end up at its starting point $x = 0$. (As in Section 1.4, we can extract this probability by integrating the distribution $p(x)$ over the interval $-0.5 \leq x \leq 0.5$.) Instead, the interval within which we can expect to find the walker (with a constant 'one-sigma' or 68% probability) steadily grows, $-D\sqrt{t} \leq x \leq D\sqrt{t}$, with characteristic dependence on the square root of the time the diffusive process lasts.

Except in the trivial cases $p = 0$ or $q = 0$, diffusion also occurs when the drift velocity is non-zero. This is shown in the two figures below, considering a low but non-zero drift velocity on the left, and a high drift velocity on the right.

$V_{dr} \propto 2p-1$

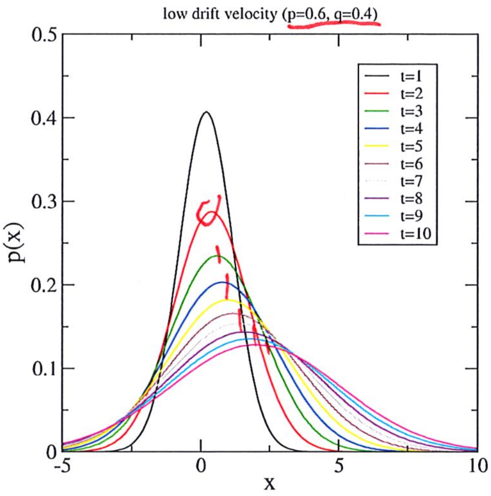

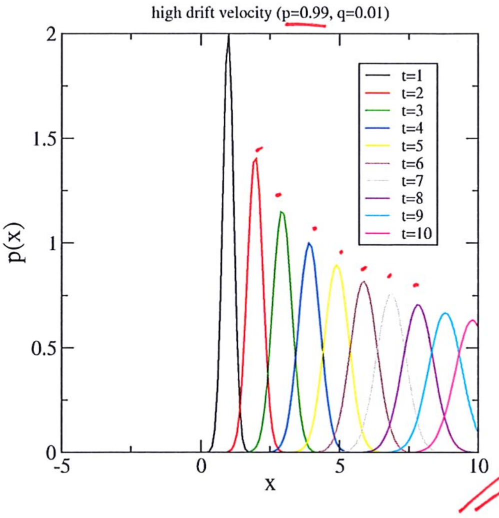

7 Feb.

8 Feb.

In the figure on the left, each individual probability distribution looks similar to the corresponding one for $v_{\mathrm{dr}} = 0$, but now their central peaks (and expectation values $\langle x \rangle$) drift steadily to the right. The distributions in the figure on the right look a bit different, but still diffuse to exhibit shorter and broader peaks as time goes on.

When $p \neq \frac{1}{2}$ so that $\langle x \rangle \neq 0$, it is interesting to compare the drift in the expectation value against the growth in fluctuations around $\langle x \rangle$ due to diffusion. We can do this by considering the following *relative* uncertainty:

$$\frac{\Delta x}{\langle x \rangle} = \frac{2\sqrt{\cancel{N}Npq}}{N(2p-1)} = \frac{1}{\sqrt{N}}\left(\frac{2\sqrt{pq}}{2p-1}\right) = \sqrt{\frac{\Delta t}{t}}\left(\frac{2\sqrt{pq}}{2p-1}\right) \propto \frac{1}{\sqrt{t}} \xrightarrow{t\to\infty} 0$$

$$\left.\begin{aligned} \Delta x &= 2\sqrt{Npq} \propto \sqrt{t} \\ \langle x \rangle &= N(2p-1) \propto t \end{aligned}\right\} \frac{1}{\sqrt{t}}$$

You should find that at large times this ratio vanishes proportionally to $1/\sqrt{t} \propto 1/\sqrt{N}$. Although the absolute uncertainty grows by diffusion, $\Delta x = D\sqrt{t}$, for $v_{\mathrm{dr}} \neq 0$ the linear drift in the expectation value becomes increasingly dominant as time goes on.

