Reed frogs!

`data(reedfrogs)`, Vonesh and Bolker 2005

density	pred	size	surv	propsurv
10	no	big	9	0.9
10	no	big	10	1.0
10	no	big	7	0.7
10	no	big	10	1.0
10	no	small	9	0.9
10	no	small	9	0.9

Group-level effect: tank

- 48 tanks, with or w/o predators
- each tank has a different # of tadpoles (`density`)
- each frog is an observation
  - live (1) / die (0)
- each tank: # of surviving tadpoles (`surv`, `propsurv`)

How to estimate survivorship?

- Pooled model (every tank has the same survivorship)
- No pooling model (every tank has a unique survivorship)
- Partially pooled model (tanks are similar, but with some variation in survivorship)

Pooled model

Estimate mean probability across tanks:

```
d$tank <- 1:nrow(d) # tank group variable
dat <- d %>% rename(S = surv, N = density) %>%
  select(S, N, tank)
m_pool <- ulam(
  alist(
    S ~ dbinom(N, p), # likelihood
    logit(p) <- a, # process model
    a ~ dnorm(0, 1.5) # prior
  ), data = dat, chains = 2, log_lik = TRUE)
```

No pooling model

Estimate unique probability for each tank:

```
d$tank <- 1:nrow(d) # tank group variable
dat <- d %>% rename(S = surv, N = density) %>%
  select(S, N, tank)
m_nopool <- ulam(
  alist(
    S ~ dbinom(N, p), # likelihood
    logit(p) <- a[tank], # process model
    a[tank] ~ dnorm(0, 1.5) # prior
  ), data = dat, chains = 2, log_lik = TRUE)
```

Partial pooling and hyperparameters

No pooling model:

$\alpha[tank] \sim \text{Normal}(0, 1.5)$

Partial pooling and hyperparameters

No pooling model:

$\alpha[tank] \sim \text{Normal}(0, 1.5)$

Partial pooling model:

$\alpha[tank] \sim \text{Normal}(\alpha_{tank}, \sigma_{tank})$

- $\alpha[tank]$ is drawn from a distribution
- $\alpha_{tank}, \sigma_{tank}$ are **hyperparameters**, each with its own prior

Partial pooling model

```r
d$tank <- 1:nrow(d) # tank group variable
dat <- d %>% rename(S = surv, N = density) %>%
  select(S, N, tank)
m_partialpool <- ulam(
  alist(
    S ~ dbinom(N, p), # likelihood
    logit(p) <- a[tank], # process model
    a[tank] ~ dnorm(mu_tank, sigma_tank), # group intercepts
    mu_tank ~ dnorm(0, 1.5), #  hyperprior
    sigma_tank ~ dexp(1) # hyperprior
  ), data = dat, chains = 2, log_lik = TRUE)
```

MCMC checks

```
#>             mean   sd 5.5% 94.5%   n_eff Rhat4
#> mu_tank    1.34 0.26 0.92  1.77 1416.24     1
#> sigma_tank 1.62 0.22 1.32  2.01  814.86     1
```

MCMC checks

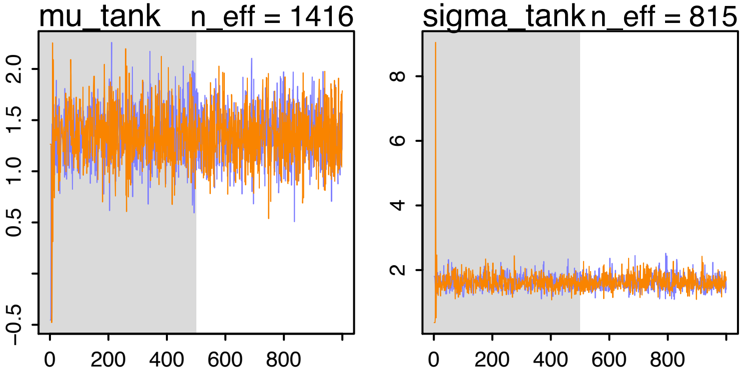
mu_tank
n_eff = 1416
sigma_tank
n_eff = 815
−0.5
0.5
1.0
1.5
2.0
2
4
6
8
0
200
400
600
800

