Notice

Page 40 of this monograph was amended after its initial release on May 27, 2016. The first two full paragraphs on page 40 have been corrected as follows:

> For most predictive modeling and machine learning applications, this is superior to a single train/test split, since *all* of the data is being used to test out-of-sample model performance as opposed to a single subset. However, it is often of limited usefulness for most insurance modeling applications, since cross validation has an important limitation: in order for it to be effective, the “training” phase of the procedure must encompass *all* the model-building steps. For a GLM, where the bulk of the model-building is the variable selection and transformation, that part would need to be included as well.

Page 46 was amended to read:

> The set of **partial residuals** for any predictor x_j in a model is defined as follows:
>
> $$r_i = (y_i - \mu_i)g'(\mu_i) + \beta_j x_{ij}, \tag{13}$$
>
> where $g'(\mu_i)$ is the first derivative of the link function. For a log link model, Equation 13 simplifies as follows:
>
> $$r_i = \frac{y_i - \mu_i}{\mu_i} + \beta_j x_{ij}. \tag{14}$$
>
> In Equation 14, the residual is calculated by subtracting the model prediction from the actual value, and then adjusted to bring it to a similar scale as the linear predictor (by dividing by μ_i). Then, $\beta_j x_{ij}$—that is, the part of μ_i that x_j is responsible for—is added back to the result.