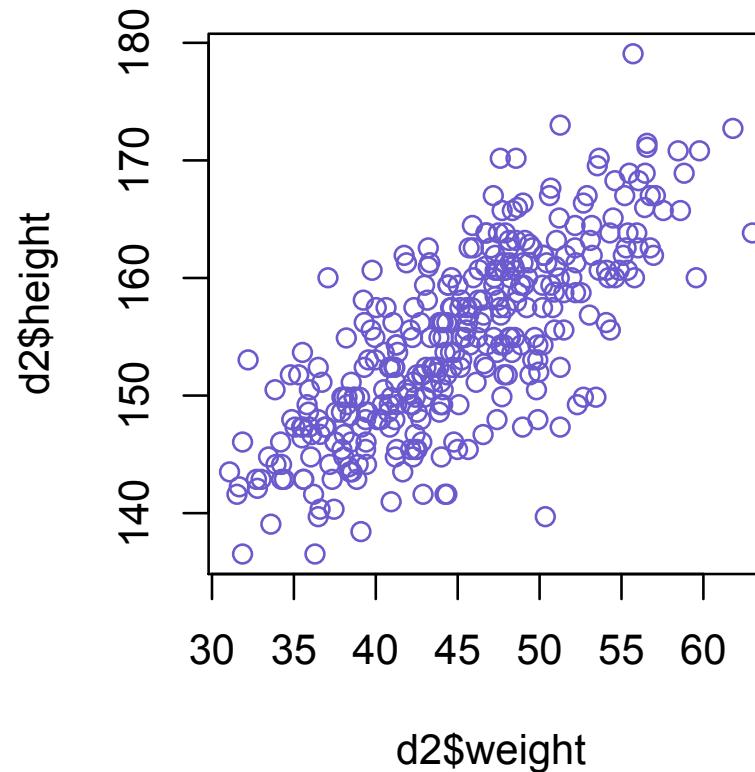


Adding a predictor variable

- How does weight describe height?



Adding a predictor variable

- Use a linear model of the mean, μ :

$$h_i \sim \text{Normal}(\mu_i, \sigma) \quad [\text{likelihood}]$$

$$\mu_i = \alpha + \beta(x_i - \bar{x}) \quad [\text{linear model}]$$

$$\alpha \sim \text{Normal}(178, 20) \quad [\alpha \text{ prior}]$$

$$\beta \sim \text{Normal}(0, 10) \quad [\beta \text{ prior}]$$

$$\sigma \sim \text{Uniform}(0, 50) \quad [\sigma \text{ prior}]$$

Adding a predictor variable

$$h_i \sim \text{Normal}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta(x_i - \bar{x})$$

mean on row i

mean when $x_i = 0$
“intercept”

change in mean, per unit change x_i
“slope”

mean weight

weight on row i

The diagram illustrates the decomposition of the regression equation $\mu_i = \alpha + \beta(x_i - \bar{x})$ into its components. A dashed box encloses the entire equation. Arrows point from each term to its corresponding interpretation: the constant term α is labeled "mean when $x_i = 0$ " and "intercept"; the coefficient β is labeled "change in mean, per unit change x_i " and "slope"; and the error term $(x_i - \bar{x})$ is labeled "mean on row i ". A separate dashed box encloses the error term $h_i \sim \text{Normal}(\mu_i, \sigma)$, with arrows pointing to "mean weight" and "weight on row i ".

Prior predictive distribution

- What do these priors mean?
- Let's simulate to find out!

$$h_i \sim \text{Normal}(\mu_i, \sigma) \quad [\text{likelihood}]$$

$$\mu_i = \alpha + \beta(x_i - \bar{x}) \quad [\text{linear model}]$$

$$\alpha \sim \text{Normal}(178, 20) \quad [\alpha \text{ prior}]$$

$$\beta \sim \text{Normal}(0, 10) \quad [\beta \text{ prior}]$$

$$\sigma \sim \text{Uniform}(0, 50) \quad [\sigma \text{ prior}]$$

Prior predictive distribution

```
set.seed(2971)
N <- 100                      # 100 lines
a <- rnorm( N , 178 , 20 )
b <- rnorm( N , 0 , 10 )
```

$$\begin{aligned}\alpha &\sim \text{Normal}(178, 20) \\ \beta &\sim \text{Normal}(0, 10)\end{aligned}$$

