PHYS 414 Problem Set 2: Turtles and aliens

This first two problems explore the common structure of dynamical theories in statistical physics as you pass from one length and time scale to another ("turtles all the way down"). Brownian motion is an excellent model system for this: in Problem 1 we move seamlessly from the stochastic description of the Fokker-Planck equation down to classical mechanics in the form of the Liouville equation; in Problem 2 we go from Liouville to the quantum scale, and enter the strange world of the quantum phase space representation. Here probabilities become quasiprobabilities, taking on negative values, and Dirac delta functions are outlawed by the Heisenberg uncertainty principle. The last problem tackles a fundamental question: are we alone in the universe?

Problem 1: From Fokker-Planck to Liouville

The derivation of the Fokker-Planck equation in Problem 2 of the first homework set was for a particular potential energy $U(x) = \frac{1}{2}k_{\text{trap}}x^2$ due to the optical tweezers, with its corresponding trap force $-U'(x) = -k_{\text{trap}}x$. If we allowed $U(x)$ to be an arbitrary function, the same method would give us the general Fokker-Planck time evolution equation for $\mathcal{P}(x, p, t)$. For later convenience we express the distribution in terms of x and $p = Mv_x$ rather than x and v_x. The Fokker-Planck equation is:

$$\frac{\partial \mathcal{P}}{\partial t} = -\frac{1}{M}\frac{\partial}{\partial x}(p\mathcal{P}) + \frac{\partial}{\partial p}\left[\left(\Gamma p + U'(x)\right)\mathcal{P}\right] + M\Gamma k_B T \frac{\partial^2 \mathcal{P}}{\partial p^2}. \tag{1}$$

Here $\Gamma = \gamma/M$, where γ is the friction coefficient and M the mass of the Brownian particle.

a) Show that Eq. (20) can be rewritten in the following form:

$$\frac{\partial \mathcal{P}}{\partial t} = -\{\mathcal{P}, H\} + \Gamma \frac{\partial}{\partial p}\left[p\mathcal{P} + Mk_B T \frac{\partial \mathcal{P}}{\partial p}\right], \tag{2}$$

where $H(x,p) = p^2/2M + U(x)$ is the Hamiltonian of the Brownian particle and $\{A, B\}$ denotes the Poisson bracket of two functions $A(x,p)$ and $B(x,p)$:

$$\{A, B\} \equiv \frac{\partial A}{\partial x}\frac{\partial B}{\partial p} - \frac{\partial A}{\partial p}\frac{\partial B}{\partial x}. \tag{3}$$

The $\Gamma \to 0$ limit of Eq. (20) is the *Liouville equation*, describing the time evolution of a probability distribution under classical mechanics. This corresponds to diluting the gas surrounding the Brownian particle until it feels no collisions. We can get the same classical limit by looking at motion on time scales $t \ll \Gamma^{-1}$, when collisions have not yet had a substantial impact on the particle. Our Liouville equation is for a single particle of mass M moving in a potential $U(x)$, but the Liouville formulation can be easily generalized to many particles and an arbitrary Hamiltonian. Eq. (2) is nice because it shows that the stochastic time evolution of the Brownian particle is just classical mechanics plus "correction" terms proportional to Γ that lead to diffusive spreading of probability distributions.

b) Convince yourself that the Liouville equation describes classical trajectories: show that the probability distribution

$$\mathcal{P}(x,p,t) = \delta(x - x_c(t))\delta(p - p_c(t)) \tag{4}$$

is a solution to Eq. (2) when $\Gamma = 0$. Here $x_c(t)$ and $p_c(t)$ are functions of time that describe the motion of a classical particle with Hamiltonian $H(x,p) = p^2/2M + U(x)$. Hence they satisfy Hamilton's equations:

$$\frac{d}{dt}x_c(t) = \frac{\partial H}{\partial p}(x_c(t), p_c(t)), \qquad \frac{d}{dt}p_c(t) = -\frac{\partial H}{\partial x}(x_c(t), p_c(t)). \tag{5}$$

