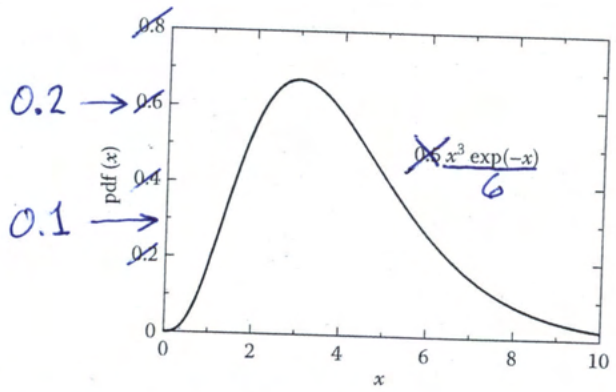

[Not normalized previously. Vertical scale/prefactor too large by factor of 3. Thanks to Huwen Kong for catching this.]

FIGURE 2.8 An asymmetric distribution. Specifically, a gamma distribution is shown. In contrast to the Gaussian distribution (Figure 2.2), values to the left and right of the mean are not equally probable.

2.3.5 Beyond Variance: Skewed Distributions and Higher Moments

While the second central moment (i.e., the variance) describes the width of a distribution, if you think of a distribution such as in Figure 2.8, it is clear that such a function is characterized by more than the mean and variance. (To put it more carefully, knowing only the mean and variance does not permit one to distinguish among distributions of *a priori* unknown functional forms.) Accordingly, statisticians frequently discuss higher moments. The most important are the skew, defined as μ_3/σ^3, and the strange-sounding kurtosis, $(\mu_4/\sigma^4 - 3)$, which measure how much the tails of the distribution bulge in comparison to the peak.

PROBLEM 2.25

Re-sketch Figure 2.8 for yourself and explain, based only on the shape of the distribution, why the mean of this distribution differs from its most probable value (the peak)? Should the mean be bigger or smaller than the peak value?

PROBLEM 2.26

What mathematical condition defines a symmetric distribution, bearing in mind that the mean need not be zero? Also, show that $\mu_3 = 0$ for symmetric distributions.

2.3.6 Error (Not Variance)

Perhaps you've heard that error decreases as the square root of the number of measurements, which is true for experimental as well as numerical data. Of course, it makes sense that error decreases with more measurement, but why is one functional form ubiquitous? We can now answer this question, and hopefully appreciate the crucial difference between variance and error.

What we are missing is what physicists would call a *scale* for the problem. You would not call a basketball player short if she were 0.002 km tall! For statistics, it is the width of a distribution which provides the natural scale—that is, the standard deviation σ.

PROBLEM 2.34

Why is the standard deviation appropriate here and not the standard error?

The Pearson correlation coefficient, r, is therefore defined using a scaled version of Equation 2.28,

$$r = \frac{\langle (x - \langle x \rangle)\,(y - \langle y \rangle) \rangle}{\sigma_x\,\sigma_y}. \tag{2.29}$$

Because the natural scales for each variable have been built into the denominator, the correlation coefficient can be compared to 1. That is, $|r| \ll 1$ implies no (or minimal) linear correlation, whereas r close to positive or negative one implies strong correlation.

Linear regression provides a highly related analysis that can be found in any statistics textbook, and will not be discussed here.

2.4.5 More Complex Correlation

As you were warned above, correlations are far from a simple business. While everything looks pretty for linearly correlated data or for uncorrelated data, it is clear from Figure 2.10b that more complex situations can arise. Note that the parabolic shape of the nonlinearly correlated data guarantees that the Pearson correlation coefficient of Equation 2.29 will be close to zero: for every value of y (and hence for every value of $\Delta y = y - \langle y \rangle$), there is a nearly equal likelihood of finding a positive or negative value of $x \simeq x - \langle x \rangle$ (since $\langle x \rangle \simeq 0$) of equal magnitude.

Things only get worse with increasingly complex systems, such as molecules that may contain hundreds or thousands of coordinates. In general, coordinates characterizing nearby atoms are correlated, and furthermore, there may be multi-dimensional correlations. In other words, any given coordinate may affect and be affected by many different coordinates and combinations thereof.