Shrinkage in frog survivorship

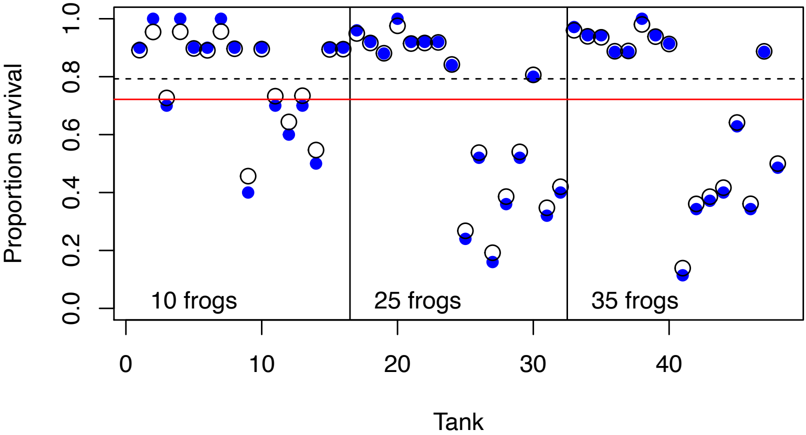

- black dashed line: population mean
- red dashed line: raw mean

Compare using `WAIC`

```
#>                  WAIC    SE  dWAIC   dSE pWAIC weight
#> m_partialpool 202.30  7.36   0.00    NA 21.97      1
#> m_nopool      214.23  4.62  11.93  4.27 25.32      0
#> m_pool        586.34 68.92 384.04 67.26 12.18      0
```

- improved fit from partial pooling
- multilevel model has fewer *effective* parameters than nopool model!
  - `m_nopool` estimated 48 intercepts
  - `m_partialpool` estimated 48 intercepts, plus 2 hyperparameters . . . but **stronger** prior!

Adaptive regularization resulted in stronger prior

```
#>             mean   sd
#> mu_tank    1.34 0.26
#> sigma_tank 1.62 0.22
```

```r
post <- extract.samples(m_partialpool)
sim_partial <- rnorm(1e4, post$mu_tank, post$sigma_tank)
sim_nopool <- rnorm(1e4, 0, 1.5)
dens(inv_logit(sim_partial), lwd = 2, adj = 0.1)
dens(inv_logit(sim_nopool), lwd = 2, adj = 0.1,
     col = "red", add = TRUE)
legend(x = 0, y = 3.5, col = c("black", "red"),
       legend = c("Partial pooling", "No pooling"),
       lty = 1, lwd = 2, bty = "n")
```

Adaptive regularization resulted in stronger prior

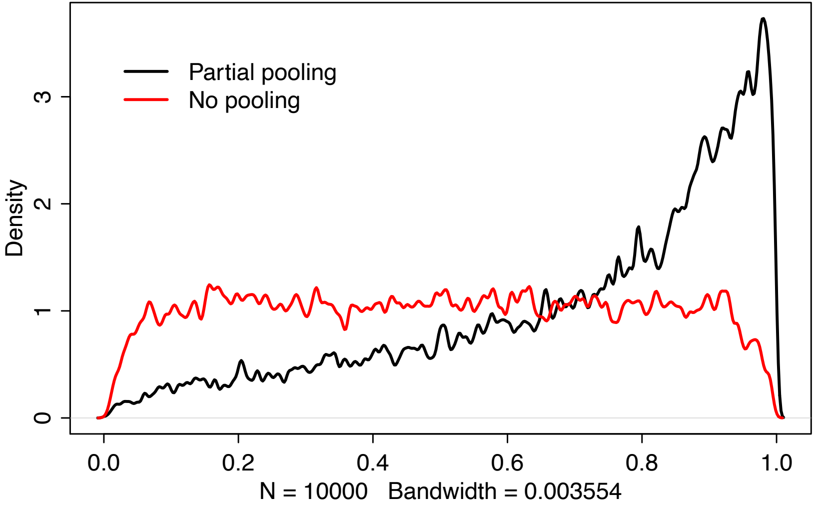

Question

Which of the following priors will produce more shrinkage in the estimates?

(a) $\alpha[tank] \sim \text{Normal}(0, 1)$

(b) $\alpha[tank] \sim \text{Normal}(0, 2)$

Attribution

This slide deck is based on Chapter 13 of Statistical Rethinking.

References

Hobbs, N Thompson, and Mevin B Hooten. 2015. *Bayesian Models: A Statistical Primer for Ecologists*. Princeton University Press.

McElreath, Richard. 2020. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Press.