

**Abstract:**

Fay-Herriot estimators are widely used to produce small area statistics when only area-level aggregate data are available. This paper investigates the conditions under which bivariate Fay-Herriot models give useful reductions in approximate prediction mean squared error (APMSE) compared to separate univariate models. The APMSE is shown to be equal under these two approaches if the sampling errors and the area-level random effects have proportional variance-covariance matrices, even if there is high correlation between the two variables of interest. The ratio of APMSEs is calculated numerically for a range of settings, and this numerical study is summarised using a novel regression tree approach. Univariate and bivariate estimators are compared using data on 30 indicators from the 2011-2012 New Zealand Health Survey, with MSEs estimated by a parametric bootstrap approach. The results suggest that bivariate modelling can be worthwhile, but only for a minority of cases.