
THERMODYNAMICS FOR IDIOTS: THIS HAS BEEN ONLY LIGHTLY PROOFREAD

Conor Rowan
Aerospace Engineering
University of Colorado Boulder
Boulder, CO 80303
conor.rowan@colorado.edu

ABSTRACT

Given that all thermodynamic phenomena can be explained with simple microscopic collision models governed by Newtonian mechanics, it is not clear why pedagogical presentations don't make this connection more explicit. Core concepts like entropy and pressure are given a straightforward, operational meaning when studied from the perspective of the microscopic model of a gas. Yet, anecdotal experience suggests that thermodynamics is poorly understood by many students and researchers in mechanics. Why not lean on the micro-model as a tool to build intuition? Of course, focus on the microscopic model is the domain of statistical mechanics, but this often comes at the expense of illuminating the macroscopic thermodynamic phenomena of interest to the engineer. In other words, the classical perspective is faced with the impossible task of explaining concepts like entropy without the micro-model, and by relying too heavily on abstract mathematics, the statistical perspective often struggles to give physical intuition for the origin of macroscopic concepts like pressure. In these notes, we try to straddle the classical and statistical perspectives in order to present some basic concepts from thermodynamics in an intuitive manner. Our perspective is that no thermodynamic concept is clear unless its explanation originates in the physics of the micro-model. Math is only to be trusted as a tool to carry out physical arguments! Additionally, our emphasis will be on showing how the formulas we discover can be used in calculations outside the cherry-picked settings of isothermal and adiabatic processes familiar from introductory texts on classical thermodynamics.

Keywords Classical thermodynamics · Statistical thermodynamics · Continuum thermodynamics

1 Model problem

We start our exposition with the simple model problem shown in Figure 1. A gas is modeled as a collection of N hard spheres of radius r and mass m in a two-dimensional box of volume V . Each gas particle has a corresponding velocity \mathbf{v}_i , where $i = 1, 2, \dots, N$ indexes the particles. The particles undergo elastic collisions with each other and the walls of the box, where an elastic collision is one that conserves kinetic energy and momentum. See Appendix A for a lazy discussion of the equations for elastic collisions. We will plan to numerically simulate the collisions of gas particles in order to "play" with some of the thermodynamic concepts we introduce. Thus, from a practical point of view, we must assume the gas particles have a finite radius r . In spite of their finite size, we will not consider rotational kinetic energy of the gas particles. The gas, understood as the collection of particles, stores energy only in the form of the particle's motion. We will define the total energy of the gas to be

$$E = \frac{1}{2}m \sum_{i=1}^N \|\mathbf{v}_i\|^2. \quad (1)$$

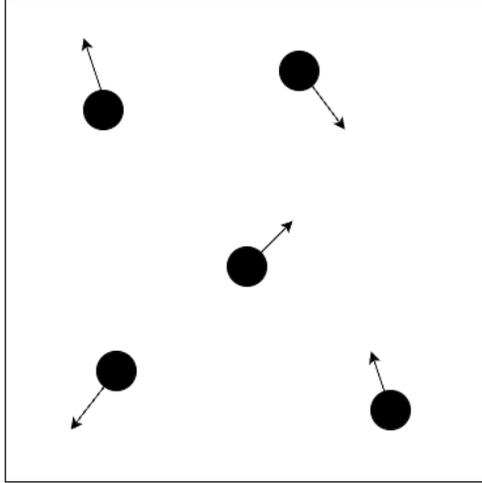


Figure 1: A gas is modeled as a large number of hard spheres contained in a box. The gas particles undergo elastic collisions with each other and the walls of the box. The goal of thermodynamics is to make statements about this system without explicitly modeling the state of each and every gas particle.

Imagine starting a numerical simulation of the collisions with a set of initial velocity vectors $\{\mathbf{v}_i^0\}_{i=1}^N$. The initial kinetic energy $E^0 = \frac{1}{2}m \sum_{i=1}^N \|\mathbf{v}_i^0\|^2$ will be conserved over time, but the momentum in each of the two coordinate directions will not. Though collisions with particles conserve both energy and momentum, collisions with the wall only conserve energy. This is a consequence of energy being a scalar quantity and momentum being a vectorial quantity.

We assume that there are two variables we can control for the system shown in Figure 1: the total energy E and the volume V . We will assume for simplicity that the number of particles N is fixed. This will be the case throughout the notes. We are interested in describing how the system depends on these quantities and how it evolves when they are varied. Note that we do not yet know what "describing the system" means. This is a very deep point. Obviously, we can describe the system by tracking the position and velocity (or momentum) of each and every particle. Even a cubic centimeter of a real-life gas would be intractably complex if described in this way. Classical thermodynamics is the result of a successful search for a simpler description of systems like the one shown in Figure 1. That temperature, pressure, and entropy are real and useful means of describing systems like this is in no way obvious. We will slowly build up to these macroscopic thermodynamic variables.

The first question we will pose to the system is as follows: is there anything we can say about the motion of the gas particles, given by the time-varying velocities $\{\mathbf{v}_i(t)\}_{i=1}^N$, other than the conservation of energy:

$$E_0 = \frac{1}{2}m \sum_{i=1}^N \|\mathbf{v}_i(t)\|^2, \quad \forall t. \quad (2)$$

Of course, we know that the dynamics of the particles are governed by Newtonian mechanics, but this does not help us make any more precise statements about the individual velocities than Eq. (2). Of course, we could simulate the collisions, but this is intractable when the number of particles is large. What if we imagine the kinetic energy of each particle as a random variable, where the collisions act like a chaotic mixing effect which introduces a kind of randomness? An example of a system governed by deterministic mechanical principles which exhibits effectively random behavior is the Galton board, shown in Figure 2. Here, deterministic chaos leads to a system which is accurately modeled as random. Let us assume that the kinetic energy K of each particle is a random variable with distribution $p(K)$. Computing the total energy of the system is akin to independently sampling the distribution N times, computing the kinetic energy, and summing the results. But we know that the total energy should equal E_0 at all times. When the velocity is a random variable, we think of the total energy at each time t as the sum of N samples. How can we guarantee that this sample always has the correct energy? Fortunately, the law of large numbers states that

$$\lim_{N \rightarrow \infty} \sum_{i=1}^N K_i \rightarrow E_0 \quad (3)$$

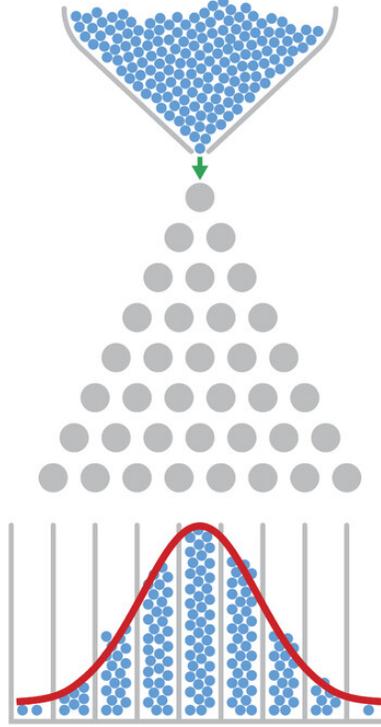


Figure 2: The Galton board is an example of a deterministic system which gives rise to effectively random behavior. In principle, whether the ball goes to the left or right at each peg could be determined, but it is accurately modeled as random. The result of the random walk of many balls through the board is an approximation of a normal distribution.

where K_i is the i -th independent sample from the distribution of kinetic energy and $\langle K \rangle = E_0$. Thus, the law of large number ensures that the sample energy converges to the true energy as the number of particles becomes large. No problem, as for the system to even be considered as a gas, the number of particles needs to be enormous. How can we determine the distribution $p(K)$? We know that it must have mean E_0 and that $K \geq 0$. One idea is to choose the distribution that maximizes uncertainty as to the value of the particle's kinetic energy subject to the two constraints. This is what is called a "maximum entropy method," as entropy is a functional defined on probability distributions that measures the extent of uncertainty in a random variable. The entropy is defined as

$$H(K) = - \int p(K) \log p(K) dK. \quad (4)$$

Note that this is the *information entropy* as opposed to the *thermodynamic entropy* which we will discuss later. For our purposes, it suffices to think of H as spitting out a large number when a probability distribution is "spread out" and a small number when it is concentrated. The information entropy measures how random a random variable is. This might appear to overlap with the notion of variance, but that entropy is a better measure of uncertainty/randomness than variance can be demonstrated with an example. Remember that the variance is defined as $\text{Var}(K) = \int K^2 p(K) dK - \langle K \rangle^2$. Consider the following distribution:

$$p(K) = \frac{1}{2} \left(\delta(K - 1) + \delta(K - 100) \right), \quad K \in [0, \infty]. \quad (5)$$

It can be verified that the variance of this distribution is large, whereas the information entropy is small. This continuous random variable only takes on two values with non-zero probability, so we expect the uncertainty to be small. But

because the variance looks at the spread around the mean, which in this case is large, it does not produce a small value. The point of finding a maximally uncertain probability distribution is to make the fewest possible commitments apart from satisfying the constraints. We are looking for a distribution that encodes no prior knowledge about the distribution of velocities. We can find this distribution by using Lagrange multipliers to enforce the constraints. The Lagrange function reads:

$$\mathcal{L} = H(K) + \lambda_1 \left(\int Kp(K)dK - E_0 \right) + \lambda_2 \left(\int p(K)dK - 1 \right) + \lambda_3 K. \quad (6)$$

The first Lagrange multiplier enforces the condition on the mean, the second enforces the normalization of the distribution, and the third handles the inequality constraint. The details of finding a stationary points of the Lagrangian are a bit messy and besides the point. The distribution which maximizes entropy and satisfies the constraints is the following:

$$p(K) = \frac{1}{E_0} e^{-K/E_0}. \quad (7)$$

Based on our maximum entropy assumption, the kinetic energy of the particles follows an exponential distribution. The larger the energy of the system, the fatter the tail of the distribution. This is what is called the "Boltzmann distribution." We note that this distribution can be transformed using the definition of the kinetic energy to obtain the "Maxwell-Boltzmann distribution" for the speed (magnitude of velocity) of the particles.

