

The issue of hierarchical data that take the form of clusters with random size has received considerable attention. Much of this has been concerned with so-called informative cluster sizes, where the cluster size is associated with the data values in the cluster. In this paper, we address several related issues that occur even when the cluster size and the data are unrelated. The first is computational: fitting models to clusters of varying sizes is often more complicated than when all clusters are of the same size. The second is more theoretical in nature: with unequal sized clusters there is typically no complete sufficient statistic \citep{ourselves1}. We study in detail a setting that is both sufficiently simple and versatile, that of normally distributed clusters that have a compound-symmetry covariance structure. In spite of the simplicity of this setting, we are able to deduce a number of results, and formulate data-analytic approaches based upon these, that simplify computations, especially in huge datasets. Maximum likelihood is compared and contrasted with sample splitting \citep{MOVEID11}, a generic pseudo-likelihood based approach that allows the data first to be analysed "by parts," the results of which are then combined using suitable weights. The findings are illustrated using data from a developmental toxicity study, where clusters are formed of fetuses within litters.