PHYSICS 210B : NONEQUILIBRIUM STATISTICAL PHYSICS
HW ASSIGNMENT #1 : PROBABILITY

(0) Study the eight worked example problems linked on the HOMEWORK web page.

(1) A six-sided die is loaded in such a way that it is twice as likely to yield an even number than an odd number when thrown.

(a) Find the distribution $\{p_n\}$ consistent with maximum entropy.

(b) Assuming the maximum entropy distribution, what is the probability that three consecutive rolls of this die will total up to seven?

(2) Show that the Poisson distribution $P_\nu(n) = \frac{1}{n!}\nu^n\,e^{-\nu}$ for the discrete variable $n \in \mathbb{Z}_{\geq 0}$ tends to a Gaussian in the limit $\nu \to \infty$.

(3) The probability density for a random variable x is given by the Lorentzian,

$$P(x) = \frac{\gamma}{\pi}\cdot\frac{1}{x^2 + \gamma^2}\ .$$

Consider the sum $X = \sum_{i=1}^{N} x_i$, where each x_i is independently distributed according to $P(x_i)$.

(a) Find the distribution $P_N(X)$. Does it satisfy the central limit theorem? Why or why not?

(b) Find the probability $\Pi_N(Y)$ that $|X_N| < Y$, where $Y > 0$ is arbitrary.

(4) Frequentist and Bayesian statistics can sometimes lead to different conclusions. You have a coin of unknown origin. You assume that flipping the coin is a Bernoulli process, *i.e.* the flips are independent and each flip has a probability p to end up heads and probability $1 - p$ to end up tails.

(a) You perform 14 flips of the coin and you observe the sequence $\{HHTHTHHHTTHHHH\}$. As a frequentist, what is your estimate of p?

(b) What is your frequentist estimate for the probability that the next two flips will each end up heads? Would you bet on this event?

(c) Now suppose you are a Bayesian. You view p as having its own distribution. The likelihood $f(\text{data}|p)$ is still given by the Bernoulli distribution with the parameter p. For the prior $\pi(p)$, assume a Beta distribution,

$$\pi(p|\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\,\Gamma(\beta)}\,p^{\alpha-1}\,(1-p)^{\beta-1}\ \ .$$

where α and β are hyperparameters. Compute the posterior distribution $\pi(p|\text{data},\alpha,\beta)$.

(d) What is the posterior predictive probability $f(\text{HH}|\text{data}, \alpha, \beta)$?

(e) Since *a priori* we don't know anything about the coin, it seems sensible to choose $\alpha = \beta = 1$ initially, corresponding to a flat prior for p. What is the numerical value of the probability to get two heads in a row? Would you bet on it?

(5) Consider the family of distributions

$$f(\boldsymbol{k}|\lambda) = \prod_{j=1}^{N} \frac{\lambda^{k_j} e^{-\lambda}}{k_j!} \, ,$$

corresponding to a set of N independent discrete events characterized by a Poisson process with Poisson parameter λ. Show that

$$\pi(\lambda|\alpha, \beta) = \frac{\alpha^{\beta}}{\Gamma(\beta)} \lambda^{\beta-1} e^{-\alpha\lambda} \, ,$$

is a family of priors, each normalized on $\lambda \in [0, \infty)$, which is conjugate to the likelihood distributions $f(\boldsymbol{k}|\lambda)$.

(6) Professor Jones begins his academic career full of hope that his postdoctoral work, on relativistic corrections to the band structure of crystalline astatine under high pressure, will eventually be recognized with a Nobel Prize in Physics. Being of Bayesian convictions, Jones initially assumes he will win the prize with probability θ, where θ is uniformly distributed on $[0, 1]$ to reflect Jones' ignorance.

(a) After N years of failing to win the prize, compute Jones's chances to win in year $N + 1$ by performing a Bayesian update on his prior distribution.

(b) Jones' graduate student points out that Jones' prior is not parameterization-independent. He suggests Jones redo his calculations, assuming initially the Jeffreys prior for the Bernoulli process. What then are Jones' chances after his N year drought?

(c) Professor Smith, of the Economics Department, joined the faculty the same year as Jones. His graduate research, which concluded that poor people have less purchasing power than rich people, was recognized with a Nobel Prize in his fifth year. Like Jones, Smith is a Bayesian, whose initial prior distribution was taken to be uniform. What is the probability he will win a second Nobel Prize in year 11? If instead Smith were a frequentist, how would he assess his chances in year 11?