


SINGLE MOLECULES IN ACTION

What does it feel like to be a molecule? How does a molecule “behave?” That is the kind of question we often ask in biochemistry classes (and in research too), not quite in such starkly anthropomorphic terms, but nevertheless in essence. For example, we might ask: how does the GCN4 transcription factor protein recognize and bind to its specific sequence? How does hexokinase sequester glucose in an anhydrous cavity to do phosphate chemistry? How does light-excited rhodopsin activate the G-protein transducin? Etcetera, etcetera.

These kinds of questions are almost always approached experimentally by observing the collective behavior of a very large number of identical molecules; a typical enzymatic assay in a test tube, for example, is carried out on, say, 0.1 mL of 1 nM enzyme — about 60 billion molecules. You observe the formation of the product of this enzyme reaction, an “initial rate” of product formation — a smoothly increasing amount of product formed as a function of time after adding those 60 billion molecules of enzyme to the test tube. This rate depends on substrate concentration, all kinds of external variables, and, of course, the amount of enzyme added. The rate is directly proportional to the number of enzyme molecules in the test tube. A typical experiment would look like this (Figure on left):



Likewise, suppose we want to know how tightly calmodulin, a Ca^{++} -binding protein, holds onto its metal ion. We could suddenly put 10^{12} molecules of the Ca^{++} -saturated protein into a Ca^{++} -free solution and “watch” the Ca^{++} dissociate from the sites on the protein. (Think, but don’t worry, about how to do this!) What would we see? Calmodulin has four Ca^{++} -binding sites, and we might see that some of the ions come off the protein quickly, and some come off more slowly. We might observe smooth, multi-exponential “dissociation kinetics” that look like the figure on the right above.


But what we do not usually think about is the fact that this “deterministic” behavior — the precisely linear increase of enzyme product formation or the multi-exponential dissociation of Ca^{++} — is **an artifact of averaging!** This nice smoothness of a “good experiment” reflects the fact that we are observing the *average behavior of many molecules*. If we could actually follow around a given individual molecule, its behavior would look very different than the familiar average behavior of the whole collection. The behavior of the individual molecule is what’s “really” going on in that test tube; what we observe in a test-tube experiment is a mere macroscopic shadow of that underlying single-molecule behavior. If you knew what the individual behavior was like, you would immediately know what the macroscopic behavior should be. But it doesn’t work the other way around: some phenomena that appear in the individual molecule’s behavior get “washed out” in the averaging that occurs in a macroscopic experiment. So if we could directly observe single macromolecules functioning, we could derive more information about their mechanisms — about how they work — than we can from the macroscopic, ensemble-averaged, test-tube experiments that we are all familiar with.

The behavior of single molecules is not just a theoretical abstraction. There are currently a variety of different types of macromolecules -- transmembrane ion channels, DNAs, transcription factors, and various kinds of enzymes -- for which postdocs, students, and even professors (!) can directly observe individual molecules in action with suitable experimental techniques. The rules that govern the observed behavior of these molecules are general principles that to all of the kinds of macromolecules we’ll be studying throughout the course. I think that it’s best to address this single-molecule behavior first — before we tackle more conventional topics like thermodynamics, structure, or folding of macromolecules — to emphasize that this is what is “really” going on at the molecular level. This section of the course is devoted to understanding the basics of how such single macromolecules behave.

The big difference between single-molecule and macroscopic behavior is the difference between randomness and determinism. Individual molecules are constantly in motion; they simply crash around randomly, buffeted by thermal agitation, the Great Dice-thrower of statistical mechanics.