2.4.5.1 The Simplest and Most General Way of Understanding Correlations

Let's return to what we first said about correlations: they are present between two variables if "one is affected when the other is varied." What does this really mean in terms of distributions? For one thing, it means that distributions of correlated variables cannot be written in the statistically independent form of Equation 2.26. But there is a simple visual test that checks this exact criterion.

Correlations are one issue where visual inspection is often better than a simplified mathematical measure like the Pearson (linear) coefficient. That is, it is easy to tell

seem that we could determine—in principle, perhaps with a super-duper-computer—the exact future dynamics of the system. Yet, we cannot.

Every system (e.g., our test tube) is in contact with the environment, which in turn, is in contact with the whole world. Therefore, to know the exact deterministic behavior of our protein, we would need to keep track of every atom in the universe! This is something I call impossible, even in principle.

Let's take a more concrete outlook. Imagine again that our system is completely classical, but that we are somehow able to make a movie (not a simulation, but a real movie) of all atoms in our test tube—but not the test tube itself (see Figure 4.3). We would indeed see all atoms interacting classically. Yet the test tube is in contact with the air at temperature T, and therefore "test-tube atoms" have the corresponding kinetic energy. If we had included these atoms in our movie, we would see them colliding with the water in the tube. But since we don't, the water molecules at the boundaries of our system appear to be experiencing random collisions—in just exactly the right way to maintain the temperature as constant. This is the random or "stochastic" element necessary for almost any realistic description of a constant-and-finite-temperature system.

The logic just described will tell you that if you look at a protein without solvent, or with just a small amount, at least the system's boundaries must be described stochastically. That is, random collisions must somehow be simulated to maintain the temperature. This is a common approach in many simulations.

We'll see more reasons in Section 4.9 why exact trajectories for molecular systems cannot be computed—that is, simulated on a computer—even in principle.

As a technical aside, note that in general, kinetic and potential energy can be exchanged, even in a Newtonian system where the total energy is maintained as constant. In a large system, then, the kinetic energy of a system may remain relatively constant even if it is isolated from a "thermal bath" of random collisions, but in a small

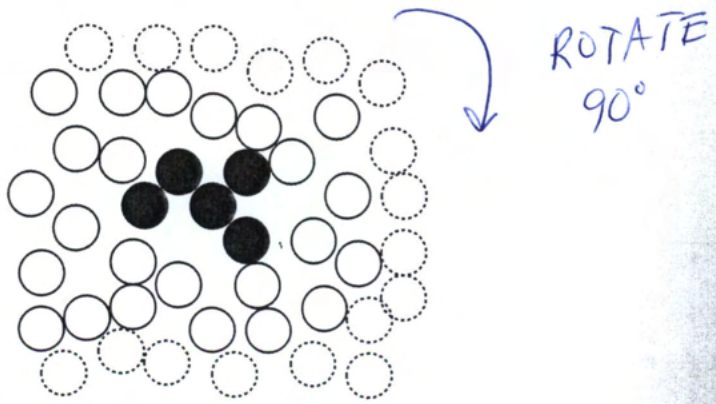

FIGURE 4.3 There is no such thing as a deterministic finite system because no system can truly be isolated. The cartoon indicates a biomolecule (filled circles) in a solvent (open circles) contained in a test tube made of atoms (dashed circles). The biomolecule interacts with the solvent, which interacts with the test tube atoms, which interact with the air.... In a biological cell, there are even more interactions.

$$\overset{(c)}{\simeq} \frac{2\pi L^2}{\lambda_A^2 \lambda_B^2} \int_0^\infty dr\, r \exp\left[-k(r-r_0)^2/2k_BT\right]$$

$$\overset{(d)}{\simeq} \frac{2\pi L^2}{\lambda_A^2 \lambda_B^2} r_0 \int_{-\infty}^{\infty} dr \exp\left[-k(r-r_0)^2/2k_BT\right]$$

$$\overset{(e)}{\sim} \frac{2\pi L^2}{\lambda_A^2 \lambda_B^2} r_0 \sqrt{\frac{k_BT}{k}} \propto T^{5/2}, \qquad (5.9)$$

where in the second line (b) we have performed the ϕ integration and also substituted the explicit form for U. In (b), we have assumed that all ϕ values are equivalent, but this is not quite true when the molecule is near the edge of the box. However, the large box ($L \gg r$) means that edge configurations constitute a negligible fraction of the total. In the third line (c), we performed the x^A and y^A integrations and assumed the value of the integral would be essentially unchanged if we set $r_{max} = \infty$.