1.5.3 Applying the central limit theorem

$S = \sum_i x_i$

Based on our work in Section 1.4, we can see how to apply the central limit theorem to analyze this fixed-step-length random walk in one dimension, for large numbers of steps N or equivalently large times $t = N\Delta t$. Each step in the random walk can be considered an independent and identically distributed random variable x_i, analogous to each spin of the roulette wheel. The corresponding probability space involves only two possible outcomes: a step of length $\ell = 1$ to the right or to the left with probability p or q, respectively. From this we can easily compute the mean and variance *of the single-step process*:

$$\mu = \langle x_i \rangle = \sum_i x_i P_i = (1)p + (-1)q = 2p - 1 = \frac{\langle x \rangle}{N}$$

$$\langle x_i^2 \rangle = p + q = 1$$

$$\sigma^2 = \langle x_i^2 \rangle - \langle x_i \rangle^2 = \cancel{1} - (4p^2 - 4p + \cancel{1}) = 4p(1-p) = 4pq = \frac{\Delta x}{N}$$

The final position x of the walker after N steps is exactly the sum over these x_i given in Eq. 9. Its probability distribution $p(x)$ from the central limit theorem is therefore obtained directly from these single-step μ and σ^2, which we can also express in terms of the drift velocity and diffusion constant:

$D^2 = 2pq/$

$V_{dr} = N(2p-1)$

$$\begin{aligned} p(x) &= \frac{1}{\sqrt{2\pi(4Npq)}} \exp\left[-\frac{(x - N(2p - \overset{\mu}{1}))^2}{8Npq}\right] \\ &= \frac{1}{\sqrt{2\pi D^2 t}} \exp\left[-\frac{(x - v_{\text{dr}}t)^2}{2D^2 t}\right], \end{aligned} \tag{17}$$

which was used to produce the three figures above. We could have jumped straight to the final line by considering Eq. 10 and noting

$$v_{\text{dr}}t = N(2p-1) = N\mu \qquad\qquad D^2 t = 4pq\frac{t}{\Delta t} = N\sigma^2. \tag{18}$$

While this dependence on p and q is specific to the particular fixed-step-length random walk we're currently considering, the results $v_{\text{dr}}t = N\mu$ and $D^2 t = N\sigma^2$ in Eq. 18 turn out to be generic. This is remarkable, because it means that the diffusive process as a whole is determined entirely by the single-step mean and variance. So long as μ and σ^2 are finite, we end up with Eq. 17 as the large-t probability distribution for any markovian random walk in a single variable x.

This result is related to the generality of the law of diffusion (Eq. 16), which we can recognize in the structure of Eq. 17. Since $t > 0$, the exponential in the gaussian distribution $p(x)$ peaks at the drifting expectation value $x = v_{\text{dr}}t = \langle x \rangle$.

The factor $(x - v_{\mathrm{dr}}t)^2$ simply quantifies the distance from this peak. As t increases, so does the factor $2D^2t$ dividing this $(x - v_{\mathrm{dr}}t)^2$, meaning that a larger distance from the peak is needed for the overall argument of the exponential to reach a given value—in other words, the peak becomes broader. This in turn requires a shorter peak, reflected in the $\frac{1}{\sqrt{t}}$ in the overall coefficient, which is set by requiring $\int p(x)\, dx = 1$. In other words, the law of diffusion holds whenever the central limit theorem is applicable. This requires that the mean and variance of the single-step process are finite, and in the computer project we will numerically investigate the *anomalous diffusion* that occurs when this condition is not satisfied.

Unit 2: Micro-canonical ensemble

2.1 Statistical ensembles and thermodynamic equilibrium

We begin this unit by formalizing the concept of *statistical ensembles*, introduced by J. Willard Gibbs in the early 1900s. Building on the probability foundations laid above, we will be interested in 'experiments' that simply allow a collection of degrees of freedom to evolve in time, subject to certain constraints. At a given time t_1, the arrangement of these degrees of freedom defines the state ω_1 of the system.