Statistics of single channel lifetimes

We are not going to discuss experimental aspects of single ion channels here. Instead, you should just know that it's possible to detect the "openings" of individual channel molecules. In other words, these proteins have the property that they undergo a conformational change from a "closed" to an "open" state, and that it is experimentally possible to observe these two conformations directly at the single-molecule level. The question we want to wrestle with here is: what kind of behavior should we expect for a protein undergoing conformational changes between two such states:




In particular, we ask: suppose that at some moment (call it zero time) the channel opens, i.e., enters state B. How long does it stay there before closing? What is the "lifetime" of the

open state? Now, since we are talking about a single molecule, you probably realize that there is no simple answer to this question, because there is an element of randomness in the single protein's behavior. Suppose we actually watched the channel open at $t=0$, and then recorded the time at which the channel closes again. Let's say the lifetime we record is 3.2 seconds. Now, let's wait until the channel opens again and *repeat* the measurement. This time, you can imagine that it would be very unlikely that we get the same lifetime. Instead, we might record a lifetime of only 0.6 seconds. Do it again: 2.9 sec. Again: 10.7 sec. Again: 0.2 sec. This can get boring, but in order to get a good idea of the channel's open-state lifetime, you will have to measure this many times because the behavior at the single-molecule level is *intrinsically statistical*.

A familiar analogy to consider is your own homely experience with light-bulbs. Suppose you go to the hardware store and buy a thousand brand-new light bulbs, and then you do an experiment. You screw one in, and start a stopwatch at the moment you flip on the light. Now you wait until the bulb burns out, and then stop the stopwatch and record the bulb's "lifetime." Let's say it's 342 hours. (You can see that this is going to be an easy but boring experiment.) Then screw in the next light bulb and do it again: 557 hours. Again: 35 hr (a bad bulb in the lot). Again: 831 hr. The lifetimes are all different, but are you surprised? Of course not -- you know that there will be some variation among the individual light bulbs. After measuring all 1000 lifetimes, you will have some kind of answer to the question: "how long is the bulb's lifetime?" But you'll realize that the answer is *intrinsically statistical*, just as with the single channel example. You can sensibly ask: what is the average lifetime of the bulb? You can ask: how much variation around the average does this lot of bulbs give? But the question: "what is the lifetime of the light bulb?" is not so meaningful (except if you are asking about the measured lifetime of a particular bulb that just burned out). What you really want to know is the statistical distribution of lifetimes. That is, what fraction of the bulbs have lifetimes between 0 and 10 hr (not many, you can imagine), between 10-20 hr, etc. Your own experience with light bulbs tells you that you might see a "bell-shaped" distribution like the one shown below, where a substantial fraction of bulbs have lifetimes in the range 400-700 hours, not many in the range 0-100 hours, and not many in the range 900-1000 hours. It turns out that this statistical distribution is the right way of asking the question "how long does the light bulb stay on?"

This statistical distribution has a name -- the "probability density function," and we have to talk about it a little more. This "pdf" is graphed below for the light bulb example. The pdf is a very powerful thing, because it contains all of the extractable information about the light bulb



lifetimes. We are now going to define this function mathematically. Let's ask the simple question: "what fraction of lifetimes falls in the range 500-600 hr?" i.e., somewhere near the peak of the distribution. In our example, let's pretend that this fraction is 0.27, 27% of all measured lifetimes being in this range of values. This fraction is actually a formal probability -- the probability of finding a lifetime, τ , in the range of 500-600 hr:

$$\text{Prob}\{\tau \in (500,600)\}, \text{ or } \text{Prob}\{\tau \in (500, 500+\Delta t)\}, \text{ where } \Delta t=100 \text{ hr.}$$

The pdf, written $p(t)$, is a function of t , and is defined in terms of the formal probability above:

$$(1) \quad p(t) = \frac{\text{Prob}\{\tau \in (t, t+\Delta t)\}}{\Delta t} \quad \text{DEFINITION!}$$

Think about this equation and try to translate it into the light bulb example above. There is a TREMENDOUS thing that you have to realize about the pdf: it is not -- repeat not -- a probability; look at the definition -- it has units of hr^{-1} , whereas probability is a pure dimensionless number between zero and one. It's best to think of the pdf here as a probability per unit time (otherwise known to kineticists as a rate constant!). In our example above,

$$p(500 \text{ hr}) = \text{Prob}\{\tau \in (500,600)\}/100 = 0.27/100 = 0.0027 \text{ hr}^{-1}.$$