So for a distribution in phase space that starts as a delta function, $\mathcal{P}(x,p,0) = \delta(x-x_0)\delta(p-p_0)$, it will remain a delta function for all $t \geq 0$ centered at the corresponding classical trajectory. If $\Gamma \neq 0$, the delta function would broaden out over time under the diffusive effects of the gas environment. *Hint:* The following Dirac delta function properties may be useful: for any function $F(a)$ that is non-singular at $a = a_0$, we can write $F(a)\delta(a-a_0) = F(a_0)\delta(a-a_0)$ and $F(a)\delta'(a-a_0) = F(a_0)\delta'(a-a_0) - F'(a_0)\delta(a-a_0)$. Here $\delta'(a)$ is the first derivative of the Dirac delta function.

c) It is useful to compare the behavior of the mean energy $\langle H \rangle(t) = \int dx\, dp\, H(x,p)\mathcal{P}(x,p,t)$ in the stochastic ($\Gamma > 0$) versus classical ($\Gamma = 0$) regimes. Using Eq. (2) and integration by parts, show that:

$$\frac{d}{dt}\langle H \rangle = -\Gamma \int dx\, dp\, p \left[\frac{p}{M}\mathcal{P} + k_B T \frac{\partial \mathcal{P}}{\partial p} \right]. \tag{6}$$

Hence when $\Gamma = 0$, $d\langle H \rangle/dt = 0$ and the mean energy does not change with time: it is a constant of motion for classical trajectories. When $\Gamma > 0$, in general $d\langle H \rangle/dt$ does not have to be zero. The Brownian particle can gain or lose energy through collisions with the surrounding gas environment. In one special case $d\langle H \rangle/dt = 0$ even when $\Gamma > 0$: show that this is true when the Brownian particle has an equilibrium distribution $\mathcal{P}_{\text{eq}} \propto \exp(-H/k_B T)$. In equilibrium the mean energy $\langle H \rangle$ stays constant, since there is no net energy exchange with the environment on average (gain is balanced by loss).

d) Imagine that in addition to the force from the potential $U(x)$, there is an external force F_{ext} on the Brownian particle (imposed for example by an experimentalist). Find the extra term that appears on the right-hand-side of Eq. (2), and show this gives the following addition to Eq. (6):

$$\frac{d}{dt}\langle H \rangle = -\Gamma \int dx\, dp\, p \left[\frac{p}{M}\mathcal{P} + k_B T \frac{\partial \mathcal{P}}{\partial p} \right] + \frac{F_{\text{ext}}}{M} \int dx\, dp\, p\mathcal{P}. \tag{7}$$

Argue that this extra term is just the average rate of work done on the Brownian particle by the external force. Notice that this term is independent of Γ: it appears both in the classical and stochastic regime.

Problem 2: From Liouville to Quantum Mechanics

"Negative energies and probabilities should not be considered as nonsense. They are well-defined concepts mathematically, like a negative of money."

Paul Dirac, 1942

Let us now zoom in even further, studying the motion of our Brownian particle at time and length scales so small that quantum effects become important. We will assume that during these minuscule time intervals the gas particles have no time to reach and collide with the bead, so we are really just dealing with a single quantum particle of mass M in a potential $U(x)$. Ideally we would like a way of modifying our Liouville equation for the classical motion of the particle to include quantum correction terms, proportional to $\hbar$. In the classical limit, $\hbar$ could be assumed negligible compared to the distance/momentum scales of interest, and we would recover the Liouville equation.

There is one problem: in the standard formulation of quantum mechanics, we never speak of a probability density $\mathcal{P}(x, p, t)$ defined over the phase space of (x, p). There is a good reason for this: a delta function probability distribution in phase space, like the classical result in Eq. (4), would violate Heisenberg's uncertainty principle, since both x and p would be known simultaneously. The answer to this problem was developed by Weyl, Wigner, Groenewold, and Moyal through the 1930's and 1940's, and came to be known as the *phase space formulation* of quantum mechanics. It is an alternative to the two better known approaches to quantization: the standard Schrödinger-Heisenberg picture of operators in a Hilbert space, and the Feynman path integral representation. Though it is formally equivalent to both of them, it languished for many years (negative probabilities are freakish!), until recently it has been resurrected as a research tool for understanding quantum optics and the decoherence of quantum systems interacting with the environment (a major issue in quantum computing). For more historical background, there is a nice article by Thomas Curtright and Cosmas Zachos at: arxiv.org/abs/1104.5269. This problem will not do full justice to the quantum phase space picture, but it will explore some of its salient features, and the elegant relationship between the quantum and classical time evolution equations.

