

Diagnosics and Model Evaluation


Lecture 03

Dr. Colin Rundel

Some more linear models

Linear model and data

```
1 ggplot(d, aes(x=x,y=y)) +  
2   geom_point() +  
3   geom_smooth(method="lm", color="blue", se = FALSE)
```



Linear model

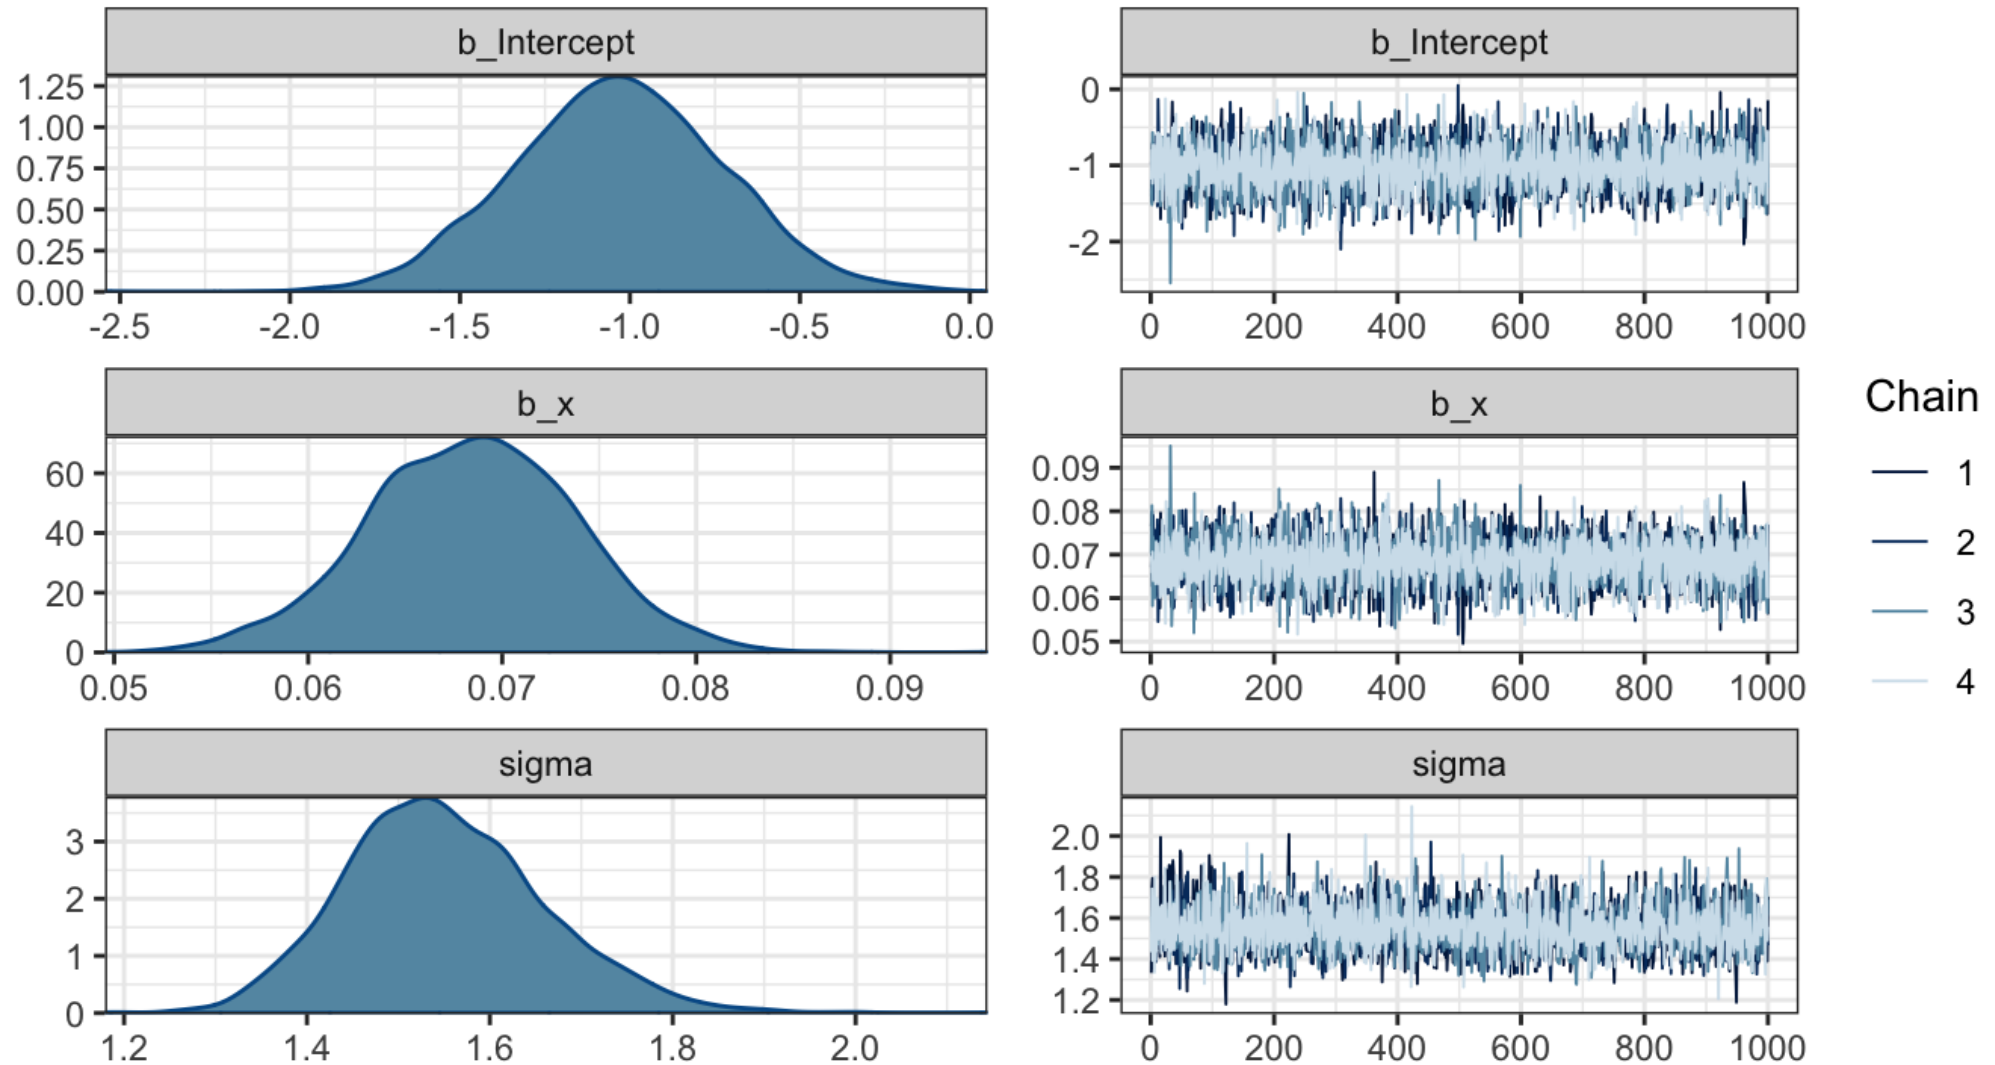
```
1 l = lm(y ~ x, data=d)
2 summary(l)
3 ##
4 ## Call:
5 ## lm(formula = y ~ x, data = d)
6 ##
7 ## Residuals:
8 ##      Min       1Q   Median       3Q      Max
9 ## -2.6041 -1.2142 -0.1973  1.1969  3.7072
10 ##
11 ## Coefficients:
12 ##              Estimate Std. Error t value Pr(>|t|)
13 ## (Intercept) -1.030315   0.310326  -3.32  0.00126 **
14 ## x           0.068409   0.005335  12.82 < 2e-16 ***
15 ## ---
16 ## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
17 ##
18 ## Residual standard error: 1.54 on 98 degrees of freedom
19 ## Multiple R-squared:  0.6266, Adjusted R-squared:  0.6227
20 ## F-statistic: 164.4 on 1 and 98 DF,  p-value: < 2.2e-16
```