R code
4.38

Prior predictive distribution

```
set.seed(2971)
N <- 100 # 100 lines
a <- rnorm( N , 178 , 20 )
b <- rnorm( N , 0 , 10 )
```

R code
4.38

$$\alpha \sim \text{Normal}(178, 20)$$
$$\beta \sim \text{Normal}(0, 10)$$

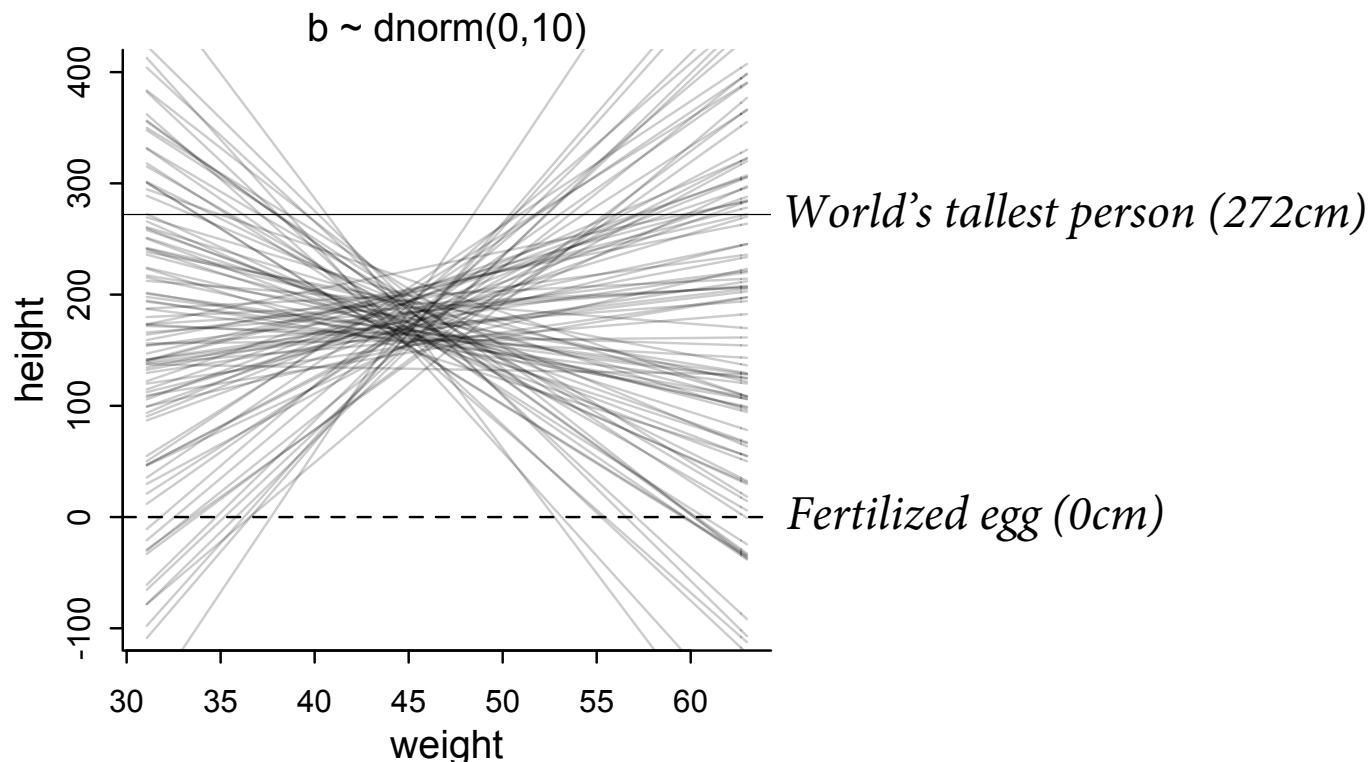


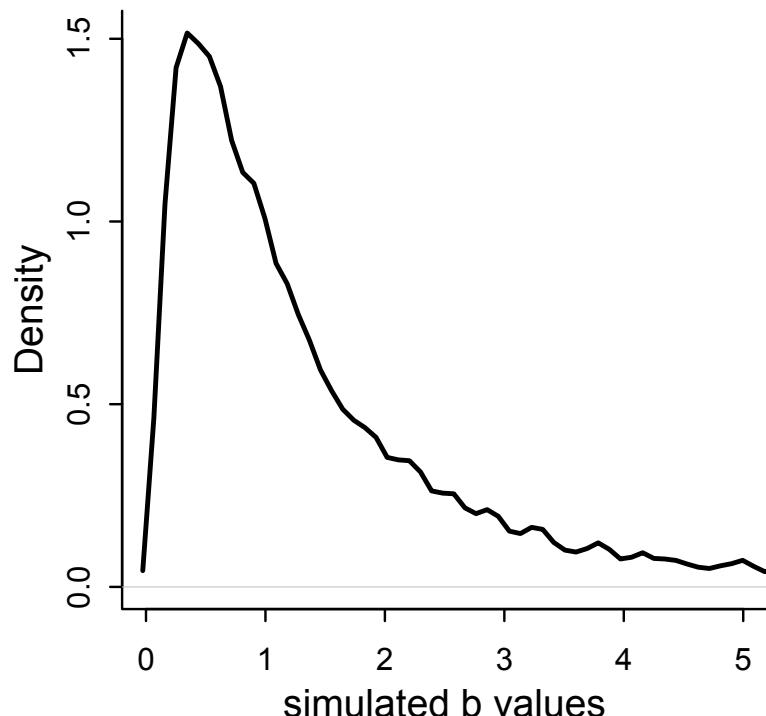
Figure 4.5

Prior predictive distribution

