

A statistical approach

So far in this class, you have learned classical thermodynamics. Starting next week, we will be studying **statistical mechanics**. Thermodynamics may look “theoretical” because it involves a lot of math, but ultimately it is an experimental science. Thermodynamics puts severe (and interesting) constraints on equations of state, but can never tell us what the equations of state actually are. Similarly, thermodynamics can allow us to measure one quantity and use it to predict the result of a very different measurement. But it could never give us the ideal gas law, or the internal energy of an ideal gas.

Statistical mechanics is the theoretical counterpart of **thermodynamics**. It’s how we can predict thermodynamic quantities from first principles. It also allows us to use thermodynamic measurements to extract microscopic properties of a system.

From quantum mechanics, you know that given a Hamiltonian describing a system, you could (in principle) solve for all the possible eigenstates and their energies. But how can you know *which* of those states a given system will be in? And given that state, how can you predict the result of interactions of the system with its surroundings, when you *don’t* know the hamiltonian or eigenstates of the surroundings? These are the questions that are answered by statistical mechanics.

Inputs to stat mech:

Energies and eigenstates of the Hamiltonian. We can actually get much of what interests us out of just the energies, just as we could compute all the thermodynamic properties from $U(S, V)$, if only we knew what it was... or from $G(T, p)$, as you did in your homework.

Output of stat mech:

Probabilities (at a given temperature) of each energy eigenstate, U , S , p , H and all thermodynamic functions. Statistical mechanics is awkward, so we will mostly want to use thermodynamics approaches when we can. e.g. if we know U , T , and S , we can just use $F = U - TS$.

Large numbers

In macroscopic systems, there are *many* atoms and molecules, typically around 10^{23} . As a result, we have no hope of actually examining every possible eigenstate, nor could we practically determine the precise *microstate* of the system. Instead, we need to examine how likely various states are. Average properties become extremely well-defined when many things are averaged.

If I flip one coin, I’ll get 50% heads, but with a pretty large uncertainty. When I flip 100 coins, I get 50 heads ± 10 coins. If I flip 10^{22} , I will get 5×10^{21} heads $\pm 10^{11}$. This is a lot of uncertainty in the total number of heads, but a very small uncertainty in the fraction of coins that will end up being heads.