1. Derive the constant of motion associated with inspiral according to the Peters equations. Hint: define $y \equiv e^2$ to get da/dt and dy/dt, then look for a constant in the form $C = af(y)$.

2. Consider a merger of two black holes of arbitrary masses and spins. Suppose that the merger takes place at an *unknown* redshift z. Show that without knowing z, the waveform of the inspiral/merger/ringdown (meaning the observed frequencies, but not the amplitudes) is not sufficient to measure the masses or angular momenta of the black holes uniquely. What combinations of masses, spins, and redshifts can be measured?

3. This problem shows the limits of order of magnitude calculations in some cases. Let's say you'd like to estimate the recoil speed of a merged black hole remnant, due to linear momentum carried away by gravitational radiation. To simplify things, suppose we have two nonrotating black holes of masses M_1 and M_2 that collide head-on, so there is no spin at any point. A theorem from black hole thermodynamics says that the square of the irreducible mass of the final black hole cannot be less than the sum of the squares of the irreducible masses of the initial black holes. For nonrotating black holes, this becomes

$$M_{\text{final}}^2 \geq M_1^2 + M_2^2 \ . \tag{1}$$

Like the increase in entropy, this is an *inequality*, but for our order of magnitude estimate we will assume $M_{\text{final}}^2 = M_1^2 + M_2^2$.

With that assumption, compute the final speed of the remnant (as a fraction of the speed of light, and as a function of M_1 and M_2) assuming that all the radiated energy is carried away in a single direction. For comparison, the best current estimate is that the speed for $M_1/M_2 \approx 10$ is ~ 30 km s^{-1} for nonrotating holes that coalesce via a quasicircular inspiral.

4. Consider the ringdown produced by two $10\,M_\odot$ black holes. Suppose that the ringdown lasts for 2 cycles and emits a total of 1% of the mass-energy of the final black hole. Assuming a nonrotating black hole ($j = 0$), what would be the frequency of the radiation and how long would it last? The frequency is in the range of human hearing (although, of course, not audible!), and sound amplitude is measured in decibels, where 0 dB has an intensity of 10^{-9} erg cm^{-2} s^{-1}. If the BH-BH merger occurs at the distance of the Virgo Cluster (about 50 million light years, or 15 million parsecs, which is $\approx 5 \times 10^{25}$ cm), compare the intensity of the ringdown at Earth with the intensity of the loudest scream ever registered (129 dB, by Jill Drake of the UK). Then, do the same calculation for GW150914, which is the loudest of the events yet detected: two $\sim 30\ M_\odot$ black holes, final spin $j = 0.7$, distance 420 Mpc.

Challenge #1: Make your own estimate of the rate per volume of BH-BH mergers (expressed in number per Gpc^3 per year), including the 90% credible interval, based on the three events reported thus far (for these purposes we assume that LVT151012 was a real event). The first Advanced LIGO run had 49 total days in which both detectors were taking data, so that will be our baseline time. Potentially relevant numbers are: GW150914 was at a distance of 420 Mpc (we'll ignore the uncertainties for simplicity) and had a signal to noise ratio of 23.7; GW151226 was at a distance of 440 Mpc and had a signal to noise ratio of 13.0; LVT151012 was at a distance of 1 Gpc and had a signal to noise ratio of 9.7. Suppose that the threshold for announcing a detection is a signal to noise ratio of 12.0 (recall that LVT151012 was a marginal detection), and remember that for a given event the distance scales as the reciprocal of the signal to noise ratio.

a) With no other information, what would be your best estimate for the rate per volume based on each of the events individually (i.e., without combining them or estimating uncertainties)?

b) How should you estimate the uncertainties for each event individually? More specifically, how would you calculate the 90% credible interval for the rate based on each event individually?

c) How should you combine the information from the three events? Do this without, then with, the uncertainties included.

d) Suppose now that you are given the information that one of the events (pick any of them) was in a direction to which Advanced LIGO was unusually sensitive. What effect, if any, would this have on your best estimate of the rate based on that event (i.e., would it decrease your best estimate, increase your best estimate, or leave it unchanged)?

e) Same question as d), but with regard to the orientation: suppose that one of the events was known to have its binary orbital axis pointed nearly towards us, which means that we see a high amplitude compared to the orientation-averaged amplitude. What effect would this have on your best estimate of the rate from that event alone?