Bayesian model (brms)

```
1 ( b = brms::brm(
2   y ~ x, data=d,
3   prior = c(
4     brms::prior(normal(0, 100), class = Intercept),
5     brms::prior(normal(0, 10), class = b),
6     brms::prior(cauchy(0, 2), class = sigma)
7   ),
8   silent = 2, refresh = 0
9 )
10 )
11 ## Running /opt/homebrew/Cellar/r/4.2.1_2/lib/R/bin/R CMD SHLIB foo.c
12 ## clang -I"/opt/homebrew/Cellar/r/4.2.1_2/lib/R/include" -DNDEBUG -I"/Users/rundel/Library/R/
13 ## In file included from <built-in>:1:
14 ## In file included from /Users/rundel/Library/R/arm64/4.2/library/StanHeaders/include/stan/math/
15 ## In file included from /Users/rundel/Library/R/arm64/4.2/library/RcppEigen/include/Eigen/Dense
16 ## In file included from /Users/rundel/Library/R/arm64/4.2/library/RcppEigen/include/Eigen/Core:
17 ## /Users/rundel/Library/R/arm64/4.2/library/RcppEigen/include/Eigen/src/Core/util/Macros.h:628:
18 ## namespace Eigen {
19 ## ^
20 ## /Users/rundel/Library/R/arm64/4.2/library/RcppEigen/include/Eigen/src/Core/util/Macros.h:628:
21 ## namespace Eigen {
22 ## ^
```

Parameter estimates

```
1 plot(b)
```



tidybayes - gather_draws (long)


```
1 b_post = b |>
2   tidybayes::gather_draws(b_Intercept, b_x, sigma)
3
4 b_post
5 ## # A tibble: 12,000 × 5
6 ## # Groups:   .variable [3]
7 ##   .chain .iteration .draw .variable .value
8 ##   <int>     <int> <int> <chr>     <dbl>
9 ## 1         1         1     1 b_Intercept -0.961
10 ## 2         1         2     2 b_Intercept -1.24
11 ## 3         1         3     3 b_Intercept -1.29
12 ## 4         1         4     4 b_Intercept -0.763
13 ## 5         1         5     5 b_Intercept -0.980
14 ## 6         1         6     6 b_Intercept -1.13
15 ## 7         1         7     7 b_Intercept -0.810
16 ## 8         1         8     8 b_Intercept -1.33
```

tidybayes - spread_draws (wide)

```
1 b_post_wide = b |>
2   tidybayes::spread_draws(b_Intercept, b_x, sigma)
3
4 b_post_wide
5 ## # A tibble: 4,000 × 6
6 ##   .chain .iteration .draw b_Intercept    b_x sigma
7 ##   <int>     <int> <int>      <dbl>  <dbl> <dbl>
8 ## 1         1     1         1    -0.961  0.0671  1.33
9 ## 2         1     2         2    -1.24   0.0726  1.75
10 ## 3         1     3         3    -1.29   0.0736  1.79
11 ## 4         1     4         4    -0.763  0.0639  1.35
12 ## 5         1     5         5    -0.980  0.0697  1.53
13 ## 6         1     6         6    -1.13   0.0692  1.53
14 ## 7         1     7         7    -0.810  0.0637  1.52
15 ## 8         1     8         8    -1.33   0.0726  1.47
16 ## 9         1     9         9    -1.22   0.0728  1.65
```



Posterior plots

```
1 b_post |>
2   ggplot(aes(fill=as.factor(.chain), group=.chain, x=.value)) +
3   geom_density(alpha=0.33, color=NA) +
4   facet_wrap(~.variable, scales = "free")
```



Trace plots

```
1 b_post %>%  
2   ggplot(aes(x=.iteration, y=.value, color=as.factor(.chain))) +  
3   geom_line(alpha=0.5) +  
4   facet_grid(.variable~.chain, scale="free_y") +  
5   geom_smooth(method="loess") + labs(color="chain")
```