At this point, we have no idea whether the choice to treat the particle's energy as random—or the assumption that such a distribution would maximize entropy—is valid. We can test this empirically with a numerical simulation. See Appendix A for discussion of the details of the collision model. We initialize a simulation with a given initial kinetic energy and $N = 250$ particles, compute the Boltzmann distribution from Eq. (3), then plot it against a histogram of the sample distribution of the particle's kinetic energies at three randomly chosen time steps. This is shown in Figure 3. The agreement with the Boltzmann distribution is quite good, which supports our choice to model the particles with an entropy-maximizing random variable.

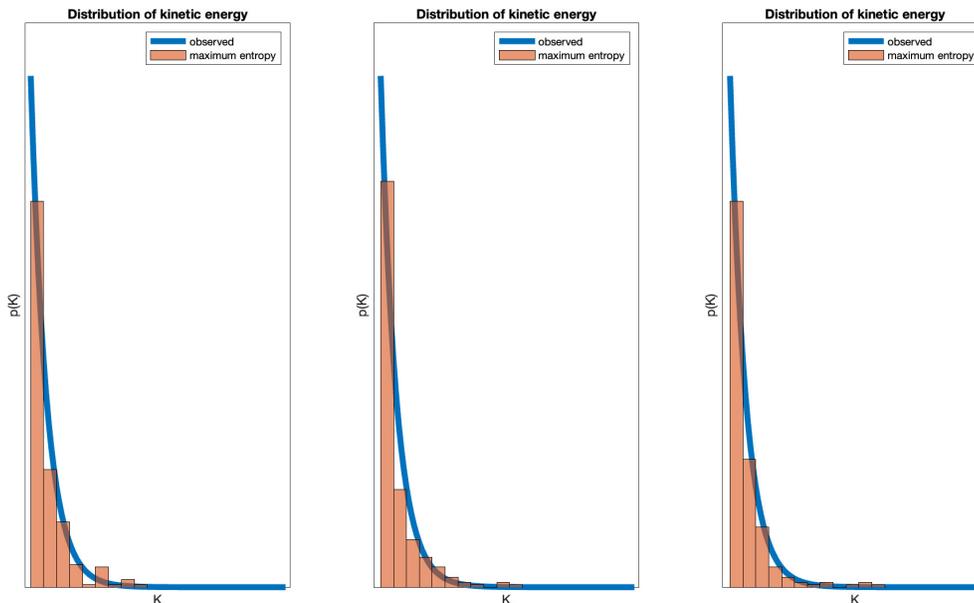


Figure 3: Checking agreement of sample distribution of kinetic energies with the Boltzmann distribution.

We have accomplished our original goal of making a statement about the behavior of individual particles from knowledge of an aggregate quantity (total energy). The price that we pay for not explicitly tracking the dynamics of each particle is that we only have statistical knowledge about their behavior. This motivates an important distinction, which will be useful to us going forward. The total energy E is a "macrostate" of the system, meaning a state that aggregates over all the particles. The individual kinetic energies $\{K_i\}_{i=1}^N$ are a "microstate" of the system, meaning a state defined on each and every particle separately. The law of large numbers in Eq. (3) showed that N samples taken from $p(K)$ represent a microstate corresponding to the macrostate of total energy E . Inherent in this statement is that there are many microstates which correspond to a given macrostate.

2 Temperature

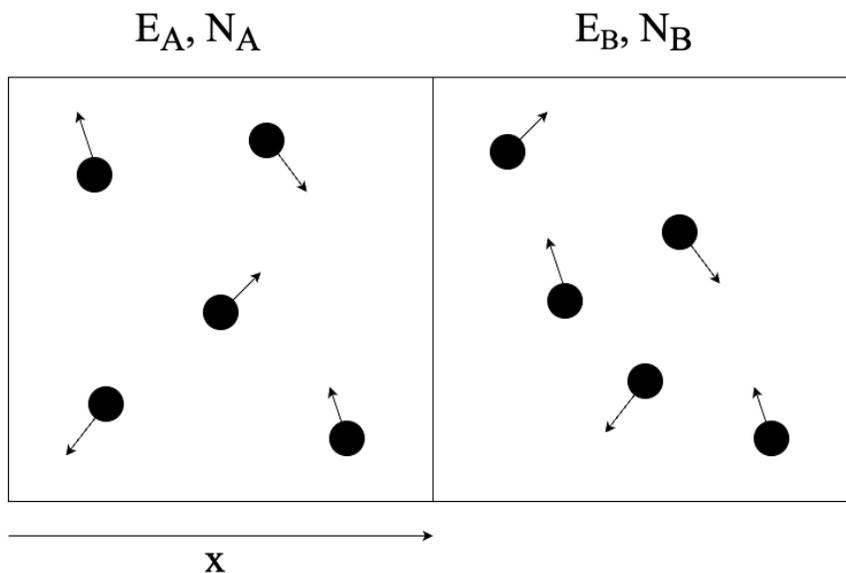


Figure 4: Model problem to motivate the concept of temperature. A movable dividing wall separates two gases from each other. The number of particles in the two boxes are N_A and N_B respectively. The goal is to find the "equilibrium" energies E_A and E_B and the position of the wall x .

We proposed using the variables of total energy and volume to describe the gas. We showed that the Boltzmann distribution for the kinetic energy of individual particles could be derived with a maximum entropy assumption and that it agreed with some simulation results. We have not made use of the system's volume yet. We suggest that the total energy and volume are the most intuitive ways to characterize a gas, as they both can be obtained with simple measurements. We know, of course, that temperature is an important thermodynamic variable. In this section, we will suggest answers to the questions: what is temperature? and why is it useful? To this end, consider the system shown in Figure 4. Two gases are separated by a dividing wall at position x . The gas in box A (on the left) has N_A particles and box B (on the right) has N_B particles. Say that the two gases exchange energy and volume with each other through the dividing wall at position x , and that there is no energy loss to the surroundings. The total energy of the system is taken to be the known value E_T . Our goal is to determine the "equilibrium" configuration of the wall position and the two energies. Equilibrium means the state that the system converges to for any given energy and volume initial conditions. What principle governs this equilibrium configuration?

Remember we said in the previous section that for a given energy macrostate, there were many corresponding microstates. In our new system, we now have both energy and volume macrostates. If we take the box to be the unit square, the volume macrostate of box A is $V_A = x$, and the volume macrostate of box B is $V_B = 1 - x$. These are macrostates because there are many sets of particle positions that correspond to the gas occupying a certain volume. The velocities are not relevant to the volume. This correspondence between individual particle positions and macroscopic volume is analogous to the idea that there are many sets of particle kinetic energies that correspond to a given total energy. With this idea of multiplicity of microstates in hand, we can make the following assertion: the equilibrium

configuration of the system is the one which maximizes the number of microstates corresponding to the macrostates E_A, V_A, E_B, V_B . In order to compute this configuration, we must formalize the notion of multiplicity of microstates. We require some way of "counting" the number of microstates corresponding to a macrostate.

