

Answers to Odd-Numbered Exercises for Chapter 25  
Fox, *Applied Regression Analysis  
and Generalized Linear Models*,  
*Third Edition* (Sage, 2016)

John Fox

Answers last modified: 2023-03-20

This document provides worked-out answers to the odd-numbered exercises in bonus Chapter 25 on Bayesian estimation of regression models (excluding the data-analysis exercises on the website for the text). The answers to both odd- and even-numbered questions, are available from the author for instructors using the text in a college or university class.

Although the text is written to be software-neutral, some of the exercises require the use of statistical software. I employed the R statistical computing environment (R Core Team, 2021) for these exercises, but you should feel free to substitute other appropriate statistical software for R. See Fox and Weisberg (2019) for an introduction to R in the context of regression analysis that largely overlaps with the text.

## Exercises for Chapter 25

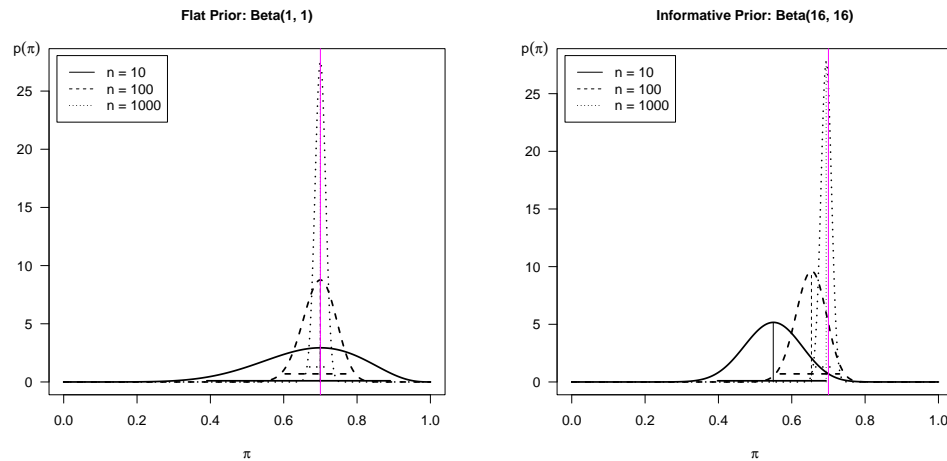
### Exercise 25.1

According to Bayes's theorem,  $\Pr(A|P) = \Pr(P|A)\Pr(A)/\Pr(P)$ , where  $\Pr(P) = \Pr(P|A)\Pr(A) + \Pr(P|\bar{A})\Pr(\bar{A}) = 1 \times .1 + .1 \times .9 = .19$ , and so  $\Pr(A|P) = (1 \times .1)/.19 = .52$ , which, indeed, is about half of  $\Pr(P|A) = 1$ .

To understand this result (and to take the hint in the problem), imagine a population in which no one has antibodies, so that instead of  $\Pr(A) = .1$  (i.e., small),  $\Pr(A) = 0$ . Then anyone who tests positive is necessarily a false positive, because no one is really positive, and  $\Pr(A|P) = 0$ , as Bayes's theorem tells us for this case. Thus  $\Pr(A|P)$  depends on the prevalence  $P(A)$  of antibodies in the population as well as on the sensitivity and specificity of the test.

### Exercise 25.3

- (a) Here are graphs of the posterior distribution of  $\pi$  for the flat Beta( $a = 1, b = 1$ ) prior (on the left) and the informative Beta( $a = 16, b = 16$ ) prior (on the right):



In both cases, the solid lines show the posterior distribution, the posterior mode (vertical line), and 95% central posterior interval (horizontal line) for  $n = 10$ , the broken lines for  $n = 100$ , and the dotted lines for  $n = 1000$ . The vertical magenta line is at the proportion of heads in the data, which is  $.7$  in all cases. Each posterior distribution is Beta( $h + a, n - h + b$ ), where  $h$  is the number of heads in the sample (e.g.,  $7$  when  $n = 10$ ).

