



FIGURE 4. Median Absolute Errors by validation month, comparing losses on the validation set for all models. There is one dot for each set of hyperparameters tested. The horizontal axis corresponds to hyperparameter values and the vertical axis to errors. The dots are ordered from lowest median absolute error to highest for each of the model families of elastic net, random forests, and gradient boosting. The length of the lines differs on the x-axis, because some models had more model-specific hyperparameter sets than others. The length of the lines differs on the y-axis because some model families had different ranges of errors. In every validation month, one of the gradient boosting models performed best. The best-performing random forests model had estimated generalization errors 1 to 7 percent higher than the best gradient boosting model. The best elastic net model had errors that were often twice the errors of the best gradient boosting model.