a) The basic tool to derive all the properties of the phase space representation is the *Wigner transformation*, a map W that converts any Hilbert space operator $\hat{A}$ in the standard picture of quantum mechanics to a corresponding scalar function $A(x, p) = W\{\hat{A}\}$ of x and p (which are the real-valued position and momentum of the particle):

$$A(x,p) = W\{\hat{A}\} \equiv 2\int_{-\infty}^{\infty} \langle x+y|\hat{A}|x-y\rangle e^{-2ipy/\hbar} dy. \tag{8}$$

Prove that $A(x, p)$ is real-valued if $\hat{A}$ is Hermitian. Also show that $A(x, p)$ can be expressed equivalently as an integral over momentum instead of position:

$$A(x,p) = W\{\hat{A}\} = 2\int_{-\infty}^{\infty} \langle p+q|\hat{A}|p-q\rangle e^{2ixq/\hbar} dq. \tag{9}$$

Hint: Standard quantum mechanics in a nutshell: $|x\rangle$ and $|p\rangle$ are eigenstates of the $\hat{x}$ and $\hat{p}$ operators respectively. In other words, $\hat{x}|x\rangle = x|x\rangle$, $\hat{p}|p\rangle = p|p\rangle$. Here are several useful properties of the eigenstates:

$$\begin{aligned}
&\langle x|p\rangle = \langle p|x\rangle^* = \frac{e^{ipx/\hbar}}{\sqrt{2\pi\hbar}}, \qquad 1 = \int dx\, |x\rangle\langle x| = \int dp\, |p\rangle\langle p| \\
&\langle x|x'\rangle = \int_{-\infty}^{\infty} dp\, \langle x|p\rangle\langle p|x'\rangle = \int_{-\infty}^{\infty} \frac{dp}{2\pi\hbar}\, e^{ip(x-x')/\hbar} = \delta(x-x'), \qquad \langle p|p'\rangle = \delta(p-p')
\end{aligned} \tag{10}$$

You may also find the following identity useful: $\delta(ax) = |a|^{-1}\delta(x)$ for any constant a.

b) For a particle described by some quantum state $|\Psi\rangle$, define a Hermitian operator $\hat{\mathcal{P}} \equiv (2\pi\hbar)^{-1}|\Psi\rangle\langle\Psi|$ (this is proportional to what we will later call the *density operator*). The Wigner transformation $\mathcal{P}(x,p) = W\{\hat{\mathcal{P}}\}$ is the central quantity in the phase space formulation, and in the classical limit corresponds to the familiar phase space probability density. However, we have to be careful, because at the quantum level $\mathcal{P}(x,p)$ is almost (but not quite) a probability distribution. Hence it is called a *quasiprobability distribution*. To understand this, let us first discuss the good news. Derive the following properties of $\mathcal{P}(x,p)$:

$$\int dp\, \mathcal{P}(x,p) = |\langle x|\Psi\rangle|^2, \qquad \int dx\, \mathcal{P}(x,p) = |\langle p|\Psi\rangle|^2, \qquad \int dx\, dp\, \mathcal{P}(x,p) = 1. \tag{11}$$

So far everything looks great: the marginal probability density of finding the particle at position x is $|\langle x|\Psi\rangle|^2$, exactly as standard quantum mechanics predicts, and similarly the marginal probability density of finding the particle with momentum p is $|\langle p|\Psi\rangle|^2$. Moreover, these two properties guarantee that $\mathcal{P}(x,p)$ is properly normalized over all phase space.

c) Now the strangeness begins: from the definition of the Wigner transformation in Eq. (8), note that there is no guarantee that $\mathcal{P}(x,p)$ is actually positive. (All that you know from part b is that the integrals over $\mathcal{P}(x,p)$ in either coordinate have to be positive.) As it turns out, $\mathcal{P}(x,p)$ can take on negative values, though the negative regions are small (on the order of $\hbar$) and hence will not have observable consequences in the classical limit, where you look at phase space at scales $\gg \hbar$. To see this for yourself, calculate the Wigner transforms of the first two eigenstates $|\Psi_0\rangle$ and $|\Psi_1\rangle$ of the quantum harmonic oscillator, which has Hamiltonian $\hat{H} = \hat{p}^2/2m + M\omega^2\hat{x}^2/2$. These eigenstates with energies $E_0 = \hbar\omega/2$ and $E_1 = 3\hbar\omega/2$ have the x-space representation:

$$\langle x|\Psi_0\rangle = \left(\frac{\alpha}{\pi}\right)^{1/4} e^{-\alpha x^2/2}, \qquad \langle x|\Psi_1\rangle = \left(\frac{\alpha}{\pi}\right)^{1/4} \sqrt{2\alpha}\, x e^{-\alpha x^2/2}, \tag{12}$$

where $\alpha \equiv M\omega/\hbar$. Show that the corresponding Wigner transforms are:

$$\mathcal{P}_0(x,p) = \frac{1}{\pi\hbar} e^{-\alpha x^2 - p^2/(\alpha\hbar^2)}, \qquad \mathcal{P}_1(x,p) = \frac{2p^2 + \alpha\hbar^2(2\alpha x^2 - 1)}{\alpha\pi\hbar^3} e^{-\alpha x^2 - p^2/(\alpha\hbar^2)}. \tag{13}$$

The function $\mathcal{P}_0(x,p)$ is everywhere positive, but $\mathcal{P}_1(x,p)$ has a pronounced negative dip around $(x,p) = (0,0)$. The fact that $\mathcal{P}(x,p)$ can become negative is one reason it is called a quasiprobability distribution.

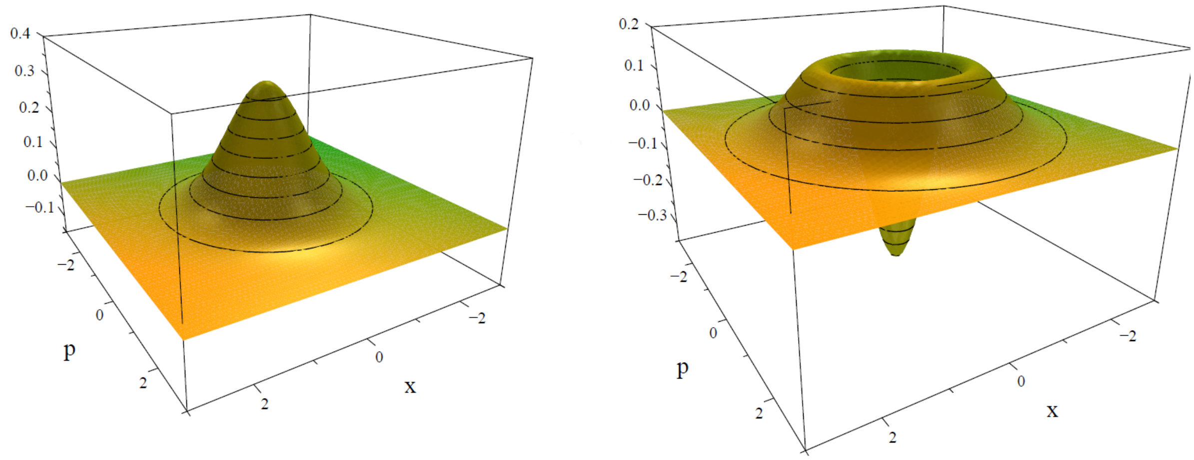

Figure 1: Simple harmonic oscillator eigenstates $\mathcal{P}_0(x,p)$ [left] and $\mathcal{P}_1(x,p)$ [right] in the phase space representation. Images courtesy of Curtright and Zachos, arxiv.org/abs/1104.5269.

d) Another reason $\mathcal{P}(x,p)$ is a quasiprobability is that it is strictly bounded in magnitude, something not true of actual probability distributions in the classical limit (think of Dirac delta functions with their infinite peaks). Show that $\mathcal{P}(x,p)$ must satisfy:

$$|\mathcal{P}(x,p)| \leq \frac{1}{\pi\hbar}. \tag{14}$$

This is a direct reflection of the Heisenberg uncertainty principle: probabilities cannot become arbitrarily concentrated (spiked) in regions of phase space on the order of $\hbar$, since that would allow both x and p to be determined simultaneously with arbitrary accuracy. As $\hbar \to 0$, these height restrictions become relaxed. *Hint:* Use the Cauchy-Schwarz inequality, which states that for any square-integrable complex functions $F(x)$ and $G(x)$, the following holds:

$$\left|\int dx\, F(x)G^*(x)\right|^2 \leq \left(\int dx\, |F(x)|^2\right)\left(\int dx\, |G(x)|^2\right). \tag{15}$$