We will use the variable $\Omega = \Omega(E, V)$ to store the number of microstates consistent with an energy/volume macrostate. As we will see, the values of the macroscopic variables determine the quantity of potential microstates. A problem immediately arises—the energy and volume are continuous variables, and thus we cannot actually enumerate microstates. A set of particles can occupy a given volume in infinitely many ways, just as a set of particles can have a given total energy in infinitely many ways. Our workaround for this issue is to imagine the microstate of the system as a point in "phase space," meaning the high-dimensional space which specifies the position and momentum coordinates of each of the particles, and to compute Ω as the volume of phase space which is consistent with total energy E and volume V . We will denote the position vector of the i -th particle as \mathbf{q}_i and its momentum as \mathbf{p}_i . The microstate of the system is fully specified by the position and momenta of all of the particles:

$$\text{microstate} = \begin{bmatrix} \mathbf{q}_1 \\ \mathbf{p}_1 \\ \mathbf{q}_2 \\ \mathbf{p}_2 \\ \vdots \\ \mathbf{q}_N \\ \mathbf{p}_N \end{bmatrix}, \quad E = \frac{1}{2m} \sum_{i=1}^N \|\mathbf{p}_i\|^2, \quad V = \max(\{q_{i1}\}_{i=1}^N) \times \max(\{q_{i2}\}_{i=1}^N). \quad (8)$$

The formula for the volume macrostate asserts that the particles live in a box defined by $[0, a] \times [0, b]$ and is not meant to be taken too seriously. Of course, unless the number of particles is huge, the maximum of the two position components need not correspond to the limits of the container of the gas. Eq. (8) is simply meant to show how a macrostate is obtained by operating on a microstate. Now, we must compute the volume of phase space corresponding to given macroscopic variables E and V . The momentum components lie on a hypersphere with radius \sqrt{E} in a $2N$ -dimensional space. The position components live in a $2N$ -dimensional hyper-rectangle. The volume factorizes and can be written as

$$\Omega(E, V) = c(2N)E^{(2N-1)/2}a^N b^N, \quad (9)$$

where we compute the surface area of the hypersphere as a "measure" of the size of the region in phase space corresponding to the energy macrostate. The constant c depends on the dimension of the hypersphere ($c = 4\pi$ for a standard 3-dimensional sphere). It is a bit odd to use the surface area for the momentum components and the volume for the position components, as the surface area technically has zero volume. This is a bit of a weird notion of volume, admittedly. The point of the function Ω is to have *some measure* of the size of the region in phase space consistent with a macrostate, and it is fair to say that Eq. (9) accomplishes this. In statistical mechanics, Ω is called the "density of states." We can now generalize to the system shown in Figure 4. Noting again that the height and total width of the box are unity, the density of states for the two boxes are

$$\Omega(E_A, V_A) = c_A E_A^{(2N_A-1)/2} x^{N_A}, \quad \Omega(E_B, V_B) = c_B E_B^{(2N_B-1)/2} (1-x)^{N_B}. \quad (10)$$

We note that just as the volume is a conserved quantity, the total energy of the system is a conserved quantity. This means that $E_T = E_A + E_B + E_{AB}$ where E_{AB} is an interaction energy for the two systems. If no such interaction energy existed, the systems would have no way of exchanging energy with one another. In practice, we might think of E_{AB} as the kinetic energy of the dividing wall, which moves frictionlessly in response to collisions of gas particles on either side. We make the assumption $E_{AB} \approx 0$, which allows us to write $E_B = E_T - E_A$. This just says the kinetic energy of the wall is small compared to the total energy of the gas. When we assume that the microstate in each box are chosen independently of each other, the density of states for the whole system is $\Omega_T = \Omega_A \Omega_B$. Furthermore, if we assume that each microstate is equally probable, we can write a probability distribution for the macrostate of the system—which can be defined entirely in terms of the state of box A—as

$$p(E_A, x) = \frac{1}{Z} \Omega_T(E_A, x) = \frac{1}{Z} c_A c_B E_A^{(2N_A-1)/2} x^{N_A} (E_T - E_A)^{(2N_B-1)/2} (1-x)^{N_B}. \quad (11)$$

The constant Z is simply a normalization factor that enforces that this distribution integrates to 1. We have encoded the fact that certain macrostates of the system have more corresponding microstates. If all microstates are equally likely

(an assumption), then the macrostates with more corresponding microstates will occur with higher probability. This leads us to make the following claim: the equilibrium configuration of the system is the macrostate with maximum likelihood. In other words, we look for the maximizer of Eq. (11). The maximum value this density function takes on is not a probability (it could be greater than 1), rather it is a likelihood. We note that only integrals of density functions give probabilities. We can find the maximizer by computing

$$\begin{aligned} \frac{\partial p}{\partial E_A} &= \frac{2N_A - 1}{2} E_A^{(2N_A - 3)/2} (E_T - E_A)^{(2N_B - 1)/2} - E_A^{(2N_A - 1)/2} \frac{2N_B - 1}{2} (E_T - E_A)^{(2N_B - 3)/2} = 0, \\ \frac{\partial p}{\partial x} &= N_A x^{N_A - 1} (1 - x)^{N_B} - x^{N_A} N_B (1 - x)^{N_B - 1} = 0. \end{aligned} \quad (12)$$

We see that the normalization factor Z (which we never determined) and the two geometric coefficients c_A and c_B can be factored out of this expression and thus do not appear. If we arrange these two equations, we find that

$$\frac{E_A}{\frac{2N_A - 1}{2}} = \frac{E_T - E_A}{\frac{2N_B - 1}{2}}, \quad \frac{N_A}{x} = \frac{N_B}{1 - x}. \quad (13)$$

The second equation is very intuitive—the volume equilibrates so that the spatial density of particles on either side is equivalent. For example, the more particles in box A, the larger the volume of box A in equilibrium. The equation for the equilibrium configuration of energy is less obvious. If we use that $E_B = E_T - E_A$ and assume that N_A and N_B are large—an assumption which we have already made use of with the law of large numbers—we can approximate this expression as

$$\frac{E_A}{N_A} = \frac{E_B}{N_B}. \quad (14)$$

Thus, we see that equivalent average energy in the two boxes corresponds to the highest multiplicity of microstates. The equilibrium configuration for the system is thus equal average energy and equal particle densities. This motivates the concept of temperature: temperature is the energy-like quantity which equilibrates between two interacting thermodynamic systems. We will define the temperature of the two systems as

$$T_A = \frac{E_A}{N_A}, \quad T_B = \frac{E_B}{N_B}. \quad (15)$$

This contrasts slightly with traditional definitions of temperature but is similar in spirit¹. The point is that the energy itself does not come into equilibrium between the two systems, the average energy does. It is not clear to us why the approximation $(2N - 1)/2 \approx N$ needs to be made.

We simulate this system by introducing a dividing wall with a given mass that moves frictionlessly in response to momentum imparted by collisions with gas particles. The wall's motion acts as a mechanism for transferring energy from one box to the other. When there are more particles on one side of the wall, or when the particles on one side have a higher average energy, there are more frequent and more energetic collisions and the wall tends to move away from this side. By virtue of moving away from these particles, the wall causes more frequent and more energetic collisions with particles on the other side, effectively transferring energy from one side to another. Example simulation results are shown in Figure 5. We randomly initialize $N_A = 150$ particles on the left and $N_B = 100$ particles on the right with two different average kinetic energies. The simulation shows that the system tends toward the equilibrium configuration predicted by Eq. (13), though there is significant noise. We expect that as the number of particles becomes extremely large, this noise should decrease. Running the simulation for longer may also show a trend toward decreasing fluctuations around the predicted equilibrium configuration. We believe this to be a nice pedagogical example for connecting thermodynamic concepts to the underlying gas particle collision model.

¹The definition of temperature typically involves Boltzmann's constant k and is given as the average energy per velocity component. We will see how far we can get ignoring these details. Many expressions we derive will differ from textbook presentations by factors involving 2 (two dimensional gas), 3 (three dimensional gas), and k .

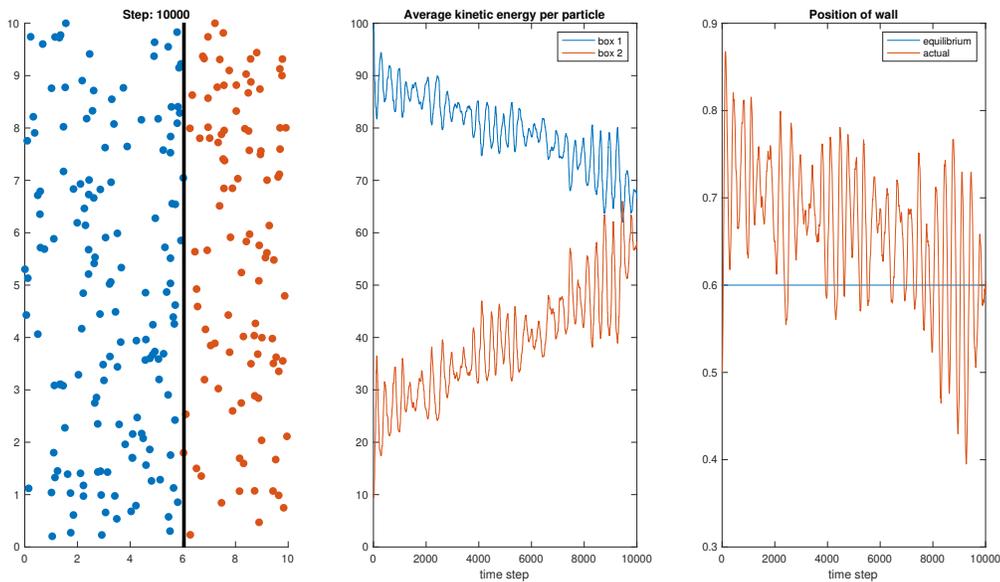


Figure 5: A snapshot of the simulation of the two box system (left), the trend toward temperature equilibrium (center), and fluctuation around equilibrium volume (right). The simulation results lend support to our assumption that the system takes on a "maximally likely" configuration.

