Measuring Lack-of-Fit of a Bayesian Model

David Fletcher[a] and **Peter Dillingham**[b]

[a]Department of Mathematics and Statistics
University of Otago
Dunedin, New Zealand
dfletcher@maths.otago.ac.nz

[b]School of Science and Technology
University of New England
Armidale, New South Wales, Australia
pdilling@une.edu.au

Keywords: Bayesian model, lack-of-fit, overdispersion, posterior predictive p-value

Abstract: When using the Bayesian framework to fit a model, it is common practice to use posterior predictive checking as a means of assessing lack-of-fit. In particular, the posterior predictive p-value is often used to help decide whether the model needs to be improved. We propose a new type of posterior predictive check that quantifies the amount of lack-of-fit of the model. We use examples to illustrate the benefits of this new tool, and compare it with the posterior predictive p-value.