R code
4.40

```
b <- rlnorm( 1e4 , 0 , 1 )  
dens( b , xlim=c(0,5) , adj=0.1 )
```

$$\beta \sim \text{Log-Normal}(0, 1)$$



Prior predictive distribution

R code
4.41

```
set.seed(2971)
N <- 100 # 100 lines
a <- rnorm( N , 178 , 20 )
b <- rlnorm( N , 0 , 1 )
```

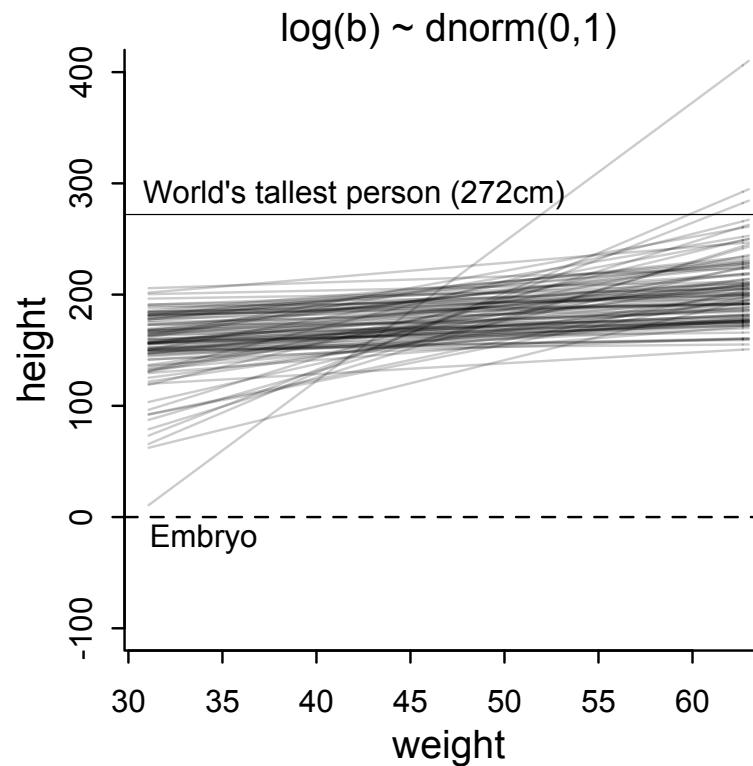
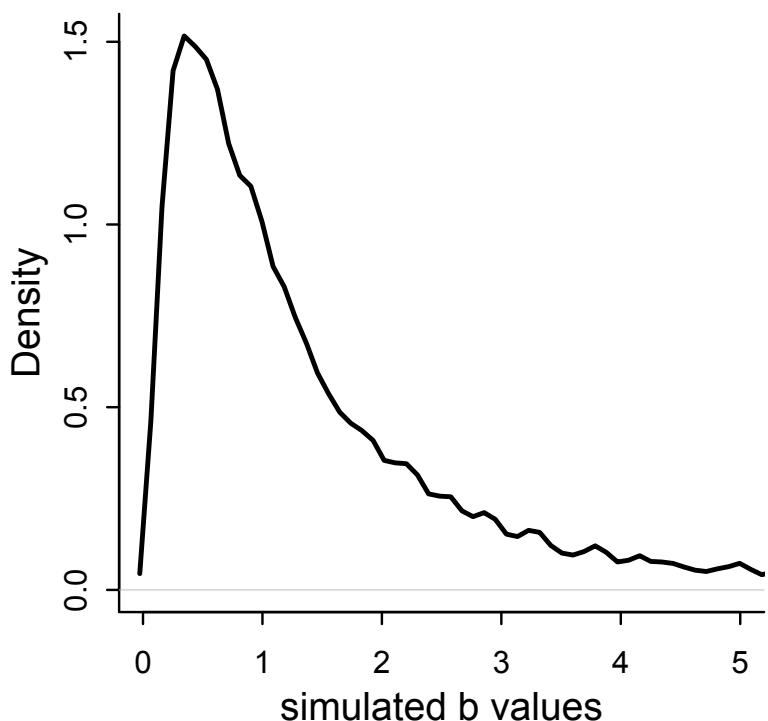