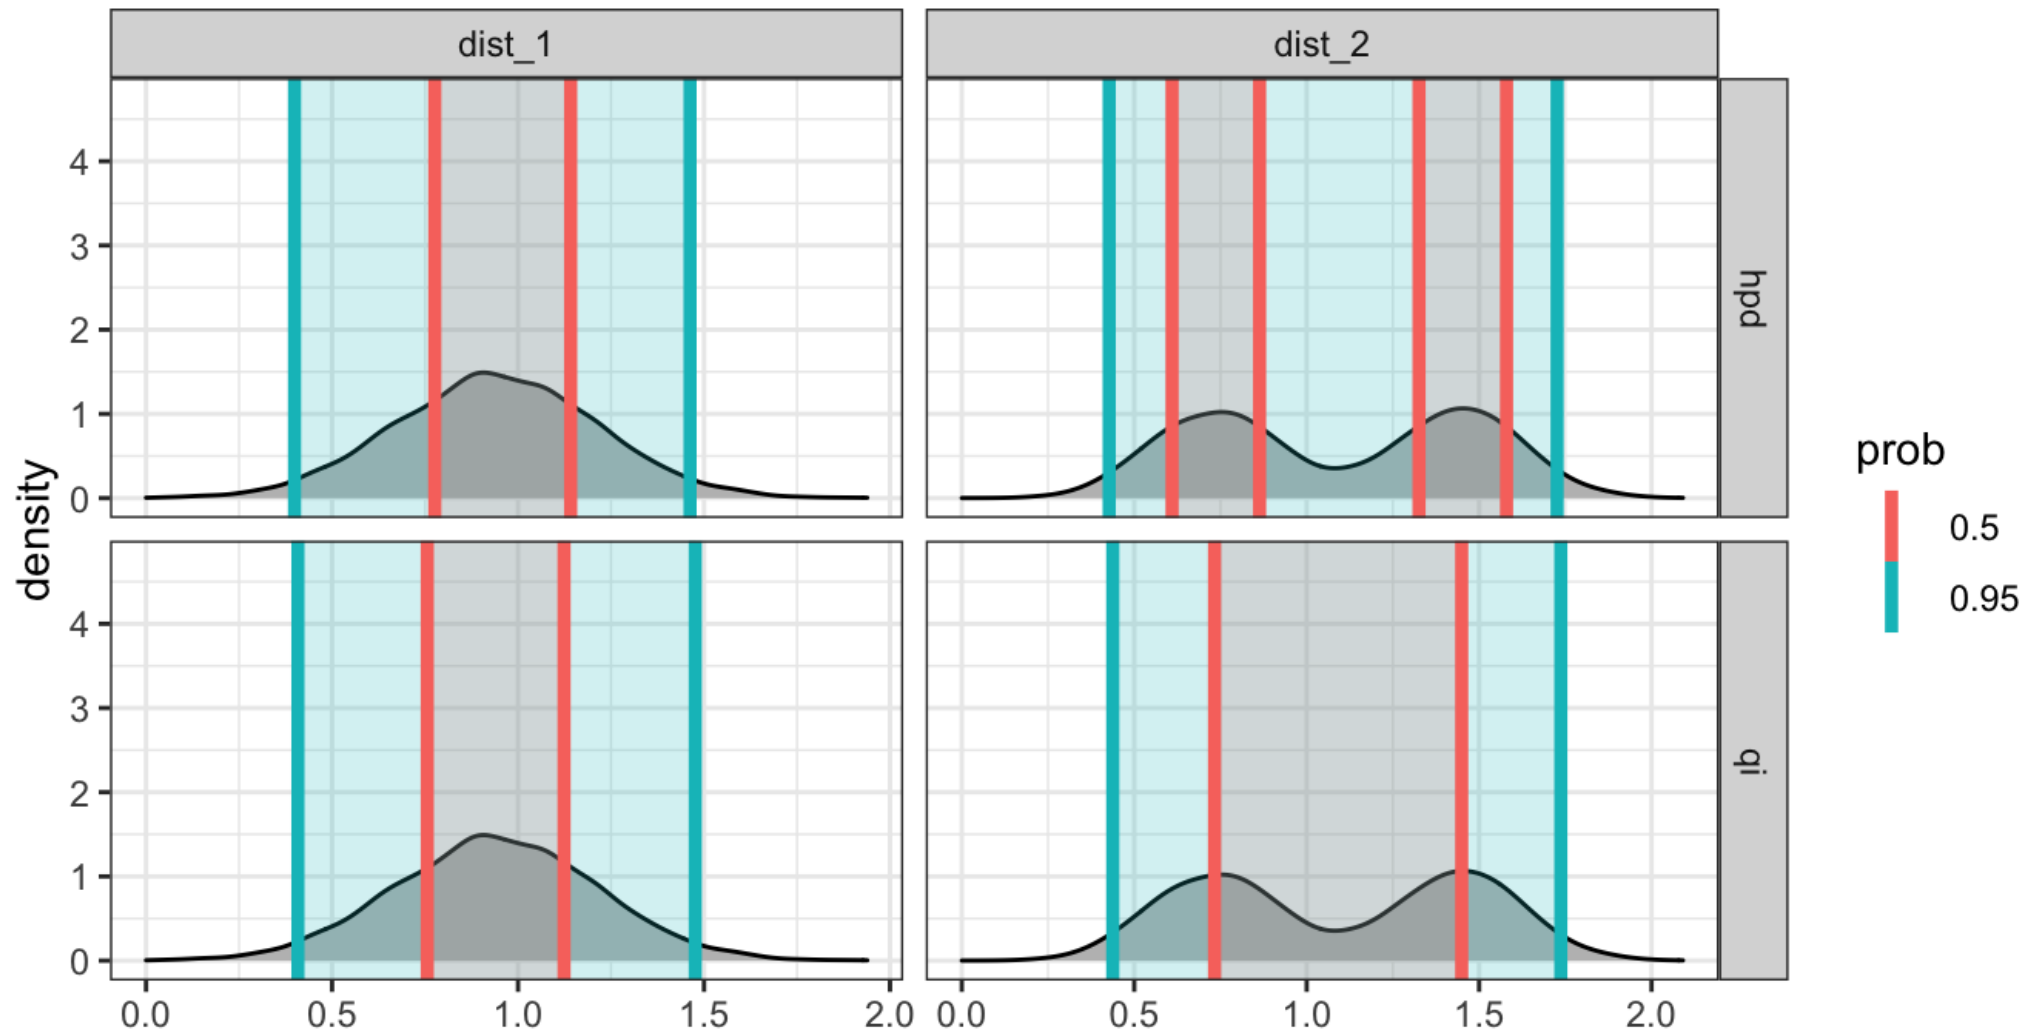


Credible Intervals

```
1 ( b_ci = b_post |>
2   group_by(.chain, .variable) |>
3   tidybayes::mean_hdi(.value, .width=c(0.8, 0.95))
4 )
5 ## # A tibble: 24 × 8
6 ##   .chain .variable   .value .lower .upper .width .point .interval
7 ##   <int> <chr>         <dbl> <dbl> <dbl> <dbl> <chr> <chr>
8 ## 1     1 1 b_Intercept -1.03  -1.40  -0.624  0.8 mean  hdi
9 ## 2     1 1 b_x          0.0685 0.0624 0.0761  0.8 mean  hdi
10 ## 3     1 1 sigma       1.55   1.40   1.69   0.8 mean  hdi
11 ## 4     2 2 b_Intercept -1.03  -1.41  -0.621  0.8 mean  hdi
12 ## 5     2 2 b_x          0.0684 0.0620 0.0753  0.8 mean  hdi
13 ## 6     2 2 sigma       1.55   1.43   1.69   0.8 mean  hdi
14 ## 7     3 3 b_Intercept -1.02  -1.40  -0.604  0.8 mean  hdi
15 ## 8     3 3 b_x          0.0684 0.0615 0.0749  0.8 mean  hdi
16 ## 9     3 3 sigma       1.55   1.42   1.69   0.8 mean  hdi
17 ## 10    4 4 b_Intercept -1.04  -1.40  -0.637  0.8 mean  hdi
18 ## # ... with 14 more rows
```


Aside - `mean_qi()` vs `mean_hdi()`

These differ in the use of the quantile interval vs. the highest-density interval.