3 Pressure

We claimed that the most natural variables to describe the state of the gas were the total energy and the volume. From here, the temperature emerged as an energy-like quantity that equilibrated for two thermodynamic systems in contact with one another. The pressure of the gas is another quantity familiar from classical thermodynamics. We will propose a microscopic basis for the pressure and then experiment with this definition using the simulation. To start, we imagine a particle labeled by i with a velocity component v_i^\perp which is perpendicular to a wall of width a . We take this wall to be rigid, unlike the movable wall of the previous section. Then, by definition, a collision with the wall results in a reversal of the perpendicular velocity component. The magnitude of force associated with an individual collision at time t_i is

$$|F_i^\perp| = \left| \frac{d}{dt}(mv_i^\perp) \right|_{t_i}. \quad (16)$$

We will not concern ourselves with the sign of the force because we know that the particle always exerts an outward force on the wall. As a result of the elastic collision, the perpendicular velocity component experience a jump discontinuity of $2m|v_i^\perp|$. The perpendicular velocity of a particle undergoing a single collision at t_i can be written as $v_i^\perp(t)(1 - 2u(t - t_i))$ where $u(t)$ is the unit step function. This means that

$$|F_i^\perp| = 2m|v_i^\perp|\delta(t - t_i). \quad (17)$$

We will hypothesize that the "pressure" on the wall is the average collision force on the wall over a suitable time interval Δt normalized by the width of the wall. The average is taken over all the particles that collide with the wall in the given time frame. This reads

$$p(t) = \frac{1}{a\Delta t} \int_t^{t+\Delta t} \sum_{j=1}^C 2m|v_j^\perp|\delta(t - t_j)dt = \frac{2m}{a\Delta t} \sum_{j=1}^C |v_j^\perp|. \quad (18)$$

Luckily, the delta function—which is a consequence of the obviously nonphysical assumption that the collisions happen over an infinitesimal time interval—is well-behaved under an integral. Eq. (18) a bit of a confusing looking, but it

expresses a simple concept. The sum is taken over all particles that collide with the wall in the given time increment, thus C represents the total number of collisions and j indexes the collision. To be clear, the same particle may undergo multiple collisions with the wall in the given time frame, so j is not necessarily an index on the particles themselves. Consider again the isolated system shown in Figure 1. For pressure to be a useful thermodynamic concept, it should neither depend on the time t (unless the state system is evolving somehow) nor the interval over which we average Δt . In other words, the pressure should be a static property of the gas like the total energy or volume. Furthermore, it should be the same for all walls of the container, regardless of their size. None of this is obvious, but we can motivate the validity of pressure as a thermodynamic property with the simulation. We take $N = 100$ particles with a given initial energy in an isolated system and compute the collision forces on each of the walls over the course of a 50 second simulation. Using Eq. (18), we compute the pressure versus time for various values of Δt . The results are shown in Figure 6. We see that as the length of the averaging time is increased, the pressure converges on a deterministic value that is independent of time. We argue that for a real thermodynamic system, where the number of particles is astronomical, practically any finite Δt will be sufficient to obtain this convergence of the pressure toward a deterministic value, as the number of collisions will be enormous. We note that the pressure can be thought of as another useful instantiation of the law of large numbers. If the collisions are thought of as random variables (with given distributions of the frequency and perpendicular velocity component), we can imagine the pressure equation as taking a "sample" mean from some function of these two random variables. When the number of samples is large, meaning that collisions are frequent, the law of large numbers says that the sample mean converges to the true mean. We also note that further numerical experiments indicate that the pressure is equivalent on all walls, even when the box is not square.

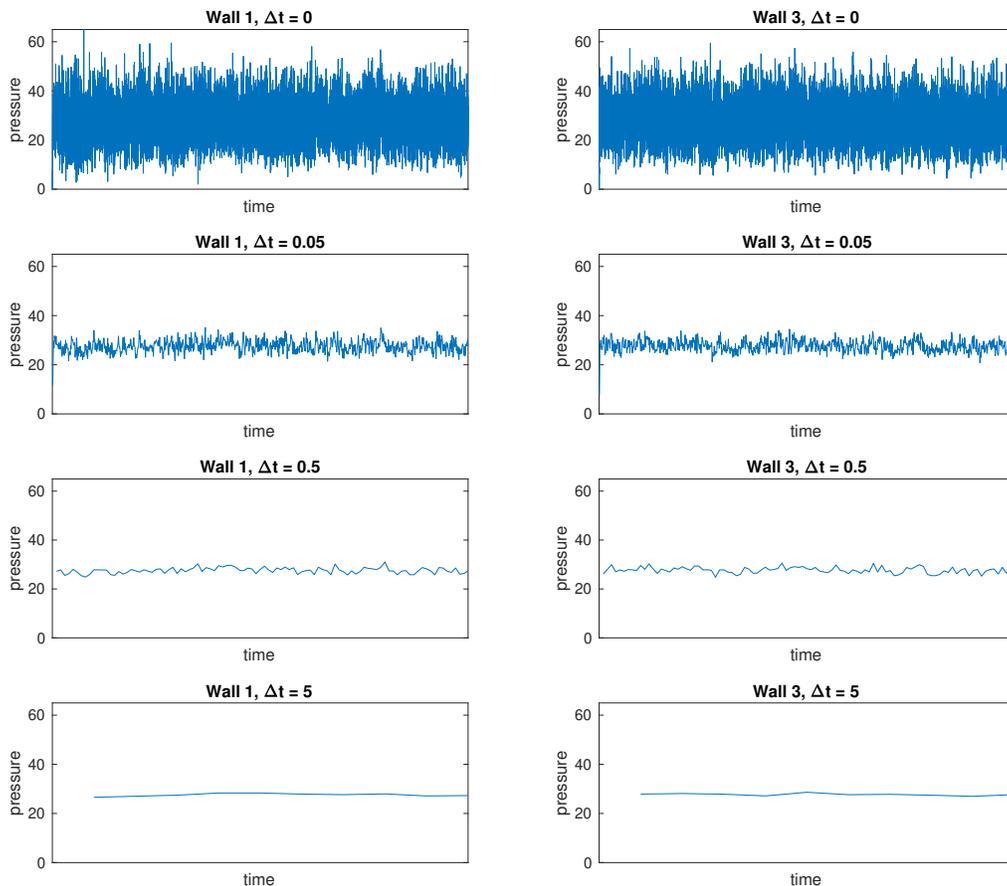


Figure 6: The pressure converges to a deterministic, time-independent value as the averaging time is increased.

4 Ideal gas law

We now have total energy, volume, temperature, and pressure in hand as concepts to make sense of the gas. It is natural to ask: how are these different quantities related, if at all? Of course, it is possible to arbitrarily specify the total energy and the volume. The total energy and the temperature are related through the number of particles. Per Eq. (18), the pressure is a function of the dimensions of the box, the number of particles (to the extent that the number of collisions depends on the number of particles), and the velocities of the particles. This suggests that the pressure may not be independent of the total energy and volume. We can perform a "back of the envelope" calculation to estimate the dependence of the pressure on the energy and volume. We will approximate the pressure as follows:

$$p = \frac{2m}{a\Delta t} \sum_{j=1}^C |v_j^\perp| \approx \frac{2mC}{a\Delta t} |v^{\text{typ}}|. \quad (19)$$

The notation " v^{typ} " is meant to indicate a typical perpendicular velocity component. This can be computed by assuming that the kinetic energy is evenly distributed among the two coordinate directions. With this assumption, we can write

$$v^{\text{typ}} = \sqrt{\frac{2}{m} \frac{E}{2N}}. \quad (20)$$

To obtain the typical velocity, we take half of the average kinetic energy to get the typical kinetic energy of one of a particle's velocity components. We again assume we are in a box of width a and height b . Given that Eq. (19) looks at the wall corresponding to the width of the box, we will say that a typical distance a particle needs to travel before a collision with the wall is $\ell = b/2$. On average, the particles are in the center of the box and are approaching the wall with a velocity of v^{typ} . These are all rough assumptions. We will use these quantities to estimate the typical time it takes a particle to collide with the wall. This is given simply by $t^{\text{typ}} = \ell/v^{\text{typ}}$. For a system with N particles, the number of collisions in a given time increment Δt is then estimated to be

$$C = \frac{\Delta t N v^{\text{typ}}}{b/2}. \quad (21)$$

We can plug this expression for the number of collisions and the typical velocity into the approximation of the pressure to obtain

$$p = \frac{2m}{a\Delta t} \frac{\Delta t N}{b/2} \sqrt{\frac{E}{mN}} \sqrt{\frac{E}{mN}} = \frac{4E}{V} = \frac{4NT}{V}. \quad (22)$$

From this very crude analysis, we have recovered the general mathematical form of the "ideal gas law," which is familiar from classical thermodynamics. At a fixed volume, the pressure is predicted to scale linearly with the total energy. At a fixed energy, the pressure is inversely proportional to the volume. We can explore these predictions with the simulation. We take $N = 300$ particles in a thermally isolated square container. The pressure is treated as a response variable and computed per Eq. (18). The volume and total energy of the particles can be controlled. See Figure 7 for results. The pressure-energy relation at a fixed volume of $V = 1$ is in fact linear, but our crude calculation does not accurately predict the slope. We fit the line and obtain an "empirical" ideal gas law as $p = 0.95E/V$. Then at a fixed energy, we can use this to predict the pressure-volume relation and compare to the simulation results. The prediction is accurate, indicating that the relation $PV \propto E$ holds for our toy model of the gas. Deriving the exact constant of proportionality is not of interest here. Without a doubt, the exact constant of proportionality can be derived from more careful analysis of the collision energies and frequencies as a function of the system's properties. We note that the true ideal gas for a two-dimensional gas in terms of pressure, volume and total energy is

$$PV = E, \quad (23)$$

which is very close to the law recovered by the simulation results. We will make use of this going forward.

