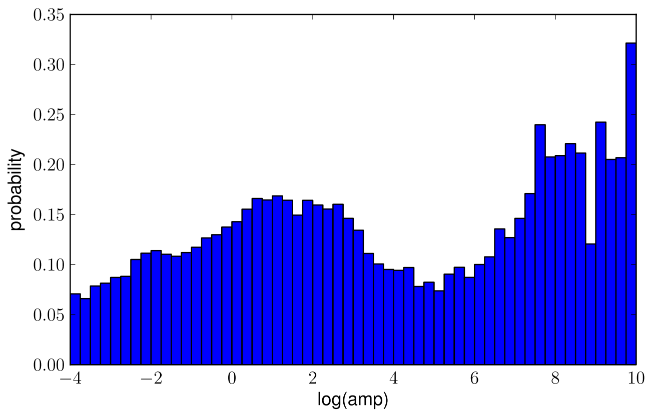

Figure 14. Coda detection probability as a function of the log amplitude of the triggering arrival. The color version of this figure is available only in the electronic edition.

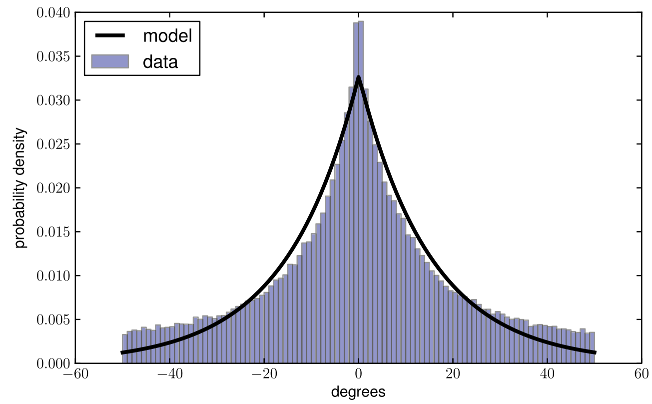

Figure 16. Coda azimuth difference versus triggering arrival. The color version of this figure is available only in the electronic edition.

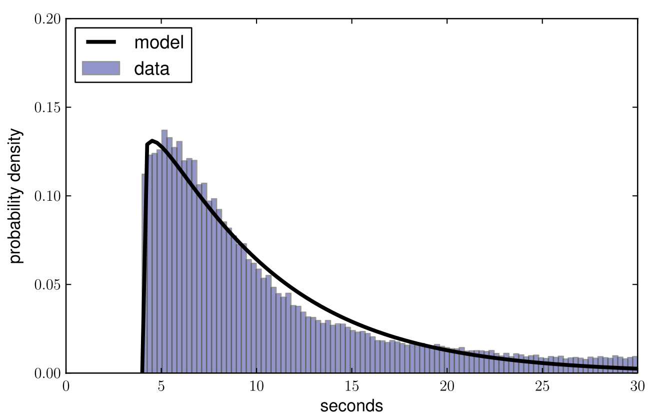

Figure 15. Time delay for coda arrival after the triggering arrival. The color version of this figure is available only in the electronic edition.

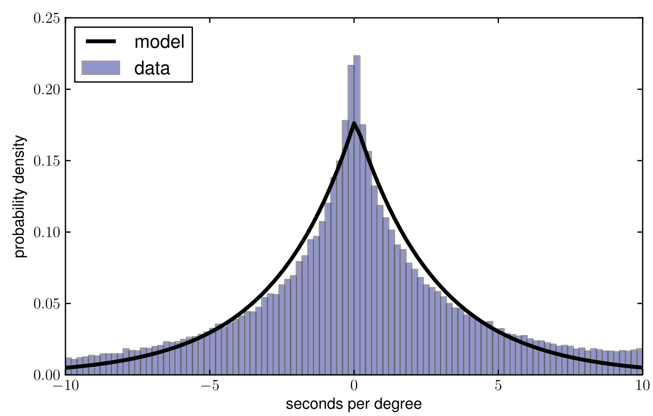

Figure 17. Coda slowness difference versus triggering arrival. The color version of this figure is available only in the electronic edition.

this type. An example of coda arrivals can be seen in Figure 1 around 20 s after the main arrival, which is at offset 100 s. One might imagine that coda arrivals can be lumped in with false arrivals, but it turns out that the attributes of coda arrivals are strongly correlated with those of the triggering arrival. If the coda arrivals are not modeled explicitly, then our inference will end up hypothesizing additional spurious events as the most likely explanation for many of the coda arrivals.

We model the parameters of the coda arrival with a distribution P_γ, which is a function of the parameters of the triggering arrival. Whether the triggering arrival was a false arrival, or caused by an event, or itself triggered by another arrival, is immaterial. We define η^a as the arrival triggered by arrival a (at the same station), or ζ if there is no such triggered arrival. Now, the probability that an arrival a triggers a coda arrival is given by $P_{\gamma,d}(a)$, which is a function of the amplitude of the arrival a. We estimate $P_{\gamma,d}$ with a nonparametric model by discretizing the log amplitude of the triggering arrival into buckets of size 0.25 between −4 and 10. This distribution is displayed in Figure 14. Any points outside these extreme values are mapped to the nearest bucket.

One problem that arises while training the model for coda arrivals is that the IDC analysts do not annotate coda arrivals in the LEB bulletin, which makes it hard to estimate the parameters of P_γ. Our solution is to heuristically annotate some of the unassociated arrivals as coda arrivals and use this annotation to learn P_γ and also P_ω. Our procedure searches the training data for any unassociated arrivals within 30 s of a prior arrival at the same station and with an azimuth and slowness within 50° and 10 s per degree, respectively, of the prior arrival's values, and to mark such arrivals as coda.

We model the distribution of the time delay between the coda arrival and the triggering arrival as a gamma distribution (see Fig. 15):

$$\eta_t^a - a_t \sim \Gamma(\rho_t, \nu_t), \quad \text{i.e.}$$

$$P_{\gamma,t}(\eta_t^a|a) = \frac{1}{\Gamma(\rho_t)\nu_t^{\rho_t}}(\eta_t^a - a_t)^{\rho_t - 1}\exp\left(-\frac{\eta_t^a - a_t}{\nu_t}\right).$$

The differences in azimuth, slowness, and log amplitude of the coda versus the triggering arrival are all modeled as Laplace distributions (Figs. 16, 17, and 18):