the fact that I does not depend on N. In other words the extensive (large) part of the problem has dropped out and only impacts the probability through the factor that multiplies I. This is a major simplification! You might have thought that the term in the exponent should have higher powers of N, but it does not.

2. **A Macroscopic Number of Spins.** Now imagine the physical scenario of making a single measurement (as opposed to repeated coin flips) of $N \gg 1$ noninteracting spin-$1/2$ particles. In that measurement, the observed z-component of each spin is up or down with equal probability.

   (i) What is the probability of observing a number $N_{\rm up}$ of up spins in a given observation? Write your answer in terms of the fraction $f = N_{\rm up}/N$.

   (ii) Although $N_{\rm up} = N/2$ is the most likely observation, a typical measurement will not yield *exactly* half the spins pointing up. For Avogadro's number of spins, $N \approx 10^{24}$, estimate the relative probability of a small deviation $\delta = 0.0000001$ from the ideal fraction, i.e., calculate $P(f = 0.5 + \delta)/P(f = 0.5)$. Your numerical answer need not be highly accurate; just determine the order of magnitude. (For this purpose, Taylor expansion of $\ln P$ about $\delta = 0$ is both permitted and a good idea).

(iii) In the large N limit, the distribution $P(f)$ becomes well-approximated by a Gaussian distribution. Derive this Gaussian using your Taylor expansion in (ii). Use that result to determine the Gaussian distribution $P(N_{\rm up})$.

(iv) What does your result imply about the reproducibility of measurements on macroscopic systems? Imagine repeating the measurement many times and getting the fraction f of up spins from each repeated experiment. Would the variance of those values of f be large or small? Specifically, how would the variance depend on N? What about the variance of the values of $N_{\rm up}$ collected from each experiment?

3. **Fun with Gaussians.** One of the most important continuous distributions is the Gaussian distribution

$$P(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

I assume you probably already have some familiarity with this distribution. These problems are either a crash course or a refresher. I do not consider it important that you "discover" the standard tricks on your own, but I do think you should be aware of them. Please talk to me or to other students if this is foreign!

(i) **Normalization:** Show that $\int dx\, P(x) = 1$. [*Standard trick: Convert to a two dimensional integral over a joint distribution of identical, independent Gaussians for the x and y coordinates. Remind yourself what a Jacobian is.*]

(ii) **Mean:** Show that $\langle x \rangle = \int_{-\infty}^{\infty} dx\, xP(x) = \mu$. [*Standard trick: Substitute $u = x - \mu$ and notice that an integral cancels by symmetry.*]

(iii) **Variance:** Show that $\left\langle (\delta x)^2 \right\rangle = \int_{-\infty}^{\infty} dx\, (x-\mu)^2 P(x) = \sigma^2$. [*Standard trick: From (i) and symmetry find the integral of $e^{-\alpha x^2}$ from 0 to ∞. Differentiate with respect to α.*]

(iv) **Generating function:** Show that $\left\langle e^{\beta x} \right\rangle = \int_{-\infty}^{\infty} dx\, e^{\beta x} P(x) = \exp\left(\beta\mu + \frac{\sigma^2\beta^2}{2}\right)$. [*Standard trick: Complete the square.*]

(v) **Cumulant generating function:** The cumulant generating function, $\ln\left\langle e^{\beta x} \right\rangle$ follows simply from (iv). Comment on how this can be used to verify that the mean and variance are μ and σ^2, respectively. What does the Gaussian distribution's cumulant generating function tell you about the value of the higher order cumulants.