You can restate this equation in words as follows: *The "probability of burnout per time" of a light bulb that has already been on for 500 hours is 0.0027 per hour -- i.e., there's a 2.7% chance that it will burn out within the next 10 hours, or about a 27% chance that it will burn out in the next 100 hours.* (Obviously you have to keep the Δt time interval small; it would be crazy to say that the bulb has a 270% chance of burnout in the next 1000 hours!) Try restating the above sentence with the phrase "rate constant of burnout" substituted for "probability of burnout per unit time." See how "natural" that sounds?

Here's another very big thing to realize. The actual probability of finding a lifetime in the interval $(t, t+\Delta t)$ depends strongly on the choice of Δt , the arbitrarily chosen "interval width." (That is, the probability of finding a lifetime in the range 500-600 hr is going to be ~twice as big as finding one in the range 500-550 hr, right?) But the probability density, $p(t)$, does not depend on the interval! The " Δt " factor has been sort of "normalized out" of the probability to give you the probability density $p(t)$. The pdf is an intrinsic property of the lot of light bulbs -- it characterizes the lifetime distribution of that particular lot of light bulbs. If you go to the store and buy a light bulb, you should ask the clerk to tell you the $p(t)$ for that lot of bulbs, because then you could answer for yourself the central question: "what's the probability that this particular light bulb I just bought will have a lifetime between 610 and 675 hours?" The answer to that question is:

$\text{Prob}\{\tau \in (610, 610+65)\} = p(610) \times 65.$ (But you can do this calculation only if the clerk has told you $p(t)$! That's why $p(t)$ is so important to know.)

If you know $p(t)$, you can also answer the important question: what's the probability that this light bulb I just bought will burn out before 1000 hours? The answer is:

$$\text{Prob}\{\tau \leq 1000\} = \int_0^{1000} p(t) dt$$

OR, you might want to know the probability that a bulb will last at least 1000 ho

$$\text{Prob}\{\tau > 1000\} = \int_{1000}^{\infty} p(t) dt$$

(Think about this! Derive these results for yourself, given the definition of $p(t)$, eq 1!)

The point I'm trying to make is that the probability density function is a fundamental way to describe the behavior of something that varies randomly, rather than deterministically. If you know $p(t)$, then you know a great deal about the behavior of the system. (Example: Remember your first-semester P-chem and think of the Maxwell-Boltzmann distribution of molecular speeds of a ideal gas. That is a pdf -- a function that tells you, at a given temperature, what fraction of the molecules are whizzing around rapidly, slowly, etc. From that pdf, you can calculate the average molecular speed as a function of temperature, for example.)

What works for light bulbs and for ideal gas molecules also works for single macromolecules. If we could know what $p(t)$ is, then we could find out many things about the random behavior of these macromolecules. But how do you find out $p(t)$ for a given situation? In the case of the light-bulbs, the sales-clerk tells you -- it was measured at the factory. But for single channels undergoing open-closed conformational transitions, there is no sales clerk -- we will have to derive an expression for the probability density function of lifetimes in the open state and in the closed state. How do we do this? We must first make up a mental picture -- a model - - of what we mean by the statement: "the channel opens at $t=0$ and then at some random time later, it closes." This is the heart of the problem: to think up a model of randomness that applies to this particular situation. We will do this problem 2 ways -- one that is rather formal (and traditional -- i.e., in the textbooks), and one that is a little easier to imagine mechanistically at the single molecule level.

Derivation #1: Open-state lifetimes for a single channel

Suppose the channel enters the open state at $t=0$. We ask: what is the probability that the channel will still be open at some time later, i.e., that it has not yet closed at time t ? Let us call this probability $S(t)$. (Here, "S" stands for the probability of "survival," or, if you like, "still open." I will call this the "survival function.") We need to think up an equation containing $S(t)$. Such an equation will not just pop out of thin air -- we have to think up a picture of randomness and then express it in words, and then translate this into an equation. Here is a simple model of randomness that might plausibly apply to this situation. Let's try it out by writing down two "axioms:"

A. The probability that the channel closes in a given small time interval Δt is simply proportional to that time interval. (Think: does this make sense?)