Figure 4.5

Approximate the posterior

R code
4.42

```
# load data again, since it's a long way back
library(rethinking)
data(Howell1)
d <- Howell1
d2 <- d[ d$age >= 18 , ]

# define the average weight, x-bar
xbar <- mean(d2$weight)

# fit model
m4.3 <- quap(
  alist(
    height ~ dnorm( mu , sigma ) ,
    mu <- a + b*( weight - xbar ) ,
    a ~ dnorm( 178 , 20 ) ,
    b ~ dlnorm( 0 , 1 ) ,
    sigma ~ dunif( 0 , 50 )
  ) ,
  data=d2 )
```

$$h_i \sim \text{Normal}(\mu_i, \sigma)$$

$$\mu_i = \alpha + \beta(x_i - \bar{x})$$

$$\alpha \sim \text{Normal}(178, 20)$$

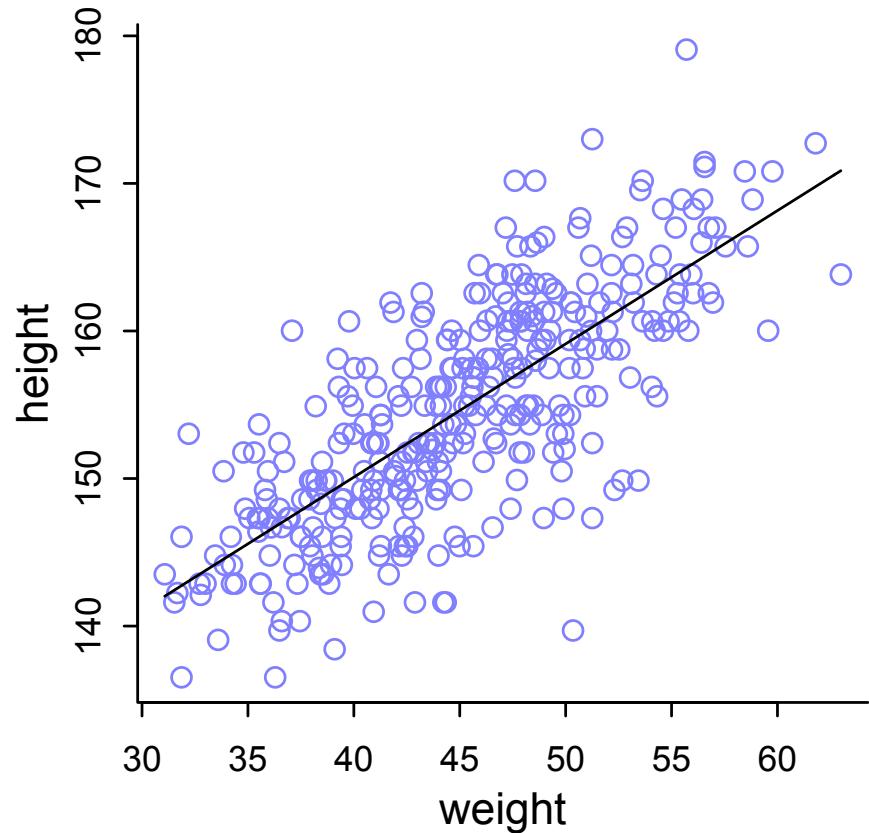
$$\beta \sim \text{Log-Normal}(0, 1)$$

$$\sigma \sim \text{Uniform}(0, 50)$$

R code
4.44

```
precis( m4.3 )
```

	mean	sd	5.5%	94.5%
a	154.60	0.27	154.17	155.03
b	0.90	0.04	0.84	0.97
sigma	5.07	0.19	4.77	5.38



R code
4.46

```
plot( height ~ weight , data=d2 , col=rangi2 )
post <- extract.samples( m4.3 )
a_map <- mean(post$a)
