Online appendices for "Some guidance for the choice of priors for Bayesian structural models in economic experiments"

James R. Bland

2025-01-30

1 Simulation details

This Appendix outlines in more detail the procedures used for the Monte Carlo exercise discussed in Section 3.3 of the main article. The replication code for this simulation can be found at: https://github.com/JamesBlandEcon/ChoiceOfPriors

The simulation was implemented in *R* (R Core Team 2021) with both Bayesian and maximum likelihood estimation calling a single *Stan* (Carpenter et al. 2017) program. Each Bayesian estimation involved four Monte Carlo chains, each with a warm-up of 5,000 iterations and 5,000 samples (this is larger than *RStan*'s default), and the `adapt_delta` tuning parameter set to 0.8 (this is *RStan*'s default). The maximum likelihood estimations used the true parameter values as initial values.

As *Stan* implements Hamiltonian Monte Carlo, there are several diagnostics used to assess whether a posterior simulation has adequately converged. If one of these diagnostics fails, then inference based on that posterior simulation may be unreliable. For the purposes of this simulation exercise, I focus just on the estimations that were performed successfully by the following metrics:

- The number of divergent transitions after warm-up must be zero,
- The Gelman-Rubin statistics (Gelman and Rubin 1992) must be less than 1.01 for all of a model's parameters, and
- The effective sample size must be greater than 1,000 for for all of a model's parameters, including the generated quantities.

For each simulation step, *all three* estimated Bayesian models (i.e. for the three priors used in the exercise) must pass these hurdles in order to be included in the simulation exercise. The simulation ran until 1,000 simulation steps had successfully been completed. That is, until all three Bayesian estimators had successfully converged by these measures 1,000 times. Table 1 summarizes the number of unsuccessful simulation steps for each of the three priors used. Almost all of the dropped simulation steps were due to failing to have zero divergent

Table 1: Counts of estimations dropped for each of the four Bayesian models.

prior	divergences	Rhat	n_eff	total
calibrated	280	0	0	280
x2	378	0	0	378
x8	494	1	6	497

transition after warmup. Note also that the model estimated with the calibrated prior failed to converge less frequently. This is an important technical issue not discussed in the main body of the text: another consideration for prior calibration can be if the sampler converges at all.

Figure 1 shows simulated parameter realizations that met these criteria (blue), and those that did not meet these criteria (red). Here we can see that the estimators typically had trouble with datasets generated with larger realizations of r or smaller realizations of λ. These issues are as much a problem with the econometric model (i.e. prior and likelihood combination) as they are with the experiment design, so if this exercise is performed *before* running the experiment, this provides us with an opportunity to re-design the experiment if these convergence properties are not satisfactory.

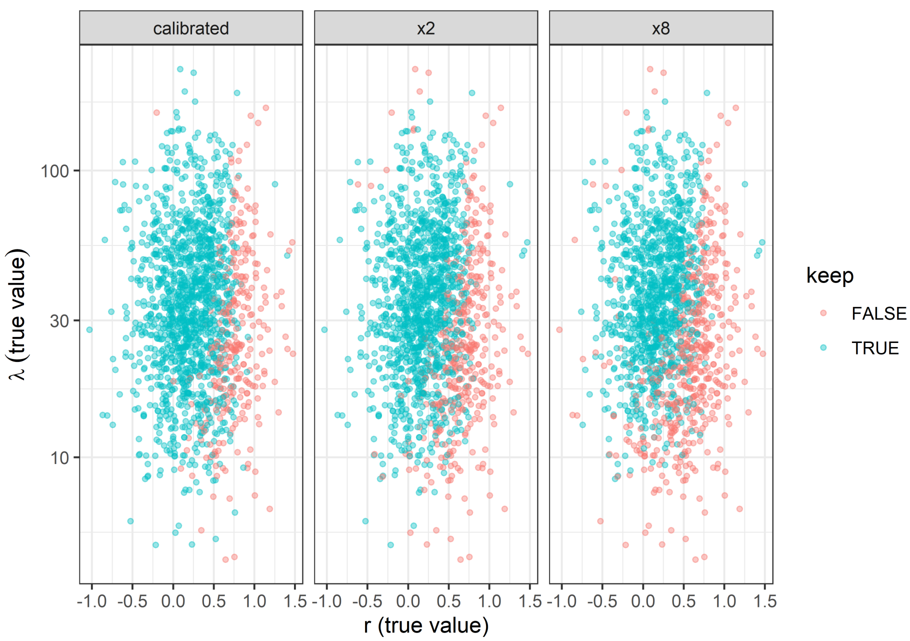

Figure 1: True parameter values for kept (blue) and dropped (red) observations in the simulation.