Caterpillar Plots

```
1 b_ci %>%  
2   ggplot(aes(x=.value, y=.chain, color=as.factor(.chain))) +  
3   facet_grid(.variable ~ .) +  
4   tidybayes::geom_pointintervalh() +  
5   ylim(0.5, 4.5)
```




Predictions

lm predictions

```
1 (l_pred = broom::augment(l, interval="confidence"))
2 ## # A tibble: 100 × 10
3 ##       y       x .fitted .lower   .upper .resid   .hat .sigma .cooksd .std.re
4 ##   <dbl> <int>   <dbl> <dbl>   <dbl> <dbl>   <dbl> <dbl>   <dbl>   <dbl>
5 ## 1 -3.24     1 -0.962 -1.57  -0.355  -2.28  0.0394  1.53  0.0467  -1.
6 ## 2 -2.00     2 -0.893 -1.49  -0.296  -1.11  0.0382  1.54  0.0107  -0.
7 ## 3 -2.59     3 -0.825 -1.41  -0.237  -1.77  0.0371  1.54  0.0264  -1.
8 ## 4 -3.07     4 -0.757 -1.34  -0.177  -2.31  0.0359  1.53  0.0436  -1.
9 ## 5 -1.88     5 -0.688 -1.26  -0.118  -1.20  0.0348  1.54  0.0113  -0.
10 ## 6 -0.807    6 -0.620 -1.18  -0.0583 -0.187  0.0338  1.55  0.000266 -0.
11 ## 7 -2.09     7 -0.551 -1.10   0.00127 -1.54  0.0327  1.54  0.0175  -1.
12 ## 8 -0.227    8 -0.483 -1.03   0.0609   0.256  0.0317  1.55  0.000466   0.
13 ## 9  0.333    9 -0.415 -0.950  0.121    0.747  0.0307  1.55  0.00384   0.
14 ## 10  1.98    10 -0.346 -0.873  0.180    2.32  0.0297  1.53  0.0358   1.
15 ## # ... with 90 more rows
```

Confidence interval

```
1 l_pred |>  
2   ggplot(aes(x=x,y=y)) +  
3     geom_point() +  
4     geom_line(aes(y=.fitted), col="blue", size=1.5) +  
5     geom_ribbon(aes(ymin=.lower, ymax=.upper), col="blue", fill="blue", a
```




brms predictions

```
1 (b_pred = predict(b))
2 ##           Estimate Est.Error           Q2.5           Q97.5
3 ## [1,] -0.96419044   1.553634 -4.04028547  2.124732
4 ## [2,] -0.89721849   1.593799 -3.96603496  2.179769
5 ## [3,] -0.79179315   1.594565 -3.96413800  2.391643
6 ## [4,] -0.73859326   1.579550 -3.83546564  2.266440
7 ## [5,] -0.68293418   1.579669 -3.79089154  2.460559
8 ## [6,] -0.61155250   1.609090 -3.81348198  2.516711
9 ## [7,] -0.56675460   1.579506 -3.66369377  2.479165
10 ## [8,] -0.45334581   1.570071 -3.47909432  2.620280
11 ## [9,] -0.47151098   1.565524 -3.49855796  2.580172
12 ## [10,] -0.35537333   1.552547 -3.30957100  2.739897
13 ## [11,] -0.26908970   1.562577 -3.30266874  2.745822
14 ## [12,] -0.20382232   1.582830 -3.19318502  2.973233
15 ## [13,] -0.18236463   1.568774 -3.25612849  3.056613
16 ## [14,] -0.03840721   1.571181 -3.10470573  3.032293
```

Credible interval

```
1 d |>
2   bind_cols(b_pred) |>
3   ggplot(aes(x=x,y=y)) +
4     geom_point() +
5     geom_line(aes(y=Estimate), col="red", size=1.5) +
6     geom_ribbon(aes(ymin=Q2.5, ymax=Q97.5), col='red', fill='red', alpha=0.5)
```



Why are the intervals different?

Raw predictions

```
1 dim( brms::posterior_predict(b) )  
2 ## [1] 4000 100  
3 dim( brms::posterior_epred(b) )  
4 ## [1] 4000 100
```