How can we justify an infinite r_{max} in the partition function (5.9)? At first glance, it seems r must be limited by $L/2$ or certainly by L. But then the situation seems even worse: What if atom A is quite near the wall of the box? Indeed, in principle, the maximum r depends on the orientation (i.e., on ϕ) and could be quite small for some A locations. This is where our two earlier conditions come in—the stiffness of the spring (large k) and the small equilibrium size of the molecule ($r_0 \ll L$). The stiffness means that for any position of A, only a small range of r near r_0 is important (i.e., contributes to the integral). The smallness means that the fraction of A positions that are close to the wall (on the order of r_0 distant) is negligible. Together, these assumptions tell us we can ignore the ϕ dependence in (b), set r_{max} to infinity in line (c), and set the linear r factor in the integrand to r_0 in (d), making only very small errors.

Given all the explanation necessary to perform a reasonable calculation on even this simplest of molecular systems, you should recognize the intrinsic difficulty of performing statistical mechanics calculations for molecules. Indeed, it is fair to argue that, even for our model, we have somewhat oversimplified. The symbol "$\sim$" was used in line (e) of Equation 5.9 to emphasize that the stiff spring assumption is somewhat inconsistent with our earlier assumption that the velocities were fully independent. After all, if the spring gets infinitely stiff, at least one of the six components of the momenta (represented in suitable relative coordinates) is not fully independent. While this book relies solely on classical mechanics to underpin statistical mechanics, diatomic molecules are usually treated quantum mechanically—see the discussions in the book by McQuarrie and Simon—where the coupling between configurational and momentum degrees of freedom is handled more naturally.

5.2.3.2 Loss of Translational Entropy

The most important physical implication of the diatomic partition function (5.9) is a loss of entropy, as compared to noninteracting atoms. Further, we will see that translational entropy is lost, which is a generic feature of "tying" together two or more particles, as we did with our spring.

The calculation of the entropy proceeds similarly to the ideal case, except that now we have $\langle U \rangle = k_BT$ because there are effectively two interacting configurational

coordinates (the relative x and y positions embedded in r and ϕ). Thus, $\langle E\rangle = \langle U\rangle + \langle \mathrm{KE}\rangle = 2k_\mathrm{B}T$. On the other hand, because $L \gg \lambda$, we expect $F = -k_\mathrm{B}T \ln Z \gg 2k_\mathrm{B}T$. Therefore, it is again a reasonable approximation to set

$$S \simeq k_\mathrm{B} \ln Z = k_\mathrm{B}\left[\ln\left(\frac{L}{\lambda_\mathrm{A}}\right)^2 + \ln\left(\frac{r_0\sqrt{k_\mathrm{B}T/k}}{\lambda_\mathrm{B}^2}\right)\right] \simeq k_\mathrm{B} \ln\left(\frac{L}{\lambda_\mathrm{A}}\right)^2, \tag{5.10}$$

where we used our stiff-spring and small r_0 assumptions. (Note that in the full entropy, both λ_A and λ_B correctly appear in a symmetric way.)

The diatomic entropy (5.10) is only marginally larger than the *single*-atom entropy (Equation 5.3) and much smaller than that for two noninteracting atoms (Equation 5.6). Physically, this makes sense. Because the two atoms are tied together, they really can only explore the configuration space available to one of the atoms. They cannot perform "translational" motion independently, and hence the entropy loss is specified in this way.

5.2.4 Lessons Learned in Two Dimensions

Later, you'll see that the 2D calculations we just performed are excellent models for more advanced calculations. Indeed, we have seen both types of multidimensional systems: a multi"molecular" system (two atoms) and a nontrivial diatomic "molecular" system. Four things are clear already: (1) noninteracting systems are easy, (2) internal coordinates make a lot of sense for describing the configuration of a molecule, and (3) molecular calculations are tricky in general, and with internal coordinates, one particular subtlety is the Jacobian factors needed in configurational integrals. Furthermore, (4) the loss of translational entropy occurring in our simple diatom qualitatively indicates what happens in protein–ligand binding.

