Statistica Sinica **33** (2023), 1-24
doi:https://doi.org/10.5705/ss.202021.0258

REFERENCE PRIORS FOR THE GENERALIZED EXTREME VALUE DISTRIBUTION

Likun Zhang and Benjamin A. Shaby

University of Missouri and Colorado State University

Abstract: We derive a collection of reference prior distributions for a Bayesian analysis under the three-parameter generalized extreme value (GEV) distribution. These priors are based on an established formal definition of noninformativeness. They depend on the ordering of the three parameters, and we show that the GEV is unusual in that some orderings fail to yield proper posteriors for any sample size. We also consider a reparametrization that explicitly regards a return level estimation, which is the most common goal of a GEV analysis, to be the most important inferential task. We investigate the properties of the derived priors using a simulation, and apply the priors to an analysis of a fire threat index in California.

Key words and phrases: Noninformative priors, objective Bayes, posterior normality.

1. Introduction

For a Bayesian analysis under the three-parameter generalized extreme value (GEV) model, a formal notion of the noninformativeness of the prior distribution can be achieved using a reference analysis. We derive reference priors under the standard parametrization of the GEV, showing that the resulting posterior distributions are improper for some, but not all orderings of the parameters. We further show that re-parametrizing to prioritize an inference on a high quantile results in the same behavior as the standard parametrization. Using a simulation, we compare the performance of the reference priors with that of two previously recommended priors: an alternative rule-based noninformative prior, and a prior based on domain knowledge, finding none to be uniformly most desirable. The tradeoffs are evident in our analysis of the extremes of a fire threat index observed in California. In the absence of specific domain knowledge about the tail of the process under investigation, particularly when an estimation about a high quantile is the goal of the analysis, the reference prior described here, which prioritizes an inference on that high quantile, might be considered a good default

Corresponding author: Likun Zhang, Department of Statistics, University of Missouri, Columbia, MO 65211, USA. E-mail: likun.zhang@missouri.edu.