As a concrete example, consider a system of *spins* — arrows that can point either 'up' or 'down' — arranged in a line. Such spin systems will appear several times in the remainder of this module, since in addition to obeying simple mathematics analogous to flipping coins, spins also serve as good models of physical systems such as magnetic molecules. What would be a representative state (or *configuration*) for a system of $N = 8$ spins? How many distinct states are there for this system?

$\uparrow\downarrow\uparrow\uparrow\uparrow\downarrow\uparrow\uparrow$

Total # : $2 \cdot 2 \cdots 2 = 2^N$

$= 2^8 = 256$

At a different time t_2, the system's state ω_2 is generally different from ω_1. However, there are some measurements we can perform that always produce the same outcome even as the system's state changes over time. These measurements define *conserved quantities*, such as the number of spins considered in the example above.

Another important conserved quantity is the total energy E *inside* an isolated (or 'closed') system,

$$E(\omega_1) = E(\omega_2).$$

The conservation of energy is presumably a familiar concept, and you may also know that it can be rigorously proven through Emmy Noether's first theorem.[3] Because statistical physics was first developed when conservation of energy was primarily an empirical observation rather than a proven result, it was given a more grandiose name: the **first law of thermodynamics**. Another way of stating the first law is that any change in the internal energy of one particular system Ω must be matched by an equal and opposite change in the energy of some other system with which Ω is in contact. This will be important when we consider thermodynamic cycles later in the term.

---

[3] There are complications when considering the dynamical space-time of general relativity, but that's beyond the scope of this module.

For now, let's return to the example above, and endow the spin system with an internal energy by placing it in a magnetic field of strength H. That is, if a spin is parallel to the field, it contributes energy $-H$ to the total energy E of the system. If a spin is anti-parallel to the field, it instead contributes energy H. For later convenience, we define a positive magnetic field $H > 0$ to point upward, and also define n_+ to be the number of spins pointing upward — parallel to the field and therefore contributing *negative* energy. Similarly, the remaining $n_- = N - n_+$ downward-pointing spins are anti-aligned with the field and contribute positive energy. What is the total energy E of the system in terms of n_+ and n_-? What is E for the representative 8-spin state you wrote down above? What fraction of the states of the spin system have this energy?

$H \Uparrow$ $\quad \uparrow \quad\quad \downarrow$

$\Delta E = -H \quad\quad \Delta E = H$

$E = -H(12-8) = -4H < 0$

$$E = -Hn_+ + H(N-n_+)$$
$$= -H(2n_+ - N)$$

$\#\ w/E = -4H : \binom{8}{6} = \frac{8!}{6! \cdot 2} = 28$

Fraction : $\frac{28}{256} = 7/64$

If instead we consider $N \sim 10^{23}$ hydrogen (H_2) molecules in a container, we can write a simple expression for the internal energy E by treating each molecule as a *point-like particle*, with no size or structure. In this case each molecule contributes only its kinetic energy to E,

$$E = \frac{m}{2}\sum_{n=1}^{N} \vec{v}_n^{\,2} = \frac{1}{2m}\sum_{n=1}^{N} \vec{p}_n^{\,2},$$

where $\vec{v}_n$ is the velocity of the nth molecule, $\vec{p}_n = m\vec{v}_n$ is its momentum, and all molecules have exactly the same mass m.

As forecast at the start of the module, we treat the time evolution of any given system as a stochastic process in which the system probabilistically adopts a sequence of states $\omega_i \in \Omega$:

$$\omega_1 \longrightarrow \omega_2 \longrightarrow \omega_3 \longrightarrow \omega_4 \longrightarrow \cdots$$

This approach is a matter of practicality rather than one of principle. In principle, Newton's laws would allow us to predict the exact time evolution of (say) 10^{23} hydrogen molecules, but only by specifying 10^{23} initial conditions and solving 10^{23} differential equations. Since we cannot hope to record so much information or carry out so many computations, we instead apply probability theory in order to analyze these systems.