5 First law of thermodynamics

The first law of thermodynamics is simply the conservation of energy. It says that

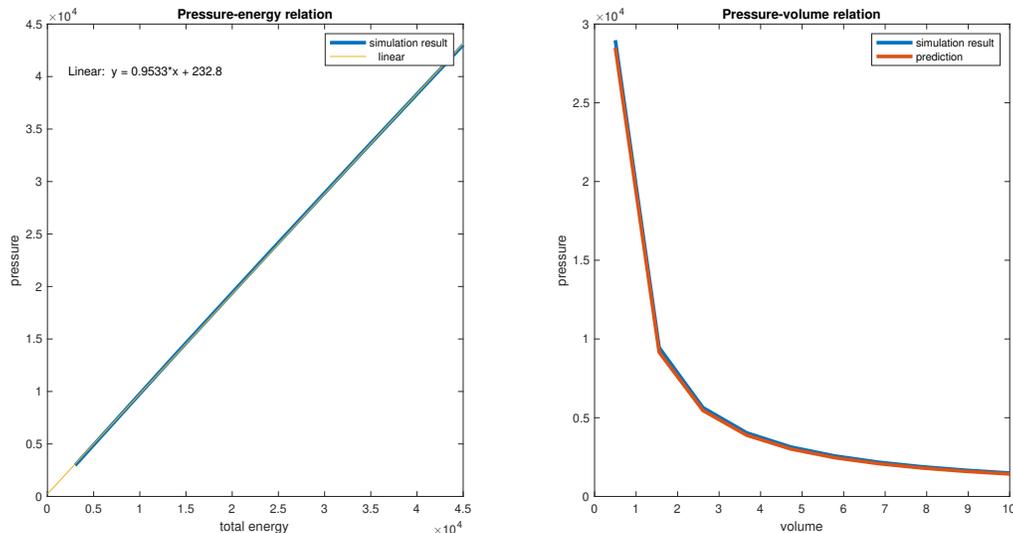


Figure 7: Our crude derivation of the ideal gas law makes specific predictions about the nature of the dependence of the pressure on the total energy and volume of the gas. Up to a constant of proportionality, these relationships are borne out by the simulation results.

$$\underbrace{\Delta E}_{\text{change in total energy}} = \underbrace{Q}_{\text{energy added}} - \underbrace{W}_{\text{work done}}. \quad (24)$$

It is simple to understand what this means when referring to the collision model of the gas. The total kinetic energy of the particles increases when energy is added to the system and decreases when the system does work. The mechanism of energy addition to the system is typically called "heat." The underlying physical mechanism of heat transfer is usually not spelled out in classical thermodynamics, though we saw with the system in Figure 4 that the addition of heat could come from a moving wall that knocked the particles in such a way that their kinetic energy tended to increase. Similarly, this same system does work, and thus decreases its total energy, when energy is transferred from the particles into the wall via a collision. Let us assume that we are comfortable with the idea that pressure is a well-defined thermodynamic quantity that is independent of time and constant on each of the container's walls. Our slipshod ideal gas law derivation and numerical experiment suggested that pressure is a function of energy and volume $p = p(E, V)$. We will use this in making sense of the first law of thermodynamics. We note that the heat added to the system is not a function of the system itself. It is something that arises from the particular environment that the system finds itself in. In other words, the heat input is not a function of the state of the system, it is something that the analyst must specify in order to do first law energy accounting on the system. This contrasts with the work done by the system, which will be dependent on the pressure. As we have seen, the pressure arises in response to the state of the system. It will be more convenient to look at the rate form of the first law, which comes from taking a time derivative of Eq. (24). This reads

$$\frac{dE}{dt} = \frac{dQ}{dt} - \frac{dW}{dt}. \quad (25)$$

We will define a thermal power input q with $Q = \int_{t_1}^{t_2} q(t) dt \implies \dot{Q}(t) = q(t)$. We imagine the energy accounting statement of Eq. (24) as summarizing the outcome of a specific process occurring over a time interval $t \in [t_1, t_2]$. The work done by the system over the course of this process is

$$W = \int_{t_1}^{t_2} p(E(t), V(t)) \frac{dV}{dt} dt. \quad (26)$$

The pressure depends on both the energy and the volume, which change over the course of the process. We are using the total energy and volume to describe the state of the system. This is an expression for the mechanical work done as a

result of volume changes. Like its thermal analogue, the mechanical power input is defined by this expression. We can thus write the rate form of the first law as

$$\frac{dE}{dt} = q(t) - p(E(t), V(t)) \frac{dV}{dt}. \quad (27)$$

Note that $q(t)$ and dV/dt are known quantities, as they define the parameterization of the process the gas undergoes. Using the ideal gas law, we can write this expression as

$$\frac{dE}{dt} + \left(\frac{1}{V(t)} \frac{dV}{dt} \right) E(t) = q(t), \quad E(t_1) = E_1, \quad t \in [t_1, t_2]. \quad (28)$$

The evolution of the energy of the system is governed by an inhomogeneous first-order linear differential equation with a time varying coefficient. The initial condition specifies the energy of the system before the process begins. The initial and final conditions on the volume are built into the parameterized path $V(t)$. We note that this presentation differs considerably from most in that there is no discussion of "exact" and "inexact" differentials (δQ vs. dE for example). We believe working with the rate form of the first law and parameterized paths for the quantities which can be directly controlled (heat input and volume) is significantly easier to understand. Eq. (28) is amenable to performing calculations on the evolution of the system's energy as a result of a parameterized process.

6 Thermodynamic entropy

We have made it through temperature, pressure, the ideal gas law, and the first law of thermodynamics. It is high time for the infamous second law of thermodynamics, about which no one seems to be able to make any sense. The second law says something like "thermodynamic entropy always increases." What is entropy? The textbook definition of the change in entropy S is usually given as

$$\Delta S = \int \frac{\delta Q}{T}. \quad (29)$$

What does this quantity physically represent, and why does it increase? Given most author's stubborn reliance on obscure and undefined concepts like "disorder," this is apparently a hard question to answer. Before we can answer either of the two questions posed above, we have to answer an even more basic one: what calculation does Eq. (29) even suggest? What does δQ mean and over what is this integral taken? We will rewrite this expression more transparently as

$$\Delta S = \int_{t_1}^{t_2} \frac{q(t)}{T(t)} dt = \int_{t_1}^{t_2} \frac{Nq(t)}{E(t)} dt. \quad (30)$$

The change in a system's entropy is a kind of summation of the process of adding heat to a system at a given temperature. We have rid of the temperature by using its definition as the average kinetic energy. We will do a calculation to motivate the concept of thermodynamic entropy. Let us consider a process whereby a system does work to expand from volume V_1 to V_2 and heat is added such that the total energy is constant $E(t) = E_1$. The volume change is parameterized by $V(t) = V_1 + (V_2 - V_1)t$ for $t \in [0, 1]$. Using Eq. (28), we can write the heat input required to keep the energy constant as

$$q(t) = \frac{1}{V_1 + (V_2 - V_1)t} (V_2 - V_1) E_1. \quad (31)$$

Heat needs to be added to keep the energy constant because the system is doing work in expanding. Plugging this into the expression for the change in entropy, we obtain

$$\Delta S = \int_{t_1}^{t_2} \frac{N}{E_1} \left(\frac{1}{V_1 + (V_2 - V_1)t} (V_2 - V_1) E_1 \right) dt = N \log \left(V_1 + (V_2 - V_1)t \right) \Big|_{t_1}^{t_2} = N \log \frac{V_2}{V_1}. \quad (32)$$

We can verify that this is consistent with the statistical thermodynamic definition of entropy. This will yield insight into what the quantity physically represents. Understanding what it represents will also help answer the question of why the second law holds. In statistical mechanics, the entropy and its change are defined as

$$S = \log \Omega, \quad \Delta S = \log \Omega_2 - \log \Omega_1 = \log \frac{\Omega_2}{\Omega_1}, \quad (33)$$

where $\Omega(E, V)$ represents the number of microstates corresponding to the energy and volume macrostate. This is the "density of states" function that measures the size of the region in phase space which is consistent with the macroscopic variables we discussed earlier. We can analyze the change in entropy of the volume expansion process using this notion of entropy. The density of states corresponding to the initial and final volume are

$$\Omega_1 = cE_1^N V_1^N, \quad \Omega_2 = cE_1^N V_2^N, \quad (34)$$

where we have again used the approximation $(2N - 1)/2 \approx N$ in computing the surface area of the hypersphere which restricts the momentum to have a fixed total kinetic energy. The entropy change computed this way agrees with the classical thermodynamic notion of entropy:

$$\Delta S = N \log \frac{V_2}{V_1}. \quad (35)$$

The statistical notion of entropy answers the question of what entropy is: it is a measure of the number of microstates corresponding to a macrostate. It is not clear that it is possible to "interpret" the classical entropy in its own terms. The "law" that entropy increases is really a claim about probabilities. If we imagine that all microstates are equally probable, as we did before, the second law codifies the fact that systems tend to evolve toward macrostates with larger numbers of corresponding microstates. This law can obviously be violated, especially when the number of particles is small. In the few particle setting, the high probability macrostates are not significantly more likely than lower probability macrostates. As the number of particles increases, the probability density governing macrostates concentrates, and the likelihood of high entropy macrostates is significantly larger. The second law says that systems tend to be found in the most likely configuration, where likely here means the largest number of microstates. It also says that if systems are initialized in a low entropy configuration, they will evolve toward high entropy configurations. We have not proven that the second law takes on a deterministic character as $N \rightarrow \infty$, but we have seen from the law of large numbers how probabilistic quantities can take on a deterministic character in this setting.

