

7.6 Error Estimates

In this example, we assume that you have successfully set up a simulation that can produce “surveys” that match the data provided for this exercise (the same, in fact, as for Exercise 7.5). Hints on how to do this are given in the solution to that exercise.

Having produced a luminosity function, there are three ways to go to estimate errors.

The simplest is just to assume that the fractional error (one standard deviation) in the luminosity function is $1/\sqrt{N}$, where N is the number of objects in the bin of interest. This is cheap and cheerful, but fairly rough.

The second method is to run a full Monte Carlo simulation, making many surveys of the same size as the one you have. Here the assumption is that you have a good enough idea of the underlying luminosity function for the error estimates to be useful. In our case, of course, we know it exactly, so this method provides a benchmark to check other methods. Finally, we can try a bootstrap; sample with replacement from the one survey you actually have.

Below are the results, for the data provided for this exercise. The Monte Carlo simulation generated 800 surveys, as did the bootstrap. The Monte Carlo is much slower than the bootstrap, because of the difficulties in drawing the required random numbers (see solution for 7.5). The numbers quoted are standard deviation of estimates in bin i , divided by the median of the estimates in bin i . For the simple method, we just quote $1/\sqrt{N}$, which out to be comparable.

log power	\sqrt{n}	Monte Carlo	bootstrap
-0.75	1.	∞	0.73
-0.25	0.58	0.56	0.49
0.25	0.33	0.44	0.57
0.75	0.14	0.42	0.50
1.25	0.10	0.42	0.51
1.75	0.18	0.49	0.55
2.25	0.58	0.95	0.93

A couple of points emerge. First, the Monte Carlo and bootstrap agree quite well, with the $1/\sqrt{N}$ method being systematically low. Second, there is the bizarre ∞ – what does this mean?

The problem is that, in the high and low bins, there are very few objects and in many realization of the Monte Carlo simulation, there are none. This gives a very peculiar distribution of estimates – see Figure 1. In the bottom bin, the median estimate is actually zero for the Monte Carlo simulation. It isn’t zero for the bootstrap, but there is the same detached peak at zero.

This shows that we need to be very careful of what sort of error estimates we quote. In other bins, the distributions are very asymmetrical as well (although not with detached peaks). The text quoted an interquartile range for the error – this is better than a standard deviation, but still misleading in some cases.

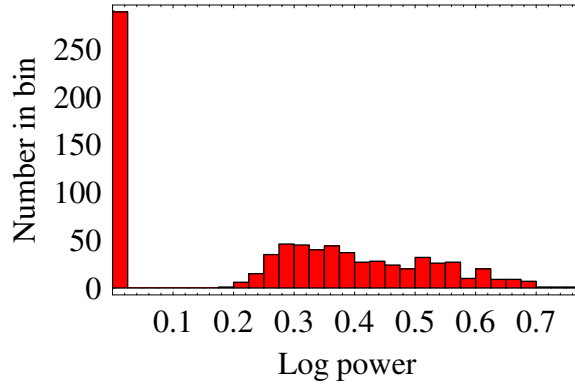


Figure 1: Distribution of estimates of the normalized density in the lowest power bin of the simulation, from 800 bootstrap repetitions. The large number of zero detections is a warning flag!

One way out is to work with percentiles of the distribution of estimates. In Figure 2 are shown error bars which extend from the 25% point to the 75% point. As we see, the error estimates are in accord and the agreement with the theoretical curve is good. The possibility of zero space density in the bottom bin is correctly captured.

One remaining difficulty is not shown in Figure 2, but should be apparent in your data - the estimate in the top bin is not very good. This is because the luminosity function is dropping away so steeply that binning inevitably introduces a lot of bias.




Figure 2: Error bars joining 25% and 75% points of the distributions of density, at each power level, for Monte Carlo (black) and bootstrap (red).