For the flat prior, the posterior mode is always equal to  $.7$  regardless of the sample size, but as the sample size grows, the posterior distribution becomes less spread out and thus the 95% posterior interval gets narrower. The informative prior pulls the posterior distribution towards its center at  $\pi = .5$ , but as the sample size grows, the influence of the prior weakens and by  $n = 1000$ , the posterior distribution is centered very close to  $\pi = .7$ ; as well, as  $n$  increases, the spread of the posterior distribution decreases and the 95% posterior interval gets narrower.

- (b) As in part (a), the posterior distribution of  $\pi$  is Beta( $h + a, n - h + b$ ), but now  $n = 10, a = 16$  and  $b = 16$ , while the number of heads  $h$  varies from  $0$  to  $10$ ; the corresponding posterior modes are

$h$	0	1	2	3	4	5
Posterior mode	.375	.400	.425	.450	.475	.500
$h$	6	7	8	9	10	
Posterior mode	.525	.550	.575	.600	.625	

When the sample proportion ( $h/10$ ) is below  $.5$ , so is the posterior mode, but because the sample size  $n = 10$  is small, the posterior mode is drawn strongly upwards towards the center of the

informative prior distribution at  $\pi = .5$ . Similarly, when the sample proportion is above .5, the posterior mode is drawn downwards towards  $\pi = .5$ .

### Exercise 25.5

- (a)  $\sigma^2/n$  is the sampling variance of the mean of an independent random sample of size  $n$ .
- (b) Precision is the inverse of variance. Consequently, the precision of the sample mean is the inverse of its sampling variance, and the precision of the prior mean is the inverse of the prior variance. Variance is a measure of dispersion or variation; if, for example, the sampling variance of the sample mean is large, then an estimate of  $\mu$  based on the sample mean will be imprecise—will vary a great deal from sample to sample. The inverse of the variance of course increases as the variance declines, and thus is a measure of precision in the ordinary sense of the word.

In Bayesian estimation, there are two sources of information about  $\mu$ , one coming from the data (the sample mean  $\bar{y}$ ), and the other from the prior (the prior mean  $\mu_0$ ). As the precision of each of these sources of information increases (i.e., as the corresponding variance declines), the amount of information about  $\mu$  increases—as it turns out, proportionally—and so it makes sense to weight each value by its precision to form a combined estimate of  $\mu$ .

- (c) Combining information about  $\mu$  from the data and the prior should produce a better estimate than using either source of information alone, and so it's sensible that the precision of the posterior mean, which combines the data with the prior, should be the sum of the two precisions and hence larger than either. Because the variance is the inverse of the precision, the posterior variance is smaller than the variance of the sample mean or the prior variance.

