



Figure 1: Comparisons of different adversarial models. Fix a round t . Consider the simple scenario wherein each of the M clients contains only one data that is drawn from Gaussian distribution. Each of the three figures shows the empirical densities of the M data points. As shown in (a), under the adversarial unavailability model, the system adversary can inspect the local data and remove up to ϵM data points that are the farthest from the true mean (i.e., the blue bars). Thus the remaining empirical density is biased. In (b), the system adversary injects up to ϵM data points increasing the empirical densities of small values (shown as red bars); however, the empirical density of good data is unchanged and still unbiased. In (c), under the Byzantine attack model, the system adversary first removes the data points in the blue bars and then injects the same amount of data points with small values.