Experimental interlude: The harmonic oscillator Wigner functions shown in Fig. 1 are extremely important in quantum optics. It turns out that quantizing the electric field leads to a Hamiltonian which has exactly the same form as a harmonic oscillator, with $\hat{x}$ and $\hat{p}$ mapped to the real and imaginary parts of the complex electric field amplitude. The ground state $\mathcal{P}_0$ is called the vacuum state: it represents a state with no photons, but there is still a finite probability of nonzero (x,p) due to quantum fluctuations. The first excited state $\mathcal{P}_1$ represents a single photon. Amazingly, this single photon Wigner function can be experimentally measured using a technique called homodyne tomography. For more details see the experimental paper by A.I. Lvovsky *et al.*, Phys. Rev. Lett. **87**, 050402 (2001) [posted on the course website]. There is also a nice overview at the group's website [http://www.iqst.ca/quantech/research/fock.php] along with a gallery of Wigner functions (check out the Schrödinger cat state!).

e) The final step in surveying the phase space picture is time evolution. First, use the time-dependent Schrödinger equation, $i\hbar(\partial/\partial t)|\Psi\rangle = \hat{H}|\Psi\rangle$, to derive the time evolution of the density operator $\hat{\mathcal{P}}$ introduced above. Show that:

$$\frac{\partial}{\partial t}\hat{\mathcal{P}} = \frac{1}{i\hbar}[\hat{H}, \hat{\mathcal{P}}]. \tag{16}$$

Here $\hat{H} = \hat{p}^2/2m + U(\hat{x})$ is the Hamiltonian operator, and $[\hat{A}, \hat{B}] = \hat{A}\hat{B} - \hat{B}\hat{A}$ is the commutator. Though this is often called the quantum version of the Liouville equation, it is not transparent what the classical limit $\hbar \to 0$ means in this operator form. You will now transform Eq. (16) into the phase space representation, where the connection to the classical Liouville equation is more apparent.

f) In order to make the transformation easier, derive the following Wigner transforms:

$$\begin{aligned} W\{\hat{H}\} &= \frac{p^2}{2m} + U(x) \equiv H(x,p), \qquad W\{[\hat{p}^2, \hat{\mathcal{P}}]\} = -2i\hbar p \frac{\partial}{\partial x}\mathcal{P}(x,p) \\ W\{[\hat{x}, \hat{\mathcal{P}}]\} &= i\hbar \frac{\partial}{\partial p}\mathcal{P}(x,p), \qquad W\{[\hat{x}^2, \hat{\mathcal{P}}]\} = 2i\hbar x \frac{\partial}{\partial p}\mathcal{P}(x,p) \\ W\{[\hat{x}^3, \hat{\mathcal{P}}]\} &= \left(3i\hbar x^2 \frac{\partial}{\partial p} - \frac{1}{4} i\hbar^3 \frac{\partial^3}{\partial p^3}\right)\mathcal{P}(x,p) \end{aligned} \tag{17}$$

When carrying out the derivations, make sure to use the appropriate definition of the Wigner transform, either Eq. (8) or Eq. (9). The former is more convenient with $\hat{x}$ operators, while the latter is easier with $\hat{p}$ operators.

g) Now apply the Wigner transform to both sides of Eq. (16). To make life simpler, expand $U(\hat{x})$ in a Taylor series to third-order, $U(\hat{x}) \approx v_0 + v_1\hat{x} + v_2\hat{x}^2 + v_3\hat{x}^3$, and ignore higher-order contributions. The end result should look like:

$$\frac{\partial \mathcal{P}}{\partial t} = -\{\mathcal{P}, H\} - \frac{v_3\hbar^2}{4}\frac{\partial^3 \mathcal{P}}{\partial p^3} + \cdots, \tag{18}$$

where $\{\cdot, \cdot\}$ is just the classical Poisson bracket of Eq. (3). If you had included more terms in the Taylor series for $U(\hat{x})$, you would end up with higher-order terms in $\hbar$. Remarkably the structure of the equation is analogous to Eq. (2): you have a classical Liouville equation plus correction terms that lead to additional "diffusive" broadening of the probability distribution. Instead of being proportional to Γ as in Problem 1, here the correction terms depend on $\hbar$. The "diffusion" term (with the odd third derivative) is not because of the environment, but rather because of the inherent stochastic nature of quantum mechanics. When the higher-order terms in Eq. (18) are included (don't try this at home!), you can get a closed-form expression for the right-hand side known as the *Wigner-Moyal equation*. If you want to see Wigner-Moyal time evolution in action, the Wikipedia article [http://en.wikipedia.org/wiki/Wigner_quasiprobability_distribution] has several instructive animations: Wigner functions are among the nicest ways to visualize the dynamics of quantum particles.