B. The probability that the channel closes in a given time-interval is totally independent of how much time has previously passed. (This is the assumption that the channel does not "remember" anything about its past history. Think about this for a while: Why should it apply to the macromolecular situation we're considering?)

From these two assumptions, we will derive $S(t)$, and then from that we'll get $p(t)$. Translate each axiom into a mathematical expression.

$$(1A) \quad \text{Prob}\{\tau \in (t, t+\Delta t)\} = \mu \Delta t \quad \text{-- where } \mu \text{ is the proportionality factor}$$

$$(1B) \quad \mu \neq \mu(t), \quad \text{-- i.e., } \mu \text{ is a constant, independent of time.}$$

Now we can ask an interesting question: What is the probability that the channel first closes at a time that is just a little bit longer than time t , i.e., what is $S(t+\Delta t)$? (Don't be upset here! I realize that -- at this point -- this is a totally unmotivated question!)

We can answer this question as follows, in English: For the channel to still be open at time $t+\Delta t$, two things must happen. First, the channel must have been open at time t , and second, the channel must stay open during the little time interval Δt . Moreover, by assumption #2, these two events are totally independent. Because of this independence, we can multiply the individual probabilities to get the joint probability of the two events:

$$(2) \quad S(t+\Delta t) = S(t) (1-\mu \Delta t).$$

Look, we've got an equation! Now solve it. Let Δt get very small, and write:

$$(3) \quad \frac{S(t+\Delta t)-S(t)}{\Delta t} = -\mu S(t)$$

In the limit of $\Delta t \rightarrow 0$, the left side of (3) is simply dS/dt by definition of the derivative,

$$(4) \quad dS/dt = -\mu S, \text{ or}$$

$$(5) \quad S(t) = \exp(-\mu t). \quad \text{BIG RESULT!}$$

This result says that the survivor function falls off exponentially. What about the pdf? It turns out that the pdf is intimately related to the survivor function for any probability distribution. In fact, the two functions contain exactly the same information -- just expressed in slightly different ways. You should prove for yourself that the relation between $p(t)$ and $S(t)$ is:

$$(6) \quad p(t) = -dS/dt.$$


Therefore, we immediately have the desired pdf:

$$(7) \quad p(t) = \mu \exp(-\mu t).$$

Here are several things to notice about this pdf expected for a single channel lifetime:

1. The probability distribution is not bell-shaped! It is skewed strongly to small values of t . It's much more likely to have short-lived channels than long-lived ones.

2. Remember that the probability density function has units of time^{-1} !



Derivation #2: Flipping coins on a single channel

Here we are going to derive directly the pdf of the single-channel lifetime, without reference to the survival function. Let the channel open at $t=0$, as before, but now we pretend that a little man sits on the channel molecule and he decides when the channel will close. He makes this decision as follows. He starts his stopwatch at $t=0$, and at every tick of the watch, he reaches into a bag that contains a huge number of white marbles and a small number of red ones. He shakes the bag, grabs a marble, looks at it, and puts it back into the bag. If the marble was white, the channel is allowed to stay open until the next “trial.” But if the marble was red, he closes the channel and records the lifetime. The little man keeps this up, making a trial at each tick of the stopwatch, until he finally comes up with a red marble and closes the channel.

You can see in this example that the channel is as “random” as you can imagine. It is stupid. It is utterly without memory of its past. At each tick of the stopwatch, there is a certain low probability that the little man will pull out a red marble and close the channel. But this probability is very small, so he has to carry out many trials before finally getting a red marble. We must now put this story into mathematical language. Let’s define some useful parameters:


k — probability on a given trial of grabbing a red marble. (Prob of closing per trial)
 j — trial number ($j = 0, 1, 2, 3, \dots$)

Let $\pi(j)$ be the probability that the channel will close on the j^{th} trial (i.e., between trials # j and $(j+1)$). If the channel closes on the j^{th} trial, this means that the little man must have grabbed a white marble on trials #0, #1, #2, ... # $(j-1)$, and then finally pulled a red marble on the j^{th} trial. (Note that the first trial is called trial #0, just for proper notation’s sake.)

