

20.5: Choosing a Prior

The impact of priors on the resulting inferences are the most controversial aspect of Bayesian statistics. What is the right prior to use? If the choice of prior determines the results (i.e., the posterior), how can you be sure your results are trustworthy? These are difficult questions, but we should not back away just because we are faced with hard questions. As we discussed previously, Bayesian analyses give us interpretable results (credible intervals, etc.). This alone should inspire us to think hard about these questions so that we can arrive with results that are reasonable and interpretable.

There are various ways to choose one's priors, which (as we saw above) can impact the resulting inferences. Sometimes we have a very specific prior, as in the case where we expected our coin to land heads 50% of the time, but in many cases we don't have such strong a starting point. *Uninformative priors* attempt to influence the resulting posterior as little as possible, as we saw in the example of the uniform prior above. It's also common to use *weakly informative priors* (or *default priors*), which influence the result only very slightly. For example, if we had used a binomial distribution based on one heads out of two coin flips, the prior would have been centered around 0.5 but fairly flat, influence the posterior only slightly.

It is also possible to use priors based on the scientific literature or pre-existing data, which we would call *empirical priors*. In general, however, we will stick to the use of uninformative/weakly informative priors, since they raise the least concern about influencing our results.

This page titled [20.5: Choosing a Prior](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Russell A. Poldrack](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.