We can do another calculation to see how the classical and statistical definitions of entropy agree. This time, we will heat the system at constant volume. We will simply take $q(t) = q_0$ for $t \in [0, 1]$. In this case, the total energy increases as $E(t) = E_1 + q_0 t$. The classical definition of entropy states that

$$\Delta S = \int_0^1 \frac{Nq_0}{E_1 + q_0 t} dt = N \log \left(\frac{E_1 + q_0}{E_1} \right). \quad (36)$$

With a fixed volume V , the density of states for the initial and final configurations are given by

$$\Omega_1 = cE_1^N V^N, \quad \Omega_2 = c(E_1 + q_0)^N V^N. \quad (37)$$

The entropy change using this definition is seen to agree with the classical definition:

$$\Delta S = \log \frac{\Omega_2}{\Omega_1} = N \log \left(\frac{E_1 + q_0}{E_1} \right). \quad (38)$$

With the rate form of the first law given in Eq. (28), it is possible to compute entropy changes for complex thermodynamic processes where the volume changes and heat is added simultaneously. It is interesting to note that "proving" the equivalence between the two notions of entropy requires showing that

$$N \log \left(\frac{E_2 V_2}{E_1 V_1} \right) = N \int \frac{q(t)}{E(t)} dt, \quad (39)$$

where $E(t)$ satisfies the first law of thermodynamics with the given initial and final conditions for the energy and volume. It is not at all clear how to do this. We can ease our burden by looking instead at the rate of entropy production for an ideal gas predicted by the two definitions:

$$\frac{dS^{\text{statistical}}}{dt} = \frac{d}{dt} N \log EV = \frac{N}{EV} \left(\frac{dE}{dt} V + E \frac{dV}{dt} \right) = \frac{1}{T} \left(\frac{dE}{dt} + \frac{1}{V} \frac{dV}{dt} E \right) = \frac{q}{T},$$

$$\frac{dS^{\text{classical}}}{dt} = \frac{d}{dt} \int \frac{q}{T} dt = \frac{q}{T}, \quad (40)$$

where the last equality follows from the rate form of the first law for an ideal gas given in Eq. (28). We have used the definition of the temperature as the average energy. This shows that the two definitions of entropy predict the same rate of entropy generation. The rate form of entropy production is perhaps more illuminating anyway—the rate at which the number of microstates consistent with macrostate increases is proportional to the power input and inversely proportional to temperature. This suggests that when the temperature is small, a given energy input increases the number of states very rapidly.

7 Specific heats

Using our presentation of the first law, we can derive "specific heats" of the gas in the model problem. The specific heat is a coefficient that describes the rate at which the energy of the gas increases for a given heat input. The "constant volume" specific heat is trivial. The energy of the gas increases in lockstep with the heat input. In other words, when the volume is constant, we have $dE = dQ$ which defines the specific heat as $c_v = 1$. The constant pressure specific heat is a bit more complicated. This is the rate at which the internal energy increases for a given heat input when the gas expands (and thus does work) in order to keep the pressure constant. Using Eq. (28) with the ideal gas law $V = E/p$, we have

$$\frac{dE}{dt} + \frac{p}{E} \frac{1}{p} \frac{dE}{dt} E = 2 \frac{dE}{dt} = q(t) \implies dE = \frac{dQ}{2}. \quad (41)$$

The energy is half as sensitive to the heat input when the volume expands in order to keep the pressure constant. This is because the gas does work and thus loses energy as it expands. This implies that $c_p = 2$.

8 State variables and thermodynamic potentials

As we have already seen, the entropy can be related to the power input and temperature in rate form with

$$\frac{dS}{dt} = \frac{q}{T} \quad (42)$$

We can plug this expression into the rate form of the first law to obtain

$$\frac{dE}{dt} = T \frac{dS}{dt} - p \frac{dV}{dt}. \quad (43)$$

We are now at the point that we can clarify our description of the state of the thermodynamic system. We have so far introduced six variables relevant to the thermodynamic analysis: the number of particles N , the energy E , the volume V , the temperature T , the pressure p , and the entropy S . The equations relating these quantities are as follows:

$$\begin{aligned} T &= \frac{E}{N} && \text{(definition of temperature)} \\ p &= \frac{E}{V} && \text{(ideal gas law)} \\ S &= N \log EV && \text{(definition of entropy)} \\ dE &= T dS - p dV && \text{(first law of thermodynamics)} \end{aligned} \quad (44)$$

The heat input can be backed out of the changes in the entropy and the temperature history. The differential form of the first law looks like a total differential of the energy with two independent variables S and V . Taking this idea seriously, we write

$$T = \left(\frac{\partial E}{\partial S} \right) \Big|_V, \quad p = - \left(\frac{\partial E}{\partial V} \right) \Big|_S. \quad (45)$$

We can see that the entropy and volume are state variables by noting that

$$\oint \left[\frac{\partial E}{\partial S} \right] \cdot \left[\frac{dS}{dV} \right] = 0. \quad (46)$$

This is the fundamental theorem of calculus for line integrals. It may be simpler to imagine this as a line integral over a closed parameterized path $[S(t), V(t)]$. This says that the internal energy at a given value of S and V is fixed—there are no paths which bring us back to the same setting of the entropy and volume such that the internal energy differs from its original value. This means that these variables uniquely determine the state of the internal energy. This is not the case with all sets of thermodynamic quantities. For example, the heat input Q and the volume V do not uniquely determine the energy. The energy depends on the particular process that the system undergoes to obtain the given heat input and volume. The problem with using the entropy and volume as state variable is that unlike the volume, the entropy is not something which is easy to control. For example, if we somehow obtained an internal energy function $E(S, V)$ and wanted to use this to compute the pressure per Eq. (45), we would have to devise a strange system that stayed at constant entropy even while undergoing volume changes. Thus, we look for a new function whose derivatives yield quantities of interest which do not require strange operating conditions such as constant entropy. This can be accomplished with a Legendre transformation. With the internal energy, we think of the entropy as a control variable and the temperature as a response. We want to reverse this logic, and a particular Legendre transform accomplishes this. One Legendre transform of the internal energy is called the "Helmholtz free energy" and is given by

$$F = U - TS. \quad (47)$$

We can see that this switches the control and response variables by taking a total differential:

$$\begin{aligned} F = E - TS \implies dF = dE - SdT - TdS &= \left(\frac{\partial E}{\partial S} \right) dS + \left(\frac{\partial E}{\partial V} \right) dV - SdT - TdS \\ &= -pdV - SdT = \left(\frac{\partial F}{\partial V} \right) \Big|_T dV + \left(\frac{\partial F}{\partial T} \right) \Big|_V dT. \end{aligned} \quad (48)$$

With the Helmholtz energy, we have that the pressure and entropy are computed as

$$p = - \left(\frac{\partial F}{\partial V} \right) \Big|_T, \quad S = - \left(\frac{\partial F}{\partial T} \right) \Big|_V. \quad (49)$$

The temperature is the control variable and the entropy is the response variable when using the Helmholtz energy. This means that the pressure is computed with a volume derivative at constant temperature. Because this is a situation which is much simpler to realize in practice, the Helmholtz energy is useful for constitutive modeling. This is especially the case when the temperature of the system is constant, and we are only concerned with the force-like quantities such as pressure. This is why the Helmholtz energy appears in the solid mechanics literature, as the system is usually assumed to be at a constant temperature. The Helmholtz energy $F(T, V)$ is thus a common tool for constitutive modeling. It is an example of a "thermodynamic potential," which simply means a Legendre transform of the internal energy that is used to switch which are the control and response variables for the thermodynamic system. The ability to do this is a consequence of having a variety of ways to describe the state of the system.