5.3 COORDINATES AND FORCEFIELDS

5.3.1 Cartesian Coordinates

"Cartesian" coordinates are just a fancy way of calling the x, y, z coordinates you are already familiar with. The only generalization from what you may be used to is that we can have many more than three coordinates defining a single point configuration space. After all, to specify the configuration of a molecule with N atoms requires $3N$ numbers—x, y, and z for each atom. We'll often denote the full set of coordinates by $\mathbf{r}^N$, which is meant to remind you that there is a long list of numbers, something like in exponentiation (though N here is not an exponent). This notation builds on the common use of $\mathbf{r}$ for the vector describing the location of a point in 3D.

To put it concretely, we have

$$\begin{aligned}\mathbf{r}^N &= (\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)\\ &= (x_1, y_1, z_1,\ x_2, y_2, z_2,\ \ldots,\ x_N, y_N, z_N),\end{aligned} \tag{5.11}$$

where $\mathbf{r}_i$ denotes the location of the atom i and similarly for x_i, y_i, z_i.

where $d\mathbf{r}^{N_X}$ indicates integration over all Cartesian coordinates of all atoms of type X and $(\mathbf{r}^{N_A}, \mathbf{r}^{N_B})$ represents the coordinates of all atoms of both types.

5.5.2 Ideal Systems—Uncorrelated by Definition

While (nearly) ideal systems, where the constituent molecules are uncorrelated, are common enough in everyday life (e.g., the air), such ideality is essentially nonexistent in biology. Everything is correlated with everything else, to a good approximation. Nevertheless, ideal gases form the key example in connecting statistical mechanics to thermodynamics, as we will see. It is therefore worth the effort to understand ideal systems—and fortunately, these are the easiest systems around.

If we consider first a monatomic gas (one of the noble elements, say), and we say the atoms are uncorrelated, then no configuration is favored (energetically) over any other. That is, the potential energy is a constant, which we can safely consider to be zero: $U(\mathbf{r}^N) = 0$. We can then evaluate the partition function exactly, since the Boltzmann factor is always one. Thus, Equation 5.20 becomes

$$N \text{ identical ideal atoms:} \quad Z = \frac{\lambda^{-3N}}{N!} V^N, \tag{5.22}$$

where V is the volume of the system. If the system consists of N ideal molecules, each containing n atoms, the partition function becomes

$$N \text{ identical molecules:} \quad Z = \frac{\lambda^{-3N}}{N!} \left(q^{\text{mol}} V\right)^N, \tag{5.23}$$

where q^{mol} is the partition function of one molecule and λ is the whole-molecule thermal wavelength.

PROBLEM 5.12

Starting with the partition function for $N \cdot n$ atoms of a system N ideal molecules, derive the partition function (5.23).

These ideal systems will be essential for understanding thermodynamics in Chapter 7, and therefore for our general understanding of biophysics.

5.5.3 Nonideal Systems

In the real world, there are no ideal systems, since atoms and molecules always exert forces on one another. In biological cells, which are crowded with solvent and macromolecules, this is certainly true. Here, we will very briefly describe some key qualitative aspects of typical nonideal systems of interest.

5.5.3.1 Solute in Solvent

A common type of system with two molecule types is when solute molecules are dissolved in a solvent. If you think of the forcefield equation (5.12), in this case, there are three basic types of interactions: between pairs of solvent molecules, between pairs

probability (i.e., "original" Boltzmann factors for the full set of coordinates $\mathbf{r}^N$) that are consistent with a given R. That is, we want to sum up Boltzmann factors for all configurations such that $\hat{R}(\mathbf{r}^N) = R$ and do this for all R to get the full distribution. Mathematically, this is equivalent to

$$\begin{aligned} e^{-\mathrm{PMF}(R)/k_\mathrm{B}T} &\propto \rho(R) \\ &\propto \int d\mathbf{r}^N \, \delta\left(R - \hat{R}(\mathbf{r}^N)\right) \, e^{-U(\mathbf{r}^N)/k_\mathrm{B}T}. \end{aligned} \tag{6.8}$$

You can readily extend this type of formulation to PMFs of multidimensional subsets of general coordinates.

The Equation 6.8 looks different from the preceding Equations 6.5 through 6.7 because all variables are integrated over in the present instance. However, the two formulations are equivalent if you consider the special properties of a delta function (Section 2.2.8).