Interestingly, if your Hamiltonian only has terms up to second order in x (like the harmonic oscillator), so $v_i = 0, \forall i \geq 3$, Eq. (18) predicts a purely classical Liouville time evolution. For such a system, the quantum effects come not from the time evolution equation, but from the fact that your initial distribution $\mathcal{P}(x,p)$ at $t=0$ has to satisfy Eq. (14) (as well as the various normalization conditions in Eq. (11)). You cannot start out with Dirac delta functions as in classical mechanics.

Problem 3: Are we alone in the universe?

Figure 2: Datapoint #1: fossilized evidence of microbial communities dating back to 3.5 billion years ago, discovered in western Australia [Noffke *et al.*, Astrobiology **13**, 1103 (2013)].

In this problem we will see how Bayesian analysis can help us estimate model parameters even in the extreme case of a single datapoint: life had to arise on Earth earlier than 3.5 Gyr (gigayears) ago (see Fig. 1 for the oldest fossilized evidence currently known). As of now we have no other datapoints of life existing anywhere in the universe (though according to a study published in January 2015 there are tantalizing indications that the Curiosity rover on Mars may be on the verge of adding another datapoint; see part f of this problem for an actual calculation of what this would imply). In general, can we say anything about the likelihood of life arising from non-living matter, a process known as *abiogenesis*? Life began early in the Earth's history: the Earth is 4.5 Gyr old, and life arose within the first 1 Gyr of its existence, though almost certainly not within the first 0.5 Gyr because conditions on the very early Earth were inhospitable. This fact seems to support the idea that abiogenesis is a typical occurrence in the universe, fueling optimism about life existing on many Earth-like exoplanets in habitable zones around Sun-like stars. The current estimate based on data from the Kepler spacecraft is that there could be roughly $\approx 10^{10}$ such planets in the Milky Way alone [Petigura *et al.*, Proc. Natl. Acad. Sci. **110**, 19273 (2013)]. If they are of comparable age to the Earth, what fraction of them harbor life? Is the optimism justified?

A more careful evaluation using Bayesian analysis was performed by David Spiegel and Edwin Turner [Proc. Natl. Acad. Sci. **109**, 395 (2012); posted on the course website]. We will derive (in simplified form) a version of their main results. The goal is to determine the conditional probability $\mathcal{P}(\mathcal{M}(x)|\mathcal{D})$. Here $\mathcal{M}(x)$ is the theoretical model for abiogenesis, which depends on some parameter(s) x (in our case it will be a single parameter). $\mathcal{D}$ is the data, which consists of humans having "measured" that life arose on earth by a time $t_{\text{emerge}} \approx 1$ Gyr after the planet's formation. Since $\mathcal{P}(\mathcal{M}(x))$ can be interpreted as the probability of the model being true for a specific value of x, the conditional probability $\mathcal{P}(\mathcal{M}(x)|\mathcal{D})$ encapsulates what we can say about x given the existing data. To evaluate it, we use Bayes's rule:

$$\mathcal{P}(\mathcal{M}(x)|\mathcal{D}) = \frac{\mathcal{P}(\mathcal{D}|\mathcal{M}(x))\mathcal{P}(\mathcal{M}(x))}{\mathcal{P}(\mathcal{D})} \tag{19}$$

The denominator $\mathcal{P}(\mathcal{D})$ is a independent of x, so we can treat it as a normalization constant ensuring that $\int dx\, \mathcal{P}(\mathcal{M}(x)|\mathcal{D}) = 1$. To complete the analysis, we need expressions for $\mathcal{P}(\mathcal{D}|\mathcal{M}(x))$ and $\mathcal{P}(\mathcal{M}(x))$. The latter represents our prior knowledge (rough guess-work!) about x. Let us find each of these expressions in turn.