### Exercise 25.7

I was able to obtain a usable solution by tightening the priors for the covariance components in the logit part of the model (although the resulting effective sample sizes for these covariance components are still relatively small), decreasing the standard deviations of the normal priors by a factor of 2:  $\log_e \psi_{Z_j} \sim N(0, 0.2^2)$  for  $j = 0, 1$ , and  $z(\rho_{Z_0 Z_1}) \sim N(0, 0.3^2)$ . On the one hand, as indicated in the problem and in the text, adjusting the priors based on the results violates the Bayesian paradigm. On the other hand, the originally specified priors had no deep justification and the modified priors are still quite vague.

```
> library("rstan")
Loading required package: StanHeaders
Loading required package: ggplot2
rstan (Version 2.21.7, GitRev: 2e1f913d3ca3)

. . .

> (ncores <- parallel::detectCores())
[1] 20
> options(mc.cores=ncores)
> rstan_options(auto_write=TRUE)

> data(Blackmore, package="carData")

> Blackmore$nonzero <- as.numeric(Blackmore$exercise > 0)
> data.blackmore <- with(Blackmore,
+   list(n = nrow(Blackmore),
+        J = nlevels(subject),
+        subject = as.numeric(subject),
+        patient_control = ifelse(group == "patient", 1/2, -1/2),
+        age = age - 8,
+        log_exercise = ifelse(exercise == 0, 0, log2(exercise)),
+        nonzero = as.numeric(exercise > 0),
+        n_nonzero = sum(nonzero),
```

```

+                                     index_nonzero = which(as.logical(nonzero)))
+ )

> # hurdle model with random intercepts and slopes

> blackmore.model <- "
+   data {
+     int<lower=0> n;                               // total number of observations
+     int<lower=0> J;                               // number of subjects
+     int subject[n];                              // subject index
+     vector[n] patient_control;                   // patient - control regressor
+     vector[n] age;                               // within-subject predictor
+     vector[n] log_exercise;                       // response variable
+
+     int<lower=0, upper=1> nonzero[n];            // exercise > 0?
+     int n_nonzero;                               // number of nonzero observations
+     int<lower=1, upper=n>
+       index_nonzero[n_nonzero];                 // indices of nonzero observations
+   }
+   parameters {
+     matrix[J, 2] B;                               // subject-specific intercepts and slopes
+                                               // (deviations)
+     vector[4] beta;                               // group-level coefficients
+     real log_sigma;                               // error standard deviation
+     vector[2] log_psi_b;                          // standard deviations of coefs
+     real z_rho_b;                                 // Fisher z of correlation of subject-specific
+                                               // coefficients
+     matrix[J, 2] Z;                               // subject-specific intercepts and slopes for
+                                               // logit part of model (deviations)
+     vector[4] zeta;                               // group-level coefficients for logit part
+                                               // of model
+     vector[2] log_psi_z;                          // standard deviations of subject-specific
+                                               // logit coefs
+     real z_rho_z;                                 // Fisher z of correlation of subject-specific
+                                               // logit coefficients
+   }
+   transformed parameters {
+     // conditional expectation for linear part of model
+     vector[n] mu = beta[1] + B[subject, 1] + beta[2]*patient_control
+       + beta[3]*age + beta[4] * patient_control .* age
+       + B[subject, 2] .* age;
+     // linear predictor for for logit part of model
+     vector[n] eta = zeta[1] + Z[subject, 1] + zeta[2]*patient_control
+       + zeta[3]*age + zeta[4] * patient_control .* age
+       + Z[subject, 2] .* age;
+     real sigma = exp(log_sigma);
+     vector[2] psi_b = exp(log_psi_b); // standard deviations of
+                                               // subject-specific coefficients
+     vector[2] psi_z = exp(log_psi_z); // standard deviations of subject-
+                                               // specific logit coefficients
+     matrix[2, 2] Psi_b;                    // covariance matrix of subject-
+                                               // coefficients
+     matrix[2, 2] Psi_z;                    // covariance matrix of subject-
+                                               // specific coefficients
+     real rho_b = tanh(z_rho_b);             // correlation between subject-specific
+                                               // intercepts and slopes
+     real rho_z = tanh(z_rho_z);             // correlation between subject-specific
+                                               // logit intercepts and slopes
+     Psi_b[1, 1] = psi_b[1]^2;

```

```

+       Psi_b[2, 2] = psi_b[2]^2;
+       Psi_b[1, 2] = psi_b[1]*psi_b[2]*rho_b;
+       Psi_b[2, 1] = psi_b[1]*psi_b[2]*rho_b;
+       Psi_z[1, 1] = psi_z[1]^2;
+       Psi_z[2, 2] = psi_z[2]^2;
+       Psi_z[1, 2] = psi_z[1]*psi_z[2]*rho_z;
+       Psi_z[2, 1] = psi_z[1]*psi_z[2]*rho_z;
+     }
+   model{
+     nonzero ~ bernoulli_logit(eta);           // logit part of model
+     log_exercise[index_nonzero]
+       ~ normal(mu[index_nonzero], sigma); // linear part of model
+
+     for (j in 1:J){
+       B[j, ] ~ multi_normal(rep_vector(0, 2), Psi_b);
+       Z[j, ] ~ multi_normal(rep_vector(0, 2), Psi_z);
+       // distribution of subject-specific coefficients
+     }
+
+     // priors:
+     beta[1] ~ normal(0, 0.5); // general intercept
+     beta[2] ~ normal(0, 0.5); // difference in intercepts
+     beta[3] ~ normal(0, 0.5); // general age slope
+     beta[4] ~ normal(0, 1);   // difference in slopes
+     log_sigma ~ normal(0, 0.4);
+     log_psi_b ~ normal(0, 0.4);
+     z_rho_b ~ normal(0, 0.6);
+
+     zeta[1] ~ normal(0, 1);   // general intercept
+     zeta[2] ~ normal(0, 0.5); // difference in intercepts
+     zeta[3] ~ normal(0, 0.5); // general age slope
+     zeta[4] ~ normal(0, 0.5); // difference in slopes
+     log_psi_z ~ normal(0, 0.2);
+     z_rho_z ~ normal(0, 0.3);
+   }
+ "

> system.time(fit.blackmore <- stan(model_code=blackmore.model, chains=4,
+   model_name= "Hurdle model for Blackmore data, not so weak priors,
+               random intercepts and slopes",
+   seed=234870879, data=data.blackmore, iter=10000))
. . .

   user  system elapsed
52.185   8.099 328.750
>
> print(fit.blackmore,
+   pars=c("beta", "zeta", "sigma", "psi_b", "rho_b", "psi_z", "rho_z"),
+   digits=4)
Inference for Stan model: Hurdle model for Blackmore data, not so weak priors,
  random intercepts and slopes.
4 chains, each with iter=10000; warmup=5000; thin=1;
post-warmup draws per chain=5000, total post-warmup draws=20000.

      mean se_mean      sd   2.5%    25%    50%    75%   97.5%
beta[1] -0.3036  0.0011 0.1055 -0.5111 -0.3742 -0.3038 -0.2321 -0.0977
beta[2] -0.1889  0.0020 0.1995 -0.5815 -0.3227 -0.1892 -0.0540  0.2028
beta[3]  0.2090  0.0001 0.0179  0.1737  0.1970  0.2089  0.2208  0.2441