PROBLEM 6.2

Show the equivalence of Equations 6.8 and 6.5—when projecting from three dimensions to one—by setting $\hat{R}(x, y, z) = x$ and integrating Equation 6.8 over x.

PROBLEM 6.3

Consider the two-dimensional potential $U(x,y) = U_1(x) + U_2(x, y)$, where $U_1(x) = E_\mathrm{b}(x^2 - 1)^2$ and $U_2(x, y) = [\kappa(x)/2]y^2$ with $\kappa(x) = \kappa_0 + \kappa_1 x^2$. This is a double-well potential in x with a harmonic potential in y whose width depends on x. Assume that at the temperature, T, of interest, the constants are equal to $E_\mathrm{b} = 5k_\mathrm{B}T$, $\kappa_0 = k_\mathrm{B}T$, and $\kappa_1 = 3k_\mathrm{B}T$. For simplicity, x and y are dimensionless in this problem.

a. Roughly sketch a contour plot of the potential.
b. Given that the potential is harmonic in y, what is the (x-dependent) standard deviation of the corresponding Gaussian Boltzmann factor.
c. Use the result from (b) to calculate the necessary Gaussian integral and determine PMF(x) up to a constant.
d. Plot U_1 and the PMF on the same graph and compare their barrier heights.
e. Explain what differs between the two curves and why.

6.2.2 Proportionality Functions for PMFs

What we have described already about PMFs embodies the essential physical understanding you should have. As Equation 6.4 indicates, simply enough, the Boltzmann factor of the PMF gives the probability distribution for the coordinates it includes. The distribution will be correct if the PMF is properly computed using standard statistical projection, as illustrated in Equations 6.5 through 6.7. Again, if you ignore constants of proportionality, you will still get PMFs that yield the right relative probabilities and are sufficient for most purposes.

PROBLEM 7.1

Show that Equation 7.8 holds for systems of arbitrary dimensionality.

In terms of notation, we will find ourselves writing Boltzmann factors so often that it is convenient to define the inverse thermal energy $\beta = 1/k_B T$ (sometimes referred to as the inverse temperature by physicists who are cavalier about units). Then the Boltzmann factor is $e^{-\beta E}$, and in fact, Equation 7.8 simplifies to

$$\langle E \rangle = -\frac{\partial}{\partial \beta} \ln Z = +\frac{\partial}{\partial \beta} \ln \beta F \quad \left(\beta = \frac{1}{k_B T}\right). \tag{7.9}$$

PROBLEM 7.2

Derive Equation 7.9.

It is often important to consider only the potential energy, and it should be fairly obvious from the derivations above that the average U can be written in terms of a derivative of the configuration integral:

$$\langle U \rangle = -\frac{\partial}{\partial \beta} \ln \left(\frac{\hat{Z}}{l_0}\right), \tag{7.10}$$

where the configurational partition function is $\hat{Z} = \int dx\, e^{-U(x)/k_B T}$ in one dimension, and l_0 is an arbitrary constant length to keep our math straight. The relation (7.10) holds for systems of arbitrary dimensionality with $l_0 \to l_0^{3N}$ for N atoms.

Later, when we study other thermodynamic conditions (e.g., constant pressure instead of constant volume as we have been doing implicitly), we'll see that other first derivatives of the appropriate free energy lead to other averages.

7.2.2.1 The Entropy

Now that we've seen the average energy extracted from the free energy by differentiation, and since we know that F is obtained from the log of Z, you may not be surprised to learn that $S = (\langle E \rangle - F)/T$ can also be calculated from differentiation. In fact, simply using the product rule, one obtains

$$\frac{\partial F}{\partial T} = \frac{\partial}{\partial T}(-k_B T \ln Z) = -k_B \ln Z - k_B T \frac{\partial}{\partial T} \ln Z. \tag{7.11}$$