a) The first ingredient is a model for abiogenesis. We start with the assumption that conditions on a planet right after its formation will not allow life, up until some minimum time t_{min} has passed. If $t = 0$ is the time of planetary formation, we will fix $t_{\text{min}} \approx 0.5$ Gyr, assuming it is comparable for all Earth-like planets. Though abiogenesis is a complex series of chemical events, we can combine them all into a single overall "reaction", which happens at an unknown constant rate λ (a Poisson process) for all times $t \geq t_{\text{min}}$. More precisely, λ is the probability per unit time of abiogenesis, so that the probability of life arising in some short interval dt is λdt (or equivalently, $1 - \lambda dt$ is the probability that life did not arise in this interval). The probabilities in each consecutive interval (i.e. t to $t + dt$ and $t + dt$ to $t + 2dt$) are independent of each other. This model does not preclude life arising independently multiple times, but we are only interested in the first instance. Given the above assumptions, use the laws of probability (and the limit $dt \to 0$) to show that the probability that no life has arisen up to time t after a planet's formation is:

$$P_{\text{no-life}}(\lambda, t) = \begin{cases} 1 & 0 \leq t < t_{\text{min}} \\ e^{-\lambda(t - t_{\text{min}})} & t \geq t_{\text{min}} \end{cases} \tag{20}$$

Hence the probability that life has arisen (at least once) before time t is $P_{\text{life}}(\lambda, t) = 1 - P_{\text{no-life}}(\lambda, t)$. This will be our main model, governed by a single parameter λ which we would like to pinpoint. (As we will see in part c, we will do this by estimating $x \equiv \log_{10} \lambda$, the overall order of magnitude.)

b) To get a sense of the physical meaning of λ, show that the above model predicts the mean time at which life arose as $\langle t \rangle = t_{\text{min}} + \lambda^{-1}$. *Hint:* Which probability distribution do you use to evaluate $\langle t \rangle$? Do not just plug in $P_{\text{life}}(\lambda, t)$, since this is a cumulative distribution: it measures the probability of life emerging at *any* time before t. How do you find the probability of life emerging just during some small interval t to $t + dt$?

c) If you assume λ is set by fundamental chemistry and is the same throughout the universe, let us get a feel for the consequences of its scale. Find the different numerical values of λ (in units of Gyr^{-1}) that would imply the following facts are true for Earth-like planets of comparable age to ours ($t_0 = 4.5$ Gyr):

- λ_1: on average, we are the only such planet at the present time in the entire observable universe where life has emerged (out of $\approx 10^{20}$ Earth-like planets of similar age in the universe)
- λ_2: on average, we are the only such planet at the present time in the Milky Way where life has emerged (out of $\approx 10^{10}$ Earth-like planets of similar age in our galaxy)
- λ_3: on average, life emerges 1 million years after t_{min}. This would virtually guarantee that every Earth-like planet of comparable age in the universe has life.

From top to bottom, these give you a sense of the immense breadth of possible λ values. Since we do not even have a grasp of its order of magnitude, our prior probability distribution $\mathcal{P}(\mathcal{M}(\lambda))$ should reflect this. Let us define $x = \log_{10} \lambda$ and say that all orders of magnitude between $x_{\text{min}} = \log_{10} \lambda_1$ and $x_{\text{max}} = \log_{10} \lambda_3$ are equally probable. Writing $\mathcal{M}(x)$ instead of

$\mathcal{M}(\lambda)$ we will choose our prior probability distribution to be:

$$\mathcal{P}(\mathcal{M}(x)) = \begin{cases} \frac{1}{x_{\max} - x_{\min}} & \text{if} \quad x_{\min} \leq x \leq x_{\max} \\ 0 & x < x_{\min} \quad \text{or} \quad x > x_{\max} \end{cases} \tag{21}$$