Table 2: Root mean squared error of posterior mean estimates and maximum likelihood estimates. The certainty equivalent is for a lottery that mixes 50-50 over prizes \$10 and \$50.

	calibrated	×2	×8	MLE
Certainty equivalent	0.91	0.99	1.46	12.64
λ	12.23	14.11	16.13	35.04
r	0.12	0.14	0.29	60.86

2 Simulation with an overly-informative prior

This Appendix reports the results of a simulation similar to the one reported in the previous section of these Online Appendices and Section 3.3 of the main text. In particular, here I simulate data as coming from the following distribution:

$$r_i \sim N(0.27, 0.36^2 \times 8^2)$$
$$\log \lambda_i \sim N(3.45, 0.59^2 \times 8^2)$$

That is, the standard deviations are inflated by a factor of 8 relative to the calibrated prior.

The estimation process is identical to that described in Section 1 of these Online Appendices. Therefore, this simulation provides information on the consequences of calibrating a prior that is more informative than the underlying population distribution.

For what follows, note that the "x8" prior is now correctly specified relative to the underlying population distribution.

Table 2 reports the root mean squared error of parameter estimates. Compare this to Table 2 in the main text, where the calibrated prior is correctly specified.

Figure 2 plots the estimated certainty equivalents (vertcal coordinate) against their true values (horizontal coordinate). Compare this plot to Figure 4 of the main text, where the calibrated prior is correctly specified.

References

Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. "Stan: A Probabilistic Programming Language." *Journal of Statistical Software* 76 (1).

Gelman, Andrew, and Donald B Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences." *Statistical Science* 7 (4): 457–72.

R Core Team. 2021. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.

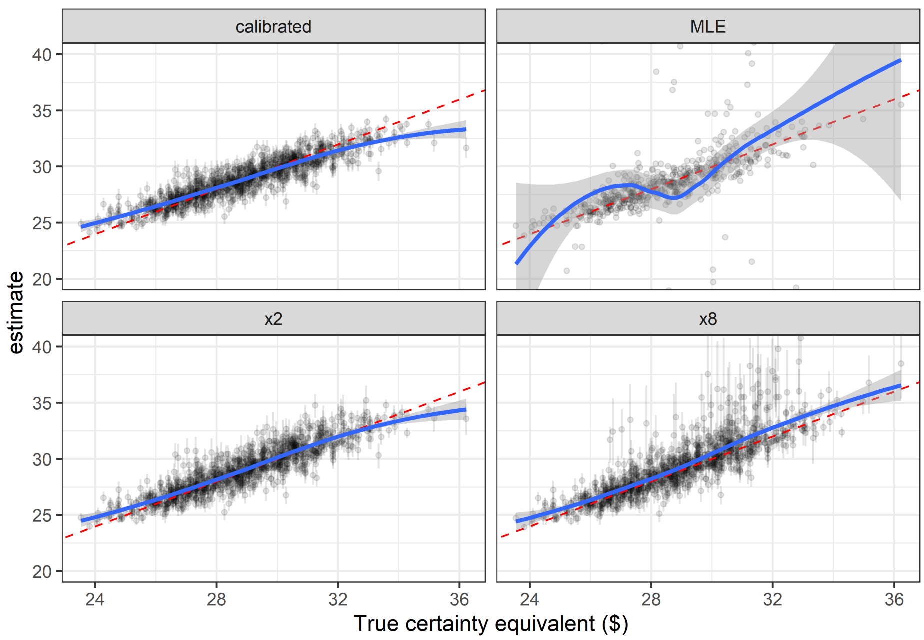

Figure 2: Simulated estimates of a certainty equivalent (vertical axis) against their true values (horizontal axis) for various prior specifications and maximum likelihood estimates. Dots show posterior means, vertical lines show a 50% Bayesian credible region (25th-75th percentile). No expression of uncertainty is shown for the maximum likelihood estimates. The red dashed line is a 45° line. The blue curve shows a smoothed mean of the posterior means.