Uninformative priors prefer simpler models HENRY MATTINGLY,
Princeton University, MICHAEL ABBOTT, Institute of Physics, Jagiellonian University, BENJAMIN MACHTA, Princeton University — The Bayesian framework
for model selection requires a prior for the probability of candidate models that is
uninformative—it minimally biases predictions with preconceptions. For parameterized models, Jeffreys' uninformative prior, p^J, weights parameter space according
to the local density of distinguishable model predictions. While p^J is rigorously
justifiable in the limit that there is infinite data, it is ill-suited to effective theories
and sloppy models. In these models, parameters are very poorly constrained by
available data, and even the number of parameters is often arbitrary. We use a principled definition of 'uninformative' as the mutual information between parameters
and their expected data and study the properties of the prior p^* which maximizes it.
When data is abundant, p^* approaches Jeffreys' prior. With finite data, however,
p^* is discrete, putting weight on a finite number of atoms in parameter space. In
addition, when data is scarce, the prior lies on model boundaries, which in many
cases correspond to interpretable models but with fewer parameters. As more data
becomes available, the prior puts weight on models with more parameters. Thus, p^*
quantifies the intuition that better data can justify the use of more complex models.

Henry Mattingly
Princeton University