This leads us to the following core definition: A **statistical ensemble** is the set of all states $\Omega = \{\omega_1, \omega_2, \cdots\}$ that a system can possibly adopt through its time evolution. Each state ω_i has some probability p_i of being adopted by the system, so we can recognize a statistical ensemble as a probability space.

Because these states ω_i depend on the 'microscopic' degrees of freedom that compose the overall system, we will refer to them as **micro-states** from now on. From the definition of probability in Section 1.1, we have the requirement $\sum_i p_i = 1$, which simply means that the system must be in *some* micro-state at any point in time. The fact that time evolution cannot change any conserved quantities, as discussed above, means that such conserved quantities *characterize* statistical ensembles. We will define different types of statistical ensemble that depend on the specific set of conserved quantities.

First we define the **micro-canonical ensemble** to be a statistical ensemble characterized by conserved internal energy E and conserved number of degrees of freedom N — which we will call **particle number** for short.

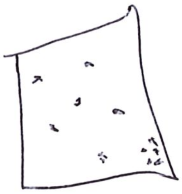

According to the discussion above, this means that a system governed by the micro-canonical ensemble is *isolated* in the sense that it cannot exchange energy or particles with any other system.

Now that the micro-canonical ensemble is defined, we can connect it to our intuition from everyday physical systems. Let's consider a collection of particles moving around and bouncing (or '*scattering*') off each other in a sealed container. To a first approximation, this should describe the behaviour of air in a room, which our lived experience indicates is spread quite uniformly throughout the room in a way that is stable as time passes. We do not expect all the air in a room to be concentrated in any one corner, nor do we expect strong collective gusts of wind without some clear external influence.

These qualitative expectations illustrate the idea of **thermodynamic equilibrium**, an axiomatic concept in statistical physics.[4] We can mathematically define thermodynamic equilibrium through the probabilities p_i that appear in the micro-canonical ensemble.

A micro-canonical system Ω with M micro-states ω_i is in thermodynamic equilibrium if and only if all probabilities p_i are equal. If M is finite, the requirement $\sum_i p_i = 1$ implies

$$p_i = \frac{1}{M}. \tag{19}$$

The full meaning and significance of this definition are not immediately obvious, and we will continue exploring them through consideration of derived quantities such as entropy and temperature. First, it's important to emphasize that this equilibrium is *dynamic*: There is not a single 'equilibrium state' that the system approaches. Instead, the equilibrium system continues probabilistically adopting all possible micro-states as it evolves in time.

---

[4] Our expectation that physical systems generically evolve towards thermodynamic equilibrium as time passes is more formally expressed as the ergodic hypothesis.

2.2 Entropy and its properties

2.2.1 Definition of entropy

We can gain further insight into thermodynamic equilibrium by considering a famous derived quantity.

The **entropy** of a statistical ensemble Ω with a countable number of micro-states M is defined to be

$$S = -\sum_{i=1}^{M} p_i \log p_i, \tag{20}$$

where p_i is the probability for micro-state ω_i to occur. Unless otherwise specified, "$\log$" indicates the natural logarithm with base e.

When the system under consideration is in thermodynamic equilibrium, we expect derived quantities such as the entropy to be stable over time, even as different micro-states are probabilistically adopted. This implies that such derived quantities are functions of the conserved quantities that are the same for all micro-states. Therefore, for the micro-canonical ensemble, the equilibrium entropy $S(E, N)$ is a function of the conserved energy and particle number.

8 Feb.

By inserting Eq. 19 into Eq. 20 you can quickly compute a simple expression for the entropy of a micro-canonical ensemble in thermodynamic equilibrium:

Your result should depend only on the number of micro-states M, and diverge as $M \to \infty$. While the energy E and particle number N are not explicit in this expression, $\{E, N, M\}$ are inter-related and might be expressed in terms of each other depending on the particular situation under consideration. For example, what is the equilibrium entropy of the system of N spins considered above, if the magnetic field is turned off, $H = 0$? What is the entropy if $E = 0$ with $H > 0$ (which requires $n_+ = n_-$)?