Since all the trials are independent events, the probability of pulling a red marble on the j^{th} trial is:

$$(8) \quad \pi(j) = (1-k)^j k \quad (\text{i.e., } j \text{ “failures” followed by a “success”})$$

This is called a “geometric” probability distribution. Think about what it says. It is a maximum at $j=0$ (the “zeroth” trial), and it falls off by a fixed fraction, $(1-k)$ for each subsequent trial. Because k is small, $(1-k)$ is close to unity (like 0.9999999...), so $\pi(j)$ doesn’t fall off very much for each trial. But after a zillion trials, that factor $(1-k)^{\text{zillion}}$ eventually gets significantly less than unity. This function, $\pi(j)$, says that in order to close on the trial # j , the channel must first have “survived” the previous j trials, each of which has a less-than 100% chance of survival. So it’s more probable to close on an earlier trial than on a later one, even though the probability of closing is the same on each trial. Here’s what an exaggerated plot of $\pi(j)$ vs. j would look like:



In this example, we envision time as increasing in an “integer” way, by clock-ticks. In fact, the geometric distribution is the “integer” analogue of a falling exponential. We must now get this problem into proper continuous time t . Suppose that the stopwatch ticks are separated by a time interval Δt seconds. Then, we can write:

$$(9) \quad t = j \Delta t.$$

Also, let the probability of closing per second be μ . Then:

$$(10) \quad k = \mu \Delta t.$$

Most importantly, $\pi(j)$ is related to the pdf $p(t)$, by:

$$\pi(j) = p(t) \Delta t \quad (\text{Make sure that you understand why!})$$

Now, substitute (9) and (10) into (8) to get rid of the “trial” parameters and introduce continuous time parameters:

$$(11) \quad \pi(j) = p(t) \Delta t = (1 - \mu \Delta t)^{j/\Delta t} \mu \Delta t, \text{ where we must take the limit } \Delta t \rightarrow 0.$$

$$(12) \quad p(t) = \lim_{\Delta t \rightarrow 0} \mu (1 - \mu \Delta t)^{t/\Delta t} = \mu \exp(-\mu t) \quad (\text{Make sure that you can take this limit!})$$

Thus, the pdf expected for a single channel lifetime is a falling exponential, just as we showed by the first method. The constant μ that describes the exponential is called the “Poisson parameter” of the pdf. It may also be called the “rate constant of closing,” as we shall see below.

Single-channel data

We have just been discussing the expected statistical behavior of a channel that opens at $t=0$. In practice, we can actually observe this. A single channel molecule is inserted into a small “patch” of cell membrane on the end of a glass micropipette, and the opening of the channel can be seen because of a huge amplification factor: every time the channel opens, about 10 million ions per second roar through it. This is equivalent to $\sim 10^{-11}$ Amps, an easily detected electric current. What we actually observe is not a single opening of the channel, but rather a sequence of many openings, each followed by a closing. Such a record might look like this:



What we are seeing in such a record is the dynamic equilibrium of the channel going back and forth between its two conformations: Closed \rightleftharpoons Open. This is a little different from how you did the light-bulb experiment; once the bulb has burned out, it's dead forever, and you just throw it out and screw in the next bulb. But with the channel, after it closes, it will eventually open again, and the moment of this re-opening can start a "new" experiment that starts a new lifetime-clock. Moreover, you've actually got two statistical measurements going here: one for the open state and one for the closed state. Each of these lifetimes — open lifetime and closed lifetime — will be exponentially distributed with its own Poisson parameter. In addition, the conformational equilibrium constant between open and closed (and hence the standard-state Gibbs free energy) can be measured right off the record from the fraction of time the channel spends in the open vs. the closed state. [*Think about this last statement! Make sure that you "get" it. See problem #5.*]