9 Inexact differentials

With a notion of the thermodynamic state in hand, we can clarify what exactly is meant by an inexact differential. Quantities such as the heat Q and work W have "inexact" differentials when they can be taken in a closed loop through space of the thermodynamic state and not integrate to zero. Let's say that we are describing the thermodynamic state in terms of the temperature and volume. It is the case that

$$\oint \delta Q \neq 0 \quad (50)$$

where the line integral is taken over a closed path in the (S, V) space. We work the entropy-volume description of the system state as these are the natural variables for the internal energy. In other words, a system can undergo a cycle returning it to the same internal energy in such a way that heat was either absorbed or emitted. This might seem obvious, as we can explicitly control the heat input to be whatever we want. However, the catch is that the system must return to the same internal energy in the cycle, so it is not totally clear the heat input can be arbitrary. We have less control over the mechanical work, as the pressure is an explicit function of the state. The mechanical work can be written as $\delta W = p(S, V)dV$. It is also the case that

$$\oint \delta W = \int p(S(t), V(t)) \frac{dV}{dt} dt \neq 0 \quad (51)$$

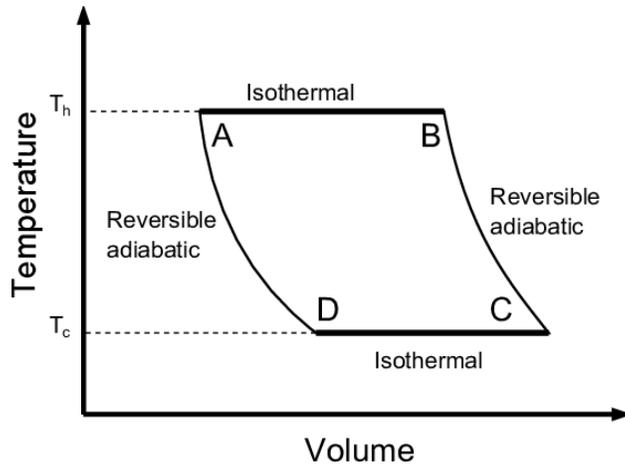


Figure 8: Undergoing a cyclic process, the Carnot cycle has a net gain of heat which is used to produce mechanical work.

for some path parameterized with the time-like variable t . Again, this means that if the system undergoes a thermodynamic process that starts and finishes at the same temperature and volume state, it need not be the case that the net work is zero. An example is the famous Carnot cycle, shown in Figure 8. The "inexact differential" notation is used to indicate that the heat and work need not be zero when integrated along a closed path.

To illustrate this, let us write the internal energy as a function of the entropy and volume. It is natural (and apparently simpler) to write $E = NT$, but we can use the definition of entropy to solve for the temperature in terms of the entropy and volume. This reads

$$S = N \log NTV^N \implies E = NT = \frac{1}{V^N} e^{S/N}. \quad (52)$$

We now compute a line integral around a closed path in the (S, V) space for the internal energy. We know that

$$\oint dE = \oint \delta Q - \delta W = 0, \quad \oint \delta Q = \oint \delta W \neq 0. \quad (53)$$

The first law says that a small change in the internal energy comes from heat and work inputs. When a line integral is taken around a closed of the thermodynamic state, the internal energy is unchanged. The net heat and work are not necessarily zero though. To illustrate this, we can actually compute line integrals of this sort. Assume that we have a closed parameterized $[S(t), V(t)]$ for $t \in [0, T]$. The line integral of the internal energy is

$$\oint dE = \int_0^T \frac{\partial E}{\partial S} \frac{dS}{dt} dt + \frac{\partial E}{\partial V} \frac{dV}{dt} dt = \int_0^T \underbrace{\frac{1}{NV^N} e^{S/N} \frac{dS}{dt}}_{\dot{Q}} dt - \underbrace{\frac{N}{V^{N+1}} e^{S/N} \frac{dV}{dt}}_{\dot{W}} dt \quad (54)$$

See this plot to experiment with this line integral. In general, the net heat input and work done are equal and opposite, though both non-zero. The differential heat and work are inexact differentials because they do not integrate to zero when taken in a closed loop. This is because they both form only part of a total differential of the internal energy function.

10 Diffusion

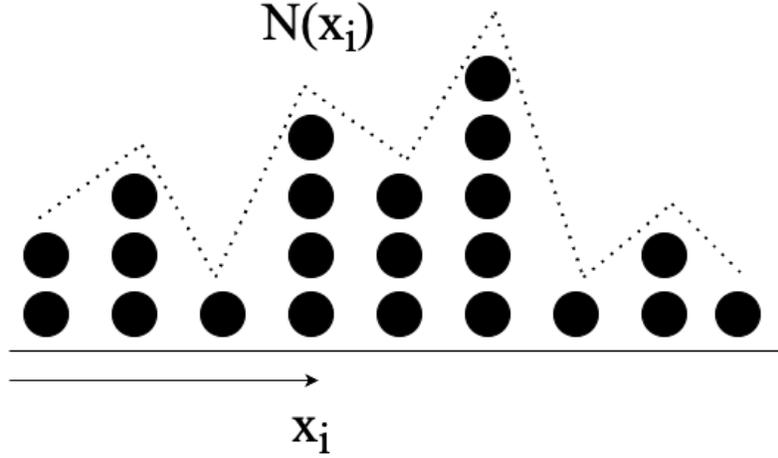


Figure 9: Continuous diffusion processes can be approximated with a simple underlying statistical model.

We can use an underlying statistical model to make sense of diffusion processes in a continuum. It is not necessary to model collisions to accomplish this. Imagine that a substance is present in a fluid medium inside of which it can move around freely, but there is no explicit forcing causing this motion. What happens? Intuitively, the substance spreads out to become uniform, as processes such as pouring milk into coffee would suggest. Using a one-dimensional micro model, it is simple to make sense of this process. Consider the problem set up shown in Figure 9. Individual particles of the substance are shown as dots, and the number of particles is counted at discrete positions with $N(x_i)$ for $i = 1, 2, \dots$. We will postulate that the diffusion process is driven by each particle undergoing a discrete random walk. The particles do not interact with each in any way, they simply take steps to occupy different positions. We will assume that the "transition probability" of each particle at position x_i is given by

$$p(x_{i-1}^{t+1}|x_i^t) = p(x_i^{t+1}|x_i^t) = p(x_{i+1}^{t+1}|x_i^t) = 1/3, \quad 0 \text{ otherwise.} \quad (55)$$

This says that at the next time step, a given particle either takes one step left, stays put, or takes one step right, all with equal probability. As $N(x_i) \rightarrow \infty$, and because the movement of the particles are independent, we can say from the law of large numbers that the average behavior of the random walk is observed deterministically. This means that $1/3$ of the particles go left, $1/3$ stay put, and $1/3$ go right. Let us assume that the each step of the random walk occurs over a time interval Δt . We note that each position x_i loses particles to its neighbors, but also acquires particles from the neighbors. Accounting for the influence from the neighbors, the change in particle count over a given time step is

$$\Delta N(x_i) = \left(N(x_{i-1}) + N(x_i) + N(x_{i+1}) \right) / 3 - N(x_i) = \frac{\Delta x^2}{3} \frac{N(x_{i-1}) - 2N(x_i) + N(x_{i+1}))}{\Delta x^2}. \quad (56)$$

Diving by the time interval Δt and taking the limit as the time interval and separation between particle positions goes to zero, we obtain

$$\lim_{\Delta t, \Delta x \rightarrow 0} \frac{\Delta N}{\Delta t} = \frac{\partial N}{\partial t} \propto \frac{\partial^2 N}{\partial x^2}. \quad (57)$$

This is the partial differential equation for diffusion. The rate of influx of particles at a given position over a given time step approximates a common finite difference approximation of the second spatial derivative. This little example shows that the limiting behavior of a large number of discrete random walk is a continuous diffusion process. Almost surely, a similar derivation for two- and three-dimensional random walks on a lattice will return similar finite difference rules for the continuous diffusion equation.

We note that another way to analyze this process is to write down a Markov chain corresponding to Eq. (55). If we assume periodic boundary conditions on the domain shown in Figure 9, it is quite simple to write down a finite Markov chain that describes this process. Consider for example the four-state Markov chain shown in Figure 10. Each state transition has a probability of $1/3$, where particles can either stay put or go to one of the two neighbors. The transition probability matrix $p(x_i^{t+1}|x_j^t)$ can be written as

$$\mathbf{P} = \begin{bmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{bmatrix}. \quad (58)$$

The stationary distribution \mathbf{s} of a Markov chain is obtained by repeatedly applying the transition matrix to some initial distribution over the states which we will call \mathbf{s}_0 . In the case of diffusion, we think of the "probability" of each state as just being the fraction of particles in that position. The stationary distribution is

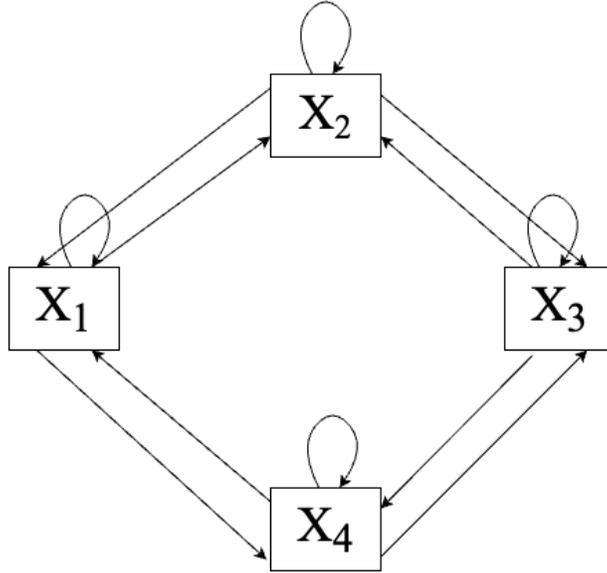


Figure 10: The diffusion process with periodic boundary conditions can be modeled as a Markov chain.