d) The implications of our single datapoint $\mathcal{D}$ are more complicated than just specifying an upper bound on Earth's abiogenesis. What $\mathcal{D}$ really states is that: "an intelligent life form on Earth was able to gather evidence at the present time ($t_0 = 4.5$ Gyr) showing that life started before a time $t_{\text{emerge}} = 1$ Gyr in the Earth's history." This presupposes that enough time has passed between t_{emerge} and the t_0 for evolution to produce a scientifically-advanced species capable of investigating fossil evidence of abiogenesis. If life on Earth emerged at $t = 4.0$ Gyr, there almost certainly would not be enough time for evolution to produce a species to collect the datapoint $\mathcal{D}$ at t_0. Let us specify a minimum time delay δt_{evolve} for the evolution of an intelligent species after abiogenesis. Then only abiogenesis events where $t_{\text{emerge}} < t_0 - \delta t_{\text{evolve}} \equiv t_{\text{required}}$ could have any possibility of being measured. Let us choose $\delta t_{\text{evolve}} = 1$ Gyr to set a rough time scale (probably on the short side) for the evolution of intelligence, so $t_{\text{required}} = 3.5$ Gyr is the cutoff for measurable abiogenesis required by evolutionary constraints. Let E be the statement "abiogenesis occurred between $t_{\min}$ and t_{emerge}", and R be the statement "abiogenesis occurred between $t_{\min}$ and t_{required}". Then we will take $\mathcal{P}(\mathcal{D}|\mathcal{M}(x))$ to mean $\mathcal{P}(E|R, \mathcal{M}(x))$, or the probability that E is true given that R and the model $\mathcal{M}(x)$ are true. Using the laws of probability and the result of part a, argue that for any measured value of t_{emerge},

$$\mathcal{P}(\mathcal{D}|\mathcal{M}(x)) = \begin{cases} \frac{P_{\text{life}}(10^x, t_{\text{emerge}})}{P_{\text{life}}(10^x, t_{\text{required}})} & \text{if } t_{\min} \leq t_{\text{emerge}} \leq t_{\text{required}} \\ 0 & \text{if } t_{\text{emerge}} < t_{\min} \text{ or } t_{\text{emerge}} > t_{\text{required}} \end{cases} \tag{22}$$

Hint: Think about the definition of conditional probability. Also note that if $t_{\min} \leq t_{\text{emerge}} \leq t_{\text{required}}$, then R is definitely true if E is true.

e) Putting the result of parts c and d together, use Bayes's rule to determine the posterior probability $\mathcal{P}(\mathcal{M}(x)|\mathcal{D})$. Make sure to normalize by choosing some appropriate numerical value for $\mathcal{P}(\mathcal{D})$. Plot $\mathcal{P}(\mathcal{M}(x)|\mathcal{D})$ versus x to see how the probability behaves. Using numerical integration, figure out the probability that x is between $x_{\min}$ and $x_{\text{mid}} = \log_{10} \lambda_2$. Let us call this probability p_{L}, where L represents extreme loneliness (we are surely alone in our galaxy, and possibly the observable universe). On the other extreme, figure out the probability p_{M} that 99% or more of Earth-like planets of comparable age to ours have seen life emerge. M represents "the more the merrier." How do you like these odds? While p_{M} is greater, p_{L} is still significant, making choosing between these options a tossup. *Hint:* you may find your numerical integrator (Mathematica!?!) gives nonsense when you try to extend the integration range down to

Figure 3: Datapoint #2 (hypothetical): the Gillespie lake outcrop on Mars exhibiting potential signs of microbial structures.

x_{min}. To resolve this, use the $\lambda \to 0$ limit of Eq. (22) (it goes to a simple constant) when integrating below x_{mid}. Use the full expression above x_{mid}.

f) Nora Noffke, the geobiologist responsible for discovering the oldest fossils on Earth (Fig. 1) published an article recently analyzing photos taken by the Curiosity rover on Mars (Fig. 2; see the write-up at: http://shar.es/1bNqS7). She makes a case that Mars exhibits structures remarkably similar to fossilized microbial mats seen on Earth. If these speculations are proven to be true, we would have a second datapoint. What would be the consequences? The Gillespie lake outcrop on Mars where these photos were taken is 3.7 Gyr old, so $t_{\text{emerge}}^{\text{Mars}} = 0.8$ Gyr (Mars has the same age as Earth). Assuming t_{min} is unchanged for Mars, and that life arose there independently of Earth, how would $\mathcal{P}(\mathcal{D}|\mathcal{M}(x))$ change with two datapoints? Recalculate p_{L} and p_{M} from part e (be careful to find the new normalization constant of the distribution first). That's a big pretty big difference, no? Stay tuned: searching for fossilized microbial mats is a major target for the upcoming Mars 2020 rover.

Note: a more complete Bayesian analysis would have allowed the other parameters like t_{min} and δt_{evolve} to vary, with appropriately chosen prior probabilities. This would be significantly more complex, beyond the scope of the problem set. If you are overly bothered by these limitations, feel free to do the analysis and write a research article!