Connecting single-molecule behavior to macroscopic measurements

Up until now, we have always referred to the "Poisson parameter" as the constant that describes the statistics of the channel lifetime. We'll now see that this is the same as a conventional rate constant measured in a macroscopic experiment. Consider the following "macroscopic" experiment with light bulbs. We have a room with 1000 new light bulbs on the ceiling, and we turn them on all at the same time by flipping the light-switch on. Now, we measure the total light shining on the floor with a photomultiplier. Initially, there will be a "100%" level of light, but after many hours, bulbs will start burning out, and the light intensity will fall gradually, until all the bulbs have burned out. This macroscopic light-vs-time curve is actually an expression of the lifetime distribution statistics of the light bulbs. In fact, you are making a direct macroscopic measurement of the survival function, except that you are not looking at light bulbs one at a time, but rather simultaneously at the whole ensemble. Nevertheless, the light-vs-time curve that you would get (plot this out!) will just parallel the survival function $S(t)$, and hence, you can derive from this the pdf, $p(t)$, from eq (6):


$$(6) \quad p(t) = -dS/dt.$$

This is a perfectly general argument, and it applies to channels as well. Suppose you had 1000 channels in a patch of membrane, and at $t=0$ you turn them all on, and just watch as they turn off. (In this thought experiment, assume that when a channel turns off it does not re-open.) Then, if you start with a "100%" level of current at $t=0$, the current will decay exponentially with time, with a time constant that is exactly equal to the time constant of the single-channel survival function, i.e., $1/\mu$. This is the connection between single-molecule and macroscopic behavior, and it explains why it is that chemical kinetics are so filled with exponential decays with time. The exponential arises from the "memoryless" property of single-molecule lifetimes! So next time you measure a nice, smooth exponential relaxation in a macroscopic experiment, remember that it is showing you the underlying statistical, random behavior of the single molecules.

Here's an example that shows you how macroscopic kinetics connects to the underlying microscopic world. Imagine that we have a membrane with a fixed number of channels in it, and we force them all to open at $t=0$. We then watch them close, never to re-open, just as described

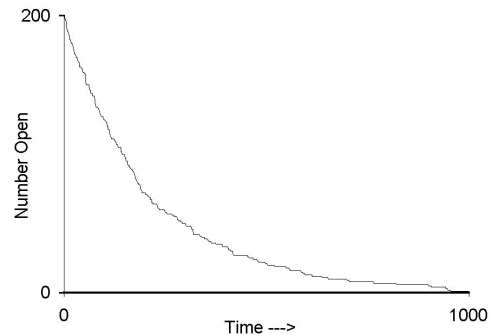

above. (Note: your neurons actually do this type of experiment all the time. In a typical synaptic transmission event, acetylcholine (Ach) gets squirted, as a 0.5 millisecond pulse, onto a membrane containing many acetylcholine receptor channels. These all open at the same time, as the pulse of Ach impinges upon the membrane. But just as quickly as the Ach appears, it disappears, as it gets hydrolyzed by acetylcholinesterase, which is found at very high concentration in the synaptic cleft. Once the Ach originally bound to the channel dissociates, the channel closes, and it never re-opens because the Ach is all eaten up within a millisecond.)

Suppose our membrane contains exactly 4 channels, each of which closes randomly, i.e., with an exponentially distributed lifetime. Let's watch these channels close with time. Such a "survival time course" will look something like this:



If you did the experiment a second time, the precise nature of the time course of the 4 channels closing would look different, because of the random nature of things, but the basic "look" of the time course would be the same.

Let's do this again with 20 channels:



Now repeat it with 200 channels (above, right). See how the time-course is getting smoother and more "deterministic" looking? Remember that this is just as random as the case with 4 channels. It just looks smoother because there are more channels to average in the time course.