$$\mathbf{s} = \lim_{n \rightarrow \infty} \mathbf{P}^n \mathbf{s}_0 = [1, 1, 1, 1]/4. \quad (59)$$

The stationary distribution can be computed by looking at eigenvectors of the transition matrix. We see that the stationary distribution is uniform. In the case of discrete probability distributions, the uniform distribution maximizes information entropy, which was our measure of uncertainty used to derive the Boltzmann distribution. Though the information entropy is not equivalent to the statistical entropy involving micro and macrostates, we see that the diffusion process (modeled as a Markov chain) tends toward a state of maximum information entropy. It can also be shown that the approach to the stationary distribution is monotonic. This limiting behavior is independent of the initial distribution over the states. Note also that we do not use the law of large numbers in the Markov chain. It turns out that if you

consider the discretized function $\{N(x_i)\}$ as a "macrostate," we can compute the statistical entropy using the same formula

$$S = \log \Omega(N(x_i)). \quad (60)$$

By treating the particles in Figure 9 as "unlabeled", Ω counts the number of configurations of particles that give rise to the same count at each position. We are not specifying how exactly this is done, but clearly it is possible. Obviously there is lots of swapping and shuffling particles which can happen while keeping the count at each position unchanged. It turns out that the equivalent of a uniform distribution—the same number of particles in each position—also maximizes this notion of entropy. We can think of diffusion as the second law of thermodynamics in action. Systems evolve toward a state of maximum entropy. This means that the number of microstates consistent with the steady state of the system is maximum. Diffusion processes tend to spread out concentrated regions of substance, which is an example of entropy maximization.

11 Shear stress in fluids

We have explored how elastic collisions of gas particles can give rise to macroscopic concepts like pressure. The simple simulation we concocted motivates the microscopic origin of forces in gases. Arguably, the same model of hard sphere elastic collisions should apply to fluids. The difference with fluids is that we are interested not only in what happens at the boundaries of the container, but what happens inside the flow. This is a significantly more complicated problem than a gas confined to a box which we only interact with through macroscopic variables like pressure and temperature. In a fluid, we care about what is happening inside the box, so to speak. Note that the micro-model of a solid is something quite a bit more complicated than collisions of hard spheres, as the particles must be connected to each other in some way in order for the material to exhibit characteristics of a solid. This means that there are mechanisms for storing both potential and kinetic energy in the micro model of the solid, which obviously complicates analysis. For now, we will be content to gain some insight into the constitutive behavior of fluids from our numerical simulation. We will investigate the origin of shear stresses in fluids. The shear stress for channel flow of a fluid is often written as $\tau = \mu \partial u / \partial y$ where u is the horizontal velocity, y is the vertical coordinate, and μ is a constant of proportionality. In other words, when layers of fluid move past each other at different speeds, there is a force of resistance which tends to decrease the velocity gradient. We can verify that the micro model predicts such a phenomenon. We introduce two "layers" of fluid on the top and bottom half of the computational domain and give them different mean horizontal velocities. We set the vertical walls to have periodic boundary conditions so fluid that leaves the domain on the right immediately re-enters on the left. Figure 11 shows the mixing of the originally stratified layers as the simulation progresses. We expect that as the slow particles diffuse into the fast layer, momentum will tend to be transferred from the fast moving particles to the slow moving particles. A similar effect is expected when fast particles diffuse into the slow layer. Intuitively, the mean velocities of the upper and lower layers should equilibrate. This is what is shown in Figure 12. We also see that if we only look at the mean velocities, there is some "unaccounted for" energy. In other words, the "shear stress" has reduced the total macroscopic energy of the flow. This energy has gone into the random component of the kinetic energy of the particles, which we think of as heat. The simulation suggests how both shear stress and dissipation can be explained with the collision model.

12 Elastic micro-model

It sure would be interesting to derive elasticity from a micro-model wouldn't it? How is it possible that entropy is a measure of microstates consistent with a macrostate, that the microstates of a solid are totally different (they probably involve bond stretching), yet the definition of entropy for a solid continuum is equivalent to that of a gas? This seems impossible!

A Collision model

We need to determine what happens when two particles collide. This is more difficult than it seems like it should be. The two particles will have initial velocities \mathbf{v}_1^i and \mathbf{v}_2^i . We want to determine their final velocities \mathbf{v}_1^f and \mathbf{v}_2^f . Supposedly this is an exercise in high school physics. The collision needs to conserve kinetic energy and momentum in both coordinate directions. We will simply say the following equations need to hold :

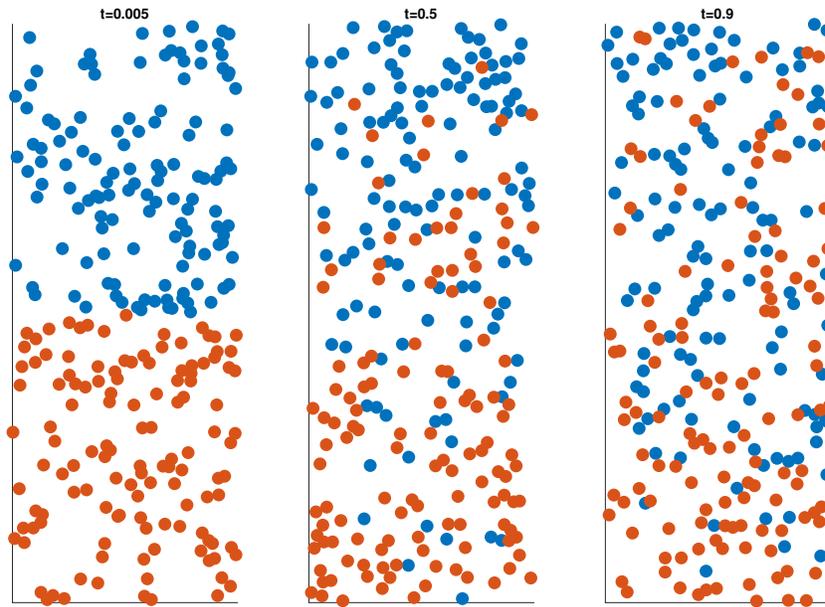


Figure 11: The red particles start with a small horizontal velocity, the blue particles with a large horizontal velocity. There are periodic boundary conditions on the vertical walls and elastic collisions with the horizontal walls. As time progresses, the particles become increasingly mixed.

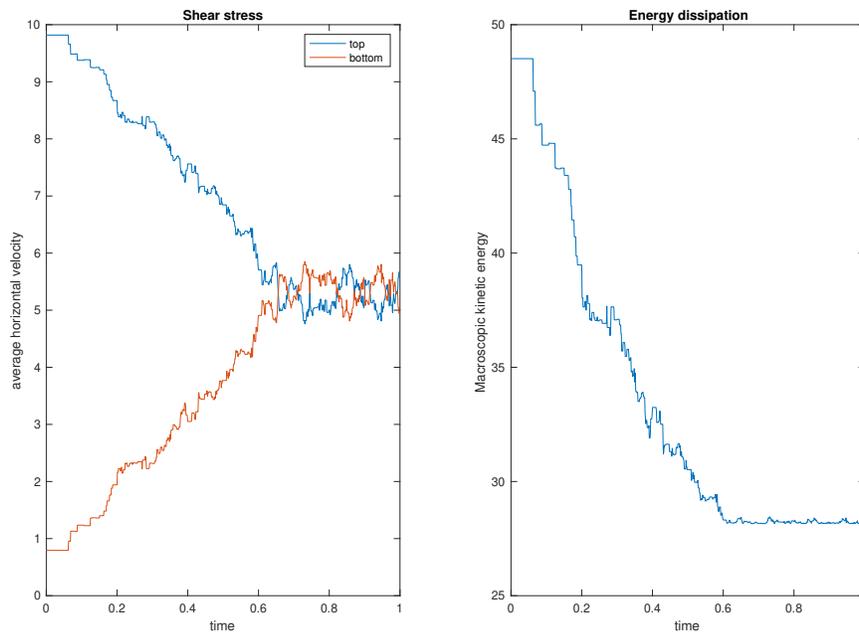


Figure 12: The fast and slow velocity equilibrate over time as a result of the collisions (left). The effective shear force converts some of the macroscopic translational energy into random motion. This is a dissipative effect, which converts mechanical energy into heat (right).

$$\begin{aligned} m\mathbf{v}_1^i + m\mathbf{v}_2^i &= m\mathbf{v}_1^f + m\mathbf{v}_2^f, \\ \frac{1}{2}m\|\mathbf{v}_1^i\|^2 + \frac{1}{2}m\|\mathbf{v}_2^i\|^2 &= \frac{1}{2}m\|\mathbf{v}_1^f\|^2 + \frac{1}{2}m\|\mathbf{v}_2^f\|^2 \end{aligned} \tag{A1}$$