Coding in advance of the April 8, 2019 class

For this assignment you will perform the first part of the analysis described in the April 8 lecture, using the full baryonic Tully-Fisher data set on the website (which contains 118 points). For this analysis, assume that the rotation speed $v_{\rm rot}$ is measured with zero uncertainty.

From the notes, the model you are fitting is

$$\log_{10} M_{\rm bary} = \tan\theta \log_{10} v_{\rm rot} + b \tag{1}$$

and thus the parameters are θ and b. For the ith rotation speed $v_{\rm rot,i}$, and for a given θ and b, the model prediction for $\log_{10} M_{\rm bary}$ is then

$$m_i = \tan\theta \log_{10} v_{\rm rot,i} + b \tag{2}$$

and the measurement is

$$d_i = \log_{10} M_{\rm bary,i} \ . \tag{3}$$

The uncertainty we are given in the table is

$$\sigma_i = \sigma_{\log_{10} M_{\rm bary}} \ , \tag{4}$$

and thus the total χ^2 for the data set is

$$\chi^2 = \sum_i \frac{(m_i - d_i)^2}{\sigma_i^2} \ . \tag{5}$$

Your task is to go through a reasonably fine grid in θ and b and to determine: (1) the minimum χ^2 over all θ and b, and (2) the values of θ (in radians) and b that minimize χ^2. Make a comment about whether this model is a good fit to the data, given the minimum chi squared and the number of degrees of freedom.

To give you advance notice, for next week I'll ask you to work with the same data set, using the same assumption that $v_{\rm rot}$ is measured with zero uncertainty, but I'll ask for more things including marginalized posterior probability densities for θ by itself and b by itself.