5. Dr. Sane has come to you with a brilliant idea. He has realized that LISA will be the ideal instrument to detect satellites around extrasolar planets. In particular, he envisions a $m = 6 \times 10^{26}$ g satellite (about 10% of Earth's mass, i.e., bigger than any satellite in the Solar System) orbiting with an orbital frequency of $f_{\rm orb} = 5 \times 10^{-5}$ Hz around a planet with mass $M = 2 \times 10^{31}$ g, about ten times Jupiter's mass. At gravitational wave frequencies $f_{\rm GW} < 10^{-3}$ Hz, LISA's expected spectral density sensitivity at signal to noise of 1 is $10^{-19}(10^{-3}\ \text{Hz}/f_{\rm GW})^2\ \text{Hz}^{-1/2}$. Assuming an observing time of 10^8 seconds, evaluate the detection prospects if the system is at a distance of 10 parsecs (about 3×10^{19} cm).

Consider a population of binaries, each of which has reduced mass μ and total mass M. Suppose they are all circular, and that the population is in steady-state, meaning that the number in a given frequency bin is simply proportional to the amount of time they spend in that bin. Also assume that the only angular momentum loss process is gravitational radiation, rather than mass transfer or other effects. For each of the following problems, derive the answers in general and then apply the numbers to WD-WD binaries, where we assume that both masses are $0.6\,M_\odot$ (note that $M_\odot = 1.989 \times 10^{33}$ g $\approx 2 \times 10^{33}$ g).

6. Using the Peters equations for circular orbits of point masses, derive the frequency $f_{\rm min}$ such that the characteristic inspiral time $T_{\rm insp} \sim 1/\left[d\ln f/dt\right]$ is equal to the Hubble time $T_H \sim 10^{10}$ yr. What is the frequency specifically for a WD-WD binary?

7. Below $f_{\rm min}$ the distribution dN/df of sources with frequency will depend on their birth population. Above it, gravitational radiation controls the distribution. Derive the dependence of dN/df on f for $f > f_{\rm min}$ (the normalization is not important).

8. Suppose there are 10^9 WD-WD binaries at frequencies $f_{\rm min} < f < 0.1$ Hz. To within a factor of 2, compute the frequency $f_{\rm res}$ above which you expect an average of less than one WD-WD binary per $df = 10^{-8}$ Hz frequency bin (this is 1/3 yr, or about the frequency resolution expected for the LISA experiment). Very roughly speaking, above $f_{\rm res}$ one can identify individual WD-WD binaries, whereas below it is the confusion limit.

9. Dr. I. M. N. Sane doesn't understand why everyone is so worried about white dwarf noise (which is supposed to be larger than the LISA instrumental noise below about 2×10^{-3} Hz). He asserts that with so many WD-WD binaries in a given bin, the total flux in gravitational waves will be very stable; in particular, he believes that from frequency bin to frequency bin, the flux will vary so little that even a weak additional source will show up easily. He comes to this conclusion by taking the square root of the flux to get a measure of the amplitude.

Show Dr. Sane the error of his ways by doing the following model problem. Let there be N sources in a given frequency bin. Suppose that they are all equally strong, but have random phases between 0 and 2π. Add the complex amplitudes based on those random phases. Take the squared magnitude of the total amplitude as a measure of the typical flux. Determine the mean and standard deviation of the flux that results. You should find that, unlike what happens when you add sources incoherently (i.e., square the amplitudes, then add), the standard deviation of the flux is comparable to the flux, so Dr. Sane's idea fails…to no one's surprise.

Challenge #2: Suppose that you are doing radio observations of a double pulsar system, in which both neutron stars are visible as pulsars. We'd like to determine, qualitatively, how overdetermined the system is. That is, we'd like to know how many aspects of the system can be measured, versus how many parameters there are. This is a deliberately vague question to get you thinking about the process of measurement. If more quantities can be measured than there are parameters, the system is overdetermined and the underlying theory can be tested.