Fluctuations and thermodynamics

In the previous section we calculated the energy and particle number fluctuations in the canonical and grand canonical ensembles. Considering how the *relative* fluctuations depend on the system size, we obtain

$$\left.\frac{\sigma_E}{\langle E\rangle}\right|_{\text{canonical ensemble}} = \frac{\sqrt{k_B T^2 \frac{\partial \langle E\rangle}{\partial T}}}{\langle E\rangle} \sim \frac{1}{\sqrt{\langle E\rangle}}$$

$$\left.\frac{\sigma_N}{\langle N\rangle}\right|_{\text{grand canonical ens.}} = \frac{\sqrt{k_B T \frac{\partial \langle N\rangle}{\partial \mu}}}{\langle N\rangle} \sim \frac{1}{\sqrt{\langle N\rangle}}$$

In both cases the relative fluctuations decay as the $-\frac{1}{2}$ power of the system size. In the $N \to \infty$ limit, called *thermodynamic limit*, the fluctuating quantities (when rescaling with the system size) become definite, not random. Thus we can replace $\langle E\rangle$ with E etc. This is why statistical mechanics is the microscopic foundation of thermodynamics.

In many cases fluctuations are the aggregate effect of many independent contributions. To consider this case more rigorously, suppose X_i are *iid* (independent, identically distributed) random variables, with $\langle X_i\rangle = \mu$ and $\mathrm{Var}(X_i) = \sigma^2$. Then the *Central limit theorem* states that

$$Z_n := \frac{\overbrace{X_1 + \cdots + X_n}^{S_n} - n\mu}{\sqrt{n}\sigma} \xrightarrow{D} \mathcal{N}(0,1)$$

Here $\mathcal{N}(0,1)$ is the distribution of standard normal (Gaussian) random variables, ie. with zero mean and unit variance. The notation $\xrightarrow{D}$ means convergence in distribution:

$$\lim_{n\to\infty} P(Z_n < z) = P(\zeta < z)$$

where ζ is a standard normal random variable. Note that this is pointwise convergence of the cumulative distribution fuction, which is weaker than the convergence of the probibility density function.

The Central limit theorem is behind the fact that the normal distribution is so prevalent: for macroscopic fluctuationns often the microscopic contributions are sufficiently independent. As we have seen before the relative fluctuations of the sum decrease as $1/\sqrt{n}$:

$$\frac{\sigma_{S_n}}{\langle S_n\rangle} \to \frac{\sqrt{n}\sigma}{n\mu} \sim \frac{1}{\sqrt{n}}$$

A simple application is the one-dimensional random walk: X_i takes values ± 1 each with probility $1/2$. The resulting trajectory, $S_n = X_1 + \cdots + X_n$ is *like* a Gaussian variable with mean zero and standard deviation $\sqrt{n}$, when sufficiently coarse grained to remove the discreteness.

Certain important cases fall outside the applicability of the Central limit theorem, like distributions where the variance (or the mean as well) is undefined. One such example is the Cauchy (or Lorentz) distribution, defined by the probability density function

$$f(x) = \frac{1}{\pi(1+x^2)} \qquad \text{or} \qquad f(x) = \frac{1}{\pi\gamma\left(1 + \left(\frac{x-x_0}{\gamma}\right)^2\right)}$$

Surprisingly the average of n iid Cauchy random variables has the same distribution as just one, which means that if one deals with such quantities, taking averages is useless.

When generalising this phenomena one arrives at the concept of *stable distributions*: these are families of distributions where the sum of such random variables is from the same family. More formally, let $\mathrm{Fam}(\Theta)$ represent a family of distributions where Θ denotes all the parameters. Suppose X_1 and X_2 are from this family. If their linear combination is also from this family:

$$X_1 \sim \mathrm{Fam}(\Theta_1), \quad X_2 \sim \mathrm{Fam}(\Theta_2) \quad \Rightarrow \quad aX_1 + bX_2 \sim \mathrm{Fam}(\Theta_3) + c$$

then we call Fam a stable distribution.

We have seen that both the normal and the Cauchy are stable distributions. One more where the probability density function can be given in closed form is the Levy distribution:

$$f(x) = \sqrt{\frac{c}{2\pi}} \frac{e^{-c/(2x)}}{x^{3/2}}$$

which can be generalised to the 4-parameter Levy-skew-α-stable family.

This distribution underpins the *Levy flight*, which is similar to a random walk, but the increments are taken from a heavy tailed distribution,

$$f(x) \sim 1/|x|^{\alpha+1} \ , \qquad \text{where } 0 < \alpha < 2 \ .$$