b_map <- mean(post$b)
curve( a_map + b_map*(x - xbar) , add=TRUE )
```

Figure 4.6

Sampling from the posterior

- Want to get uncertainty onto that graph
- Again, sample from posterior
 1. Use mean and standard deviation to approximate posterior
 2. Sample from *multivariate normal* distribution of parameters
 3. Use samples to generate predictions that “integrate over” the uncertainty

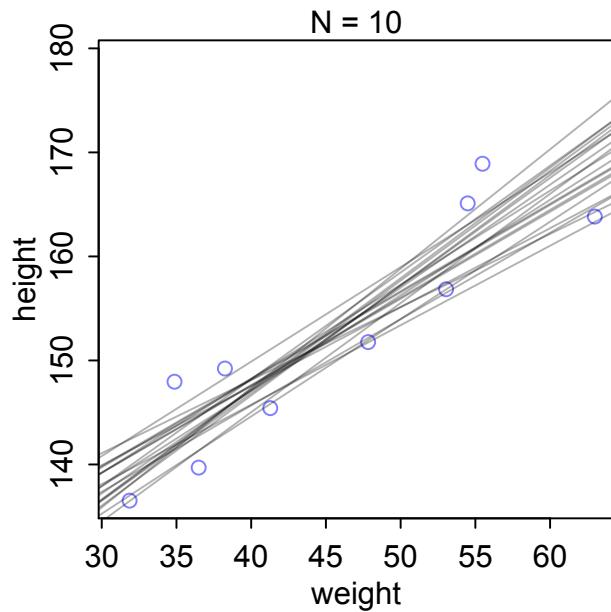
Sampling from the posterior

R code
4.47

```
post <- extract.samples( m4.3 )
post[1:5,]
```

	a	b	sigma
1	154.5505	0.9222372	5.188631
2	154.4965	0.9286227	5.278370
3	154.4794	0.9490329	4.937513
4	155.2289	0.9252048	4.869807
5	154.9545	0.8192535	5.063672

Posterior is full of lines



R code
4.47

```
post <- extract.samples( m4.3 )
post[1:5,]
```

	a	b	sigma
1	154.5505	0.9222372	5.188631
2	154.4965	0.9286227	5.278370
3	154.4794	0.9490329	4.937513
4	155.2289	0.9252048	4.869807
5	154.9545	0.8192535	5.063672