The first term is F/T by definition, and the second is $-\langle E \rangle/T$ from Equation 7.8, thus implying

$$S = -\frac{\partial F}{\partial T}. \tag{7.12}$$

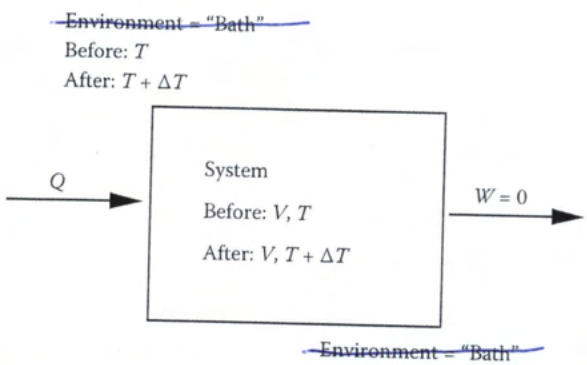

FIGURE 7.2 Heating a gas at constant volume, a process in which no mechanical work is done. Heat is added as the system temperature increases by ΔT.

thermodynamically. Thus, we will show that $c_v = Q/\Delta T$ really does describe both heat exchange and energy fluctuations in the ideal gas.

So we heat our ideal gas from temperature T to $T + \Delta T$, keeping it in a fixed-size box (constant volume). The heat added to raise the temperature is Q. Is work done? Since our gas particles don't interact with anything (except occasionally the walls of the enclosing box), the only way for them to do work is mechanically by changing the volume of the box. However, the volume is constant in this case, so no work is done ($W = 0$). Thus, from the first law, $Q = \Delta\langle E\rangle$, or from Equation 7.18, $Q = (3/2)Nk_B\Delta T$.

In other words, based on purely statistical results, we find that $Q/\Delta T = (3/2)Nk_B$ for any ΔT, in the ideal gas. (You should not make the mistake of assuming the specific heat is constant—independent of T—for any nonideal system.) But what if we use the thermodynamical definition of specific heat, namely, $c_v \equiv T\partial S/\partial T$? Well, you can differentiate our result (7.20) above for the entropy yourself and find exact agreement with the statistical result.

We have therefore shown that, at least in the ideal gas, the traditional meaning of the specific heat—the heat needed to raise the temperature—just happens to be proportional to the variance of the energy. This is an important result that you should think about.

PROBLEM 7.5

Show that you can "weigh" an ideal gas (determine the total mass) using a calorimeter, which is a machine that can simultaneously measure Q and ΔT. Assume you know the mass of each particle (atom) ahead of time, but not the number of particles.

7.4.2 Why Is It Called "Free" Energy, Anyway? The Ideal Gas Tells All

We can start with the punch line: "Free" means available. We can even give away the secret behind the punch line: a constant-temperature system exchanges energy

7.6.2 Free Energies Are "State Functions"

As we have seen more than enough times, a free energy is defined from a partition function. In turn, the partition function simply integrates to some number once the natural variables—for example, T, V, N in $Z(T, V, N)$—are fixed. Thus, the natural variables define the free energy: once you specify the variables, the free energy is some number.

A free energy is called a "state function" because it comes from a definite integral depending only on the equilibrium "state" of the system—that is, only on the natural variables. It does not depend on how that state was reached (maybe by a weird nonequilibrium process) because this doesn't change the equilibrium integral.

The notion of a state function is important in understanding free energy differences generally, and especially in understanding "thermodynamic cycles," which we shall study in the context of binding in Chapter 9.

7.6.3 First Derivatives of Free Energies Yield Averages

One of the most important ideas of this chapter is that derivatives of free energies yield statistical averages. This is a mathematical consequence of the partition-function definitions of the free energies. However, the derivative relations typically make some intuitive sense—for example, the way a free energy changes with a volume change reflects the pressure.

The relations we discussed are these:

Constant T, V, N	Constant T, P, N
$\langle E \rangle = -T^2\, \partial(F/T)/\partial T$	$H = \langle E \rangle + P\langle V \rangle = -T^2 \partial G/\partial T$
$S = -\partial F/\partial T$	$S = -\partial G/\partial T$
$P = -\partial F/\partial V$	$\langle V \rangle = \partial G/\partial P$
$\mu = \partial F/\partial N$	$\mu = \partial G/\partial N$
Constant T, V, μ: $\langle N \rangle = -\partial \Phi/\partial \mu$	

Recall that several of these are only valid in the thermodynamic limit. The partial derivatives mean that the other natural variables (other than the one being varied with the "∂" symbol) are held constant.