If you would do this with more than 1000 channels, you would see what looks very much like a perfectly smooth exponential relaxation. You are looking directly into the face of the "survival function," which is, as you now know, the probability of the channel's still being open at time t . It is directly related to the pdf. When you observe macroscopic kinetics of any kind, you are actually looking right into the face of the stochastic nature of reality!

Single Molecule Experiments Are Not Just for Channels

The single channel recording techniques, which were invented in the mid-1970s, were for many years the only experiments in all of biochemistry – indeed, in all of chemistry – in which you could directly observe the behavior of individual molecules. The single channel experiment works because two conditions are met:

1. **Isolation** -- The surface density of channel proteins on the membrane is low, and the experiment focuses on a sufficiently small patch of membrane (perhaps $1 \mu\text{m}^2$ at the end of the micropipette) that patches which isolate only one channel molecule can readily be obtained.
2. **Detection** – The apparatus is sufficiently sensitive that the electric current flowing through a single open channel can be detected. This is facilitated by the very large numbers of ions ($\sim 10^8 \text{ sec}^{-1}$) that can flow through a single open channel.

Single molecule experiments are very useful. It would be nice to be able to do them on macromolecules other than channels, but most macromolecules are not in membranes and/or do not produce big electrical signals. Fortunately, work over the last decade has shown that it is possible to analyze one molecule at a time the behavior of almost any macromolecule. Most of these experiments are done using light microscopy. Isolation of single macromolecules is achieved by attaching them to a glass surface. If the surface density is sufficiently low that the average spacing between nearest neighbor molecules is much larger than $\sim 0.5 \mu\text{m}$ (the resolution of the microscope), single molecules can be observed without interference from adjacent molecules. The detection requirement can also be met. For example, the use of electronic detectors to improve microscope sensitivity has made it possible to visualize the fluorescence of a single dye molecule attached to an individual protein molecule.

Some examples of macromolecular reactions that have been observed in single macromolecules include:

- A. Catalytic cycling of enzymes [*Science* **282**, 1877-82 (1998)] and binding and release of substrates [*Nature* **374**, 555-559 (1995)] and inhibitors [*Biochemistry* **37**, 747-57 (1998)]
- B. Binding of proteins to DNA [*Science* **267**, 378-380 (1995)]
- C. Single conformational changes in proteins [*Nature* **368**, 113-119 (1994)]
- D. Translational [*Nature* **365**, 721-727 (1993)] and rotational [*Nature* **386**, 299-302 (1997)] movements of motor enzymes
- E. Protein folding [*Science* **276**, 1112-6 (1997)].

These are exactly the types of reactions that we will discuss in macroscopic (i.e., many-molecule) thermodynamic terms later in the course. In each study cited, the single molecules exhibit random, chaotic behavior analogous to that which we have discussed for channels, and their behavior can be analyzed using the statistical methods (pdf's and survivor functions) you have learned about. When we discuss the thermodynamics, always bear in mind that the

regularity and predictability of large populations of molecules is simply a statistical property of the completely irregular and unpredictable behaviors of the individual molecules.

Big Results on Single-molecule Statistics

Here's a little summary of the main results from considering the statistics of single-channel lifetimes, considering a very simple channel that can exist in only two conformational states, open and closed:

1. The open-state lifetimes of a single channel follow an exponential distribution,

$$p(t) = \mu \exp(-\mu t),$$

where μ is the "Poisson parameter." The "exponentialness" of the distribution is a direct expression of the memoryless character of the closing process.

2. The Poisson parameter μ is related inversely to the average channel lifetime $\langle \tau \rangle$:

$$\langle \tau \rangle = 1/\mu.$$

3. In chemical-kinetic language, the Poisson parameter is, in fact, the "rate constant" of leaving the open state, i.e., the rate constant of closing.
4. The fraction of the *time* that a single channel (at equilibrium) spends in its open state is identical to the fraction of *channels* (in a large collection) in the open state (at equilibrium). This connects the usual, conventional equilibrium constant (measured on a large ensemble of channels) to single-channel statistical behavior.