Figure 4.7

Posterior is full of lines

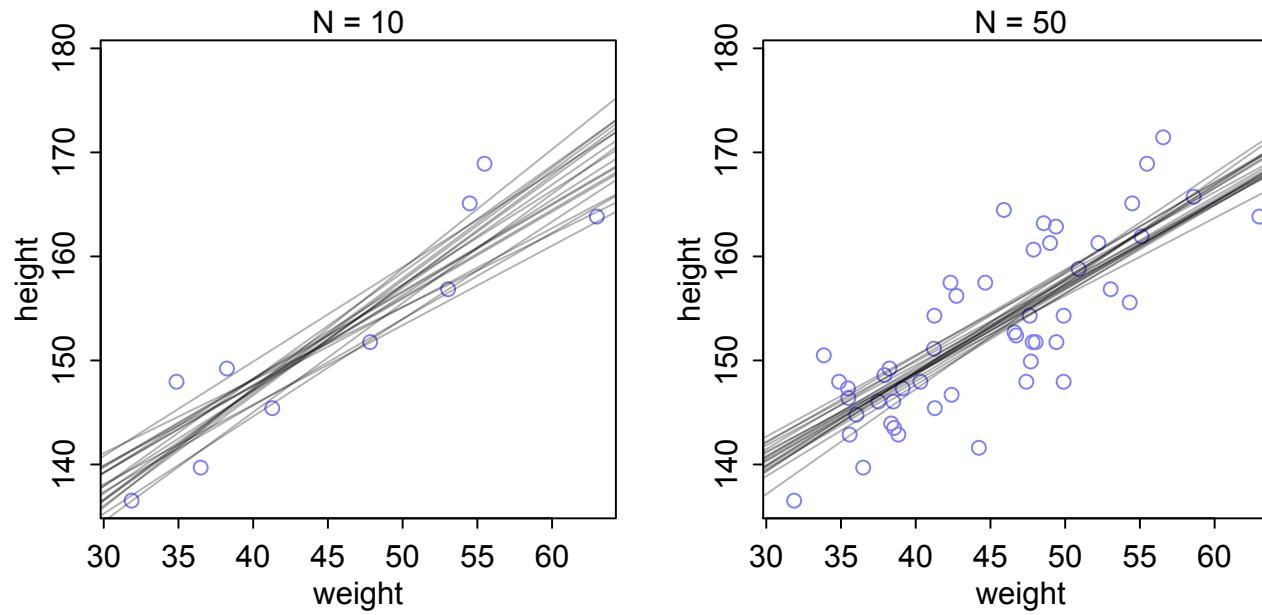