7.6.4 Second Derivatives Yield Fluctuations/Susceptibilities

Another consequence of the fundamental connection between thermodynamics and statistical mechanics is that second derivatives yield variances σ^2, in the traditional statistical sense. The relation of most interest to us arises in the specific heat

- $c_v = T\partial S/\partial T = \sigma_E^2/(k_B T)^2)$

You could derive a similar relation for the constant pressure specific heat, and analogous relations for the fluctuations (variances) in the volume and number of particles. Of special physical significance, volume fluctuations are proportional to

8.5.1 Basics of Dielectric Behavior

Dielectric behavior results from two things: First, dielectric phenomena reflect a choice to focus on a subset of charges of interest—for example, those in a protein—while ignoring other charges—for example, those of the water molecules. Thus, one makes a somewhat artificial separation between the charges of interest $(q_1, q_2, \ldots)$ and those in the rest of the system that is now considered a background "medium." Second, and here there is some real physics, the medium should be polarizable. That is, the medium should be able to change the arrangement of its charges—its charge distribution—in reaction to the charges of interest $q_1, q_2, \ldots$. Typically, a medium will consist of molecules that individually can change their dipole vectors.

The first point above may sound like semantics, but it is fundamental. It is very possible to represent the electrostatics of charges in a dielectric medium without using a dielectric constant. One does so simply by summing over all forces, including those from charges in the medium $q_1', q_2', \ldots$—which could be the partial charges on water molecules. If there are just two charges of interest, q_1 and q_2 and an additional M charges in the medium, then the net force acting on q_1 has two parts given by

$$\vec{f}_1 = k\frac{q_1 q_2}{r^2} + \sum_{i=1}^{M} k\frac{q_1 q_i'}{r_{1i}^2}, \tag{8.6}$$

where r_{1i} is the distance between q_1 and q_i'. By comparison to Equation 8.5, we can see that $\epsilon = 1$ here, but there should be no mystery about this. The basic rule of physics is that electrostatics is completely described by the Poisson force law with $\epsilon = 1$ so long as all charges are included. Dielectric theory approximates all of the effects of the sum in Equation 8.6 into a parameter $\epsilon > 1$ in Equation 8.5.

8.5.1.1 One or the Other

We have an either/or situation: you must use one description or the other, but do not mix them. That is, either you focus on a subset of charges of interest and represent the effects of all other charges by the single dielectric constant ϵ—as in Equation 8.5—or you sum over all electrostatic forces to get the exact, complicated answer—as in Equation 8.6 (see Figure 8.5). Do not sum over all charges and use a dielectric constant other than $\epsilon = 1$.

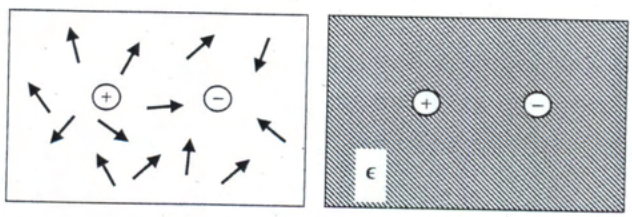

FIGURE 8.5 The two mutually exclusive descriptions of a dielectric system. The dielectric constant approximates the effects of all the dipoles (arrows) in the explicit system at left.

binding. Also, the details and sign of the change may change from system to system.

- Ligand enthalpy change, ΔH_L. A discussion analogous to that for ΔH_R applies here.
- Solvent enthalpy change, ΔH_S. What about interactions "internal" to the solvent? Although water molecules may be stripped from the surfaces of receptor and ligand, such surface waters tend to maximize their hydrogen bonds, as we discussed in the context of the hydrophobic effect (Chapter 8). Therefore, a minimal change is expected in the enthalpy of water. On the other hand, the solvent also contains salt, which might undergo more significant changes as the result of the possible burying of surface charges or polar groups upon complexation. Because such changes will be highly system specific and complicated, we will not discuss them here.
- Receptor entropy change, ΔS_R. We really need to speak of "internal" and "external" entropy. Internal entropy refers here to the change in the fluctuations of the receptor atoms—that is, the entropy change if the receptor were considered the whole system. Typically, one expects complexation with a ligand to make the receptor somewhat stiffer and decrease the internal entropy—but this is not required in principle. External entropy refers to the overall translational entropy of the receptor molecule, which decreases by definition upon complexation. Two molecules moving separately have more entropy than the same two moving together (see Section 5.2.3). The change in external entropy clearly is shared by the ligand and so part of the entropy loss can be assigned to each.
- Ligand entropy change, ΔS_L. Ligands typically exhibit some flexibility in solution (due to rotatable bonds) that will be suppressed on binding. Thus, ΔS_L generally will oppose binding. This has interesting consequences for drug design. A ligand that is designed to be as "stiff" as possible in solution (few rotatable bonds) will suffer a minimal ΔS_L penalty on binding. Nevertheless, the ligand always loses some translational entropy.
- Solvent entropy change, ΔS_S. Because water molecules will be released from the surfaces that get buried in the complex, one can expect the entropy of those solvent molecules to increase. As we discussed in the context of the hydrophobic effect, the entropy of a water molecule in bulk solution is expected to be substantially higher than a surface water.