```

beta[4]	0.2108	0.0002	0.0352	0.1404	0.1873	0.2106	0.2343	0.2790
zeta[1]	2.1796	0.0048	0.2304	1.7553	2.0213	2.1706	2.3277	2.6660
zeta[2]	-0.4772	0.0017	0.3138	-1.0869	-0.6897	-0.4774	-0.2651	0.1379
zeta[3]	0.4112	0.0022	0.1020	0.2229	0.3403	0.4073	0.4770	0.6233
zeta[4]	0.2143	0.0010	0.1299	-0.0372	0.1259	0.2140	0.3003	0.4748
sigma	1.0214	0.0003	0.0346	0.9556	0.9981	1.0207	1.0445	1.0915
psi_b[1]	1.3524	0.0010	0.0927	1.1782	1.2886	1.3503	1.4140	1.5414
psi_b[2]	0.1521	0.0004	0.0197	0.1141	0.1387	0.1518	0.1654	0.1918
rho_b	-0.4787	0.0013	0.0994	-0.6518	-0.5489	-0.4864	-0.4162	-0.2651
psi_z[1]	1.1675	0.0080	0.2094	0.7928	1.0189	1.1578	1.3002	1.6110
psi_z[2]	0.5160	0.0018	0.0739	0.3859	0.4639	0.5115	0.5620	0.6757
rho_z	0.2141	0.0065	0.2056	-0.2022	0.0732	0.2196	0.3620	0.5918
	n_eff	Rhat						
beta[1]	8574	1.0005						
beta[2]	10218	1.0002						
beta[3]	17402	1.0000						
beta[4]	20391	1.0000						
zeta[1]	2291	1.0024						
zeta[2]	32242	1.0000						
zeta[3]	2144	0.9999						
zeta[4]	16249	1.0001						
sigma	11601	0.9999						
psi_b[1]	7920	1.0003						
psi_b[2]	3052	1.0000						
rho_b	6118	1.0005						
psi_z[1]	691	1.0077						
psi_z[2]	1690	1.0008						
rho_z	990	1.0057						

Samples were drawn using NUTS(diag\_e) at Wed Mar 15 15:55:37 2023.  
 For each parameter, n\_eff is a crude measure of effective sample size,  
 and Rhat is the potential scale reduction factor on split chains (at  
 convergence, Rhat=1).

The fixed-effect estimates, in which our interest centers, are very similar to those for the model with only random slopes, reported in the text (cf., Table 25.5 on page 63). As was the case with mixed-effects models estimated by ML or REML, estimates of fixed effects are rarely sensitive to details of the specification of the random effects.

## References

- J. Fox and S. Weisberg. *An R Companion to Applied Regression*. Sage, Thousand Oaks CA, third edition, 2019.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL <https://www.R-project.org/>.