Figure 4.7

Posterior is full of lines

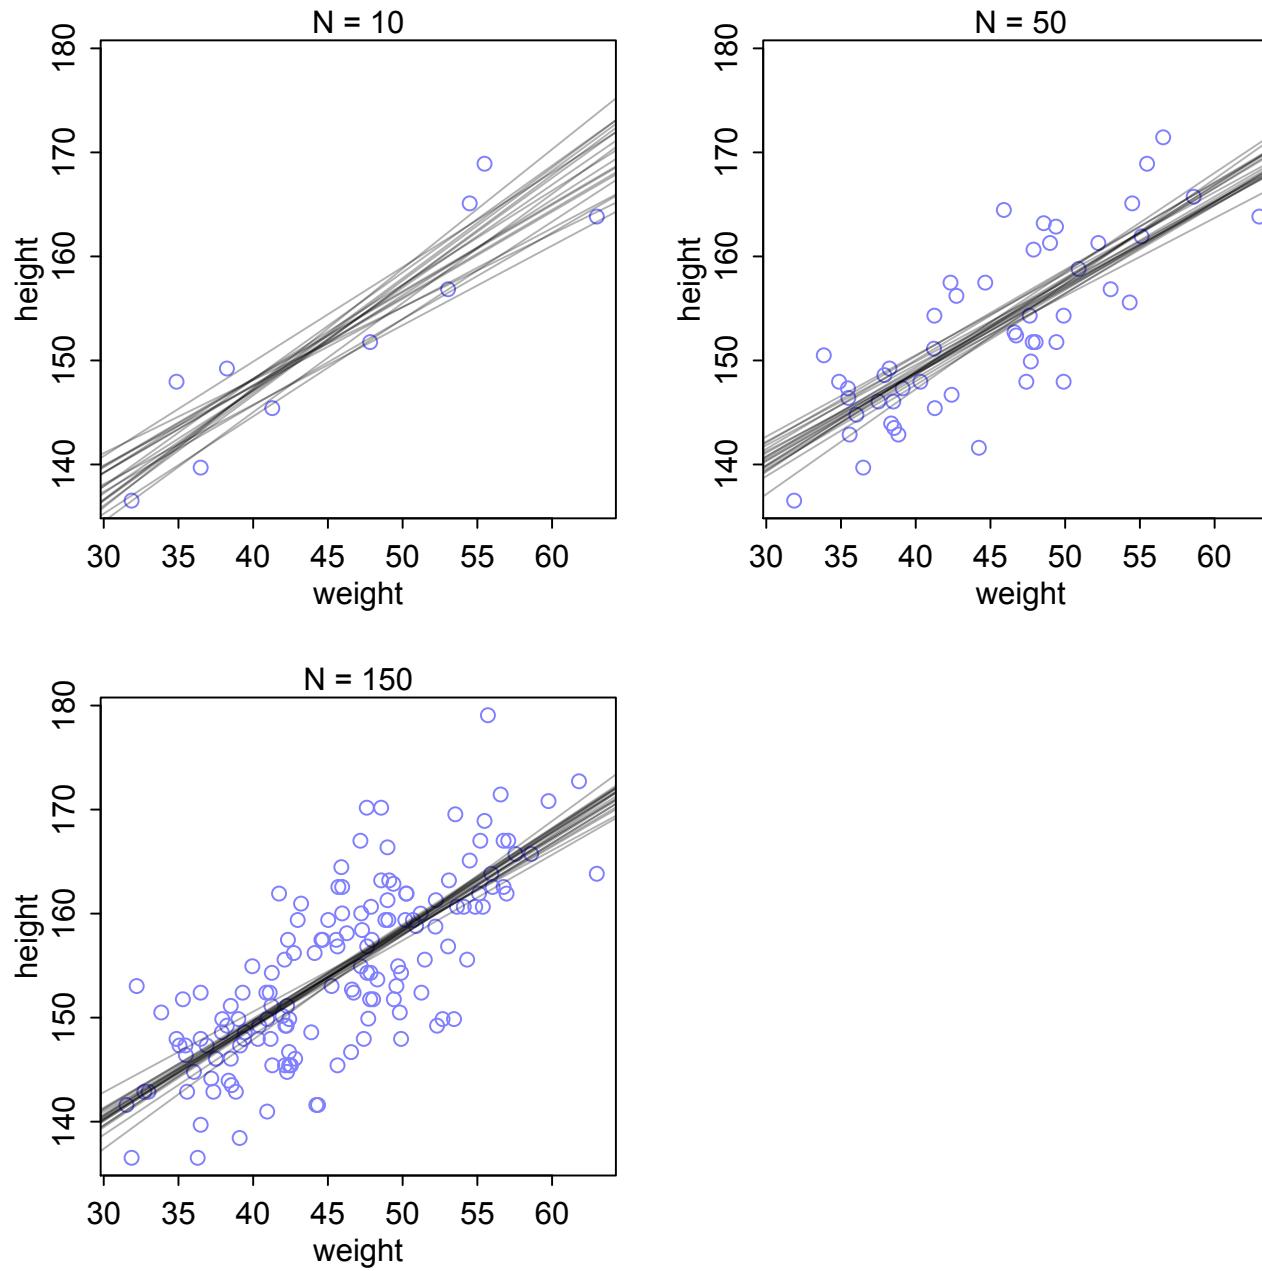


Figure 4.7