9.3.3 Entropy–Enthalpy Compensation

The idea that entropy and enthalpy "compensate" each other—to yield an overall modest binding affinity—is quite important in biological binding. In a living cell, after all, not only must binding occur, but so too must unbinding. In signaling processes, for example, there are often feedback loops to slow down a process after it has produced the desired effect. If binding were not readily reversible, it would be hard to regulate cellular processes precisely.

The overall picture of binding, as we have seen, is complex. Further, the entropy–enthalpy compensation should be expected to apply to the net or overall binding

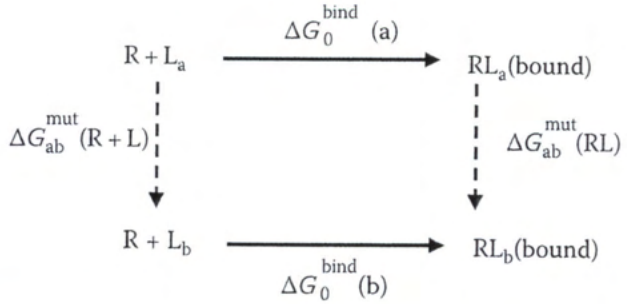

FIGURE 9.5 A thermodynamic cycle that explains relative binding affinities ($\Delta\Delta G^{\text{bind}}$) of two ligands L_a and L_b to one receptor R. Each corner of the cycle corresponds to a well-defined thermodynamic (equilibrium) state. For instance, the upper-left corner indicates the state composed of all unbound configurations of the receptor R and the ligand L_a at fixed reference concentrations. Other corners are defined similarly. Note that the horizontal arrows connote measurable binding affinities ΔG_0^{bind}. This cycle was employed by Tembe and McCammon for the calculation of relative binding affinities using "alchemical" methods.

proteins—that is, if the "ligands" are themselves proteins. Indeed R, L_a, and L_b could be three arbitrary molecules, as will be shown in the problems below.

PROBLEM 9.4

Draw a thermodynamic cycle embodying the binding of a single ligand to both a "wild-type" receptor and to a mutated receptor.

PROBLEM 9.5

Draw a thermodynamic cycle representing the relative ~~solubility~~ solvation free energy of two different molecules. ~~Solubility is the free energy~~ Solvation refers to the difference between a molecule in water and in vacuum ("gas phase").

The key quantitative point of a thermodynamic cycle is that once the conditions are specified, the absolute free energy G of each state (corner of the cycle) is fully determined. This is the state function idea of Section 7.6.2. Because G depends only on the state, the free energy differences ΔG must sum to zero going around the cycle. This sum will depend on the way the ΔG values are defined—that is, by the direction of the arrow between two states. The convention is that the arrow points from the initial state to the final state, so the free energy difference is that of the final state less that of the initial: $\Delta G(\text{state X to state Y}) = G(\text{Y}) - G(\text{X})$. Although the vertical arrows correspond to mutations (ΔG^{mut}) that could not be enacted in an experiment, such "alchemical" processes will turn out to be important in computations. Regardless of this, the free energy differences ΔG^{mut} are well defined—based on simple subtraction—since the absolute free energy is defined for each corner.

For the cycle of Figure 9.5, we therefore have

$$\Delta G_0^{\text{bind}}(\text{a}) + \Delta G_{ab}^{\text{mut}}(\text{bound}) - \Delta G_0^{\text{bind}}(\text{b}) - \Delta G_{ab}^{\text{mut}}(\text{free}) = 0, \qquad (9.10)$$