Posterior is full of lines

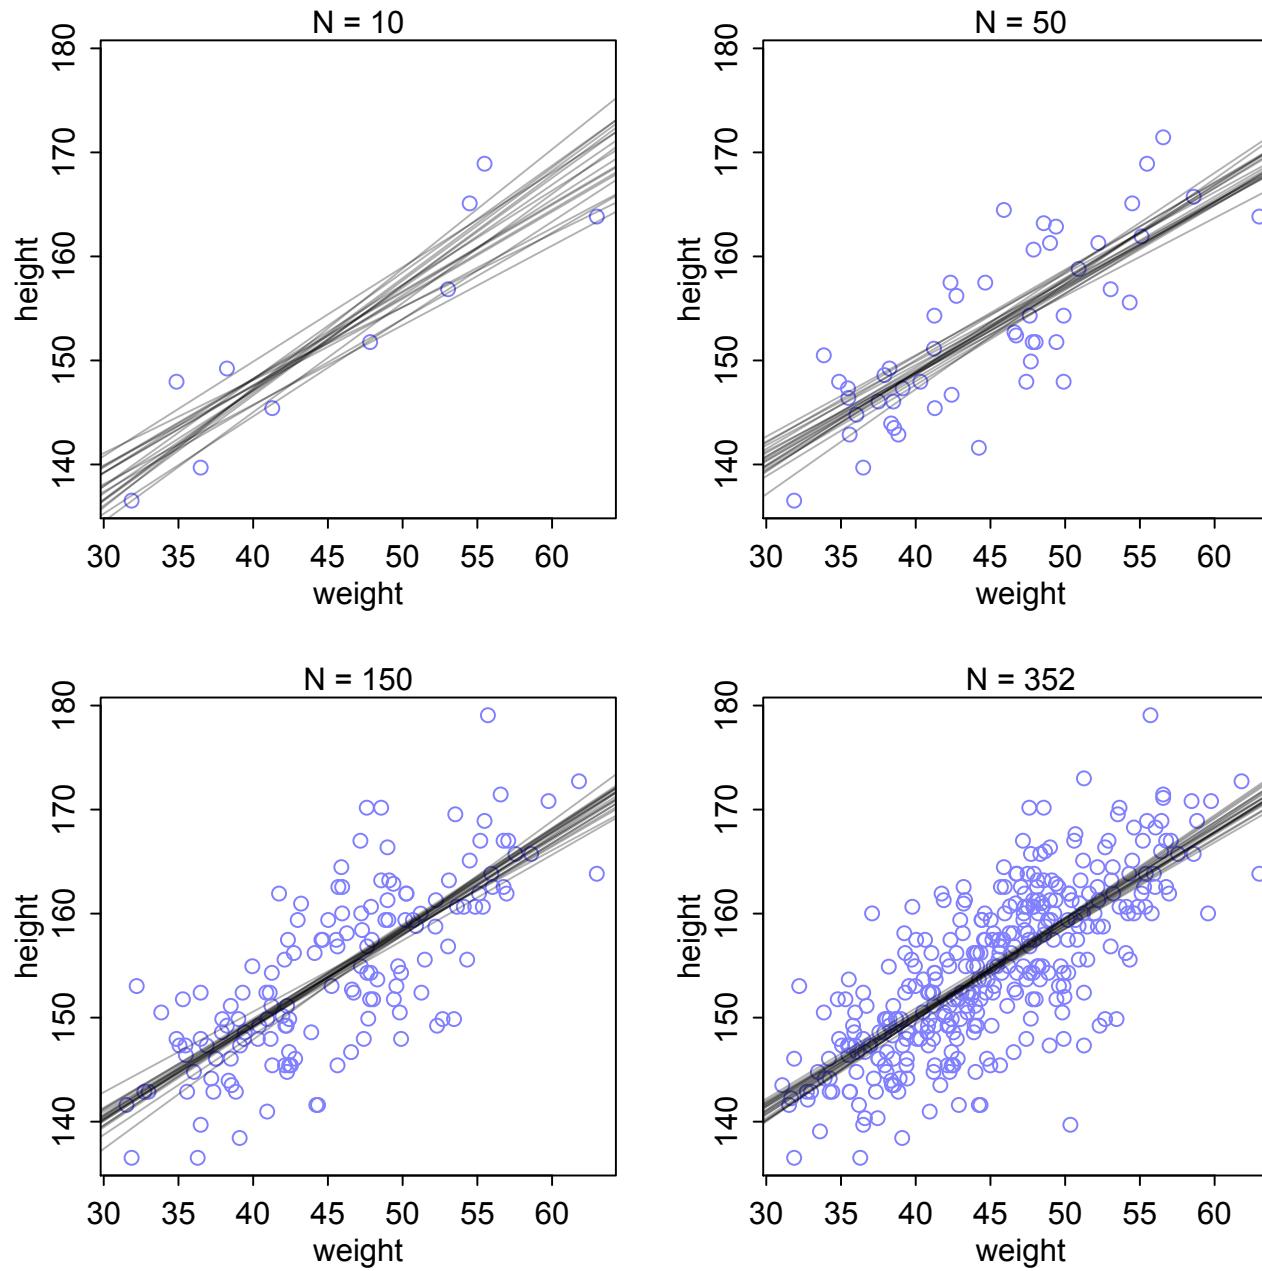


Figure 4.7