Tidy raw predictions

```
1 ( b_post_pred = tidybayes::predicted_draws(
2   b, newdata=d
3 ) )
4 ## # A tibble: 400,000 × 7
5 ## # Groups:   x, y, .row [100]
6 ##       x     y .row .chain .iteration .draw
7 ##   <int> <dbl> <int> <int>     <int> <int>
8 ## 1     1 -3.24     1     NA       NA
9 ## 2     1 -3.24     1     NA       NA
10 ## 3     1 -3.24     1     NA       NA
11 ## 4     1 -3.24     1     NA       NA
12 ## 5     1 -3.24     1     NA       NA
13 ## 6     1 -3.24     1     NA       NA
14 ## 7     1 -3.24     1     NA       NA
15 ## 8     1 -3.24     1     NA       NA
16 ## 9     1 -3.24     1     NA       NA
17 ## 10    1 -3.24     1     NA       NA
18 ## # ... with 399,990 more rows
```


```
1 ( b_post_epred = tidybayes::epred_draws(
2   b, newdata=d
3 ) )
4 ## # A tibble: 400,000 × 7
5 ## # Groups:   x, y, .row [100]
6 ##       x     y .row .chain .iteration .draw
7 ##   <int> <dbl> <int> <int>     <int> <int>
8 ## 1     1 -3.24     1     NA       NA
9 ## 2     1 -3.24     1     NA       NA
10 ## 3     1 -3.24     1     NA       NA
11 ## 4     1 -3.24     1     NA       NA
12 ## 5     1 -3.24     1     NA       NA
13 ## 6     1 -3.24     1     NA       NA
14 ## 7     1 -3.24     1     NA       NA
15 ## 8     1 -3.24     1     NA       NA
16 ## 9     1 -3.24     1     NA       NA
17 ## 10    1 -3.24     1     NA       NA
18 ## # ... with 399,990 more rows
```

Posterior predictions vs Expected posterior predictions

```
1 b_post_pred |>
2   filter(.draw <= 25) |>
3   ggplot(aes(x=x,y=y)) +
4     geom_point() +
5     geom_line(
6       aes(y=.prediction, group=.draw),
7       alpha=0.33
8     )
```





```
1 b_post_epred |>
2   filter(.draw <= 25) |>
3   ggplot(aes(x=x,y=y)) +
4     geom_point() +
5     geom_line(
6       aes(y=.epred, group=.draw),
7       alpha=0.33
8     )
```



Credible intervals


```
1 b_post_pred |>  
2   ggplot(aes(x=x, y=y)) +  
3     geom_point() +  
4     tidybayes::stat_lineribbon(  
5       aes(y=.prediction), alpha=0.25  
6     )
```

```
1 b_post_epred |>  
2   ggplot(aes(x=x, y=y)) +  
3     geom_point() +  
4     tidybayes::stat_lineribbon(  
5       aes(y=.epred), alpha=0.25  
6     )
```




Posterior predictive checks

```
1 brms::pp_check(b, ndraws = 25)
```




Residuals - lm

```
1 1 |>  
2 broom::augment() |>  
3 ggplot(aes(x=x, y=.resid)) +  
4   geom_point() +  
5   geom_hline(yintercept=0, color='grey', linetype=2)
```



Residual posteriors - brms

```
1 b %>%  
2 tidybayes::residual_draws(newdata=d) |>  
3 ggplot(aes(x=x, y=.residual, group=x)) +  
4   tidybayes::stat_pointinterval() +  
5   geom_hline(yintercept = 0, color='grey', linetype=2)
```



Model Evaluation

Model assessment

If we remember back to our first regression class, one common option is R^2 which gives us the variability in y explained by our model.

Quick review:

$$\sum_{i=1}^n (y_i - \bar{y})^2 = \sum_{i=1}^n (\hat{y}_i - \bar{y})^2 + \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

Total Model Error

$$R^2 = \frac{SS_{\text{model}}}{SS_{\text{total}}} = \frac{\sum_{i=1}^n (\hat{Y}_i - \bar{Y})^2}{\sum_{i=1}^n (Y_i - \bar{Y})^2} = \frac{\text{Var}(\hat{Y})}{\text{Var}(Y)} = \text{Cor}(Y, \hat{Y})^2$$

Some data prep


```
1 ( b_post_full = b |>
2   tidybayes::spread_draws(b_Intercept, b_x, sigma) |>
3   tidyr::expand_grid(d) |>
4   mutate(
5     y_hat = b_Intercept + b_x * x,
6     resid = y - y_hat
7   )
8 )
9 ## # A tibble: 400,000 × 10
10 ##   .chain .iteration .draw b_Intercept    b_x sigma      x      y  y_hat  resid
11 ##   <int>     <int> <int>     <dbl>  <dbl> <dbl> <int>  <dbl> <dbl> <dbl>
12 ## 1         1         1         1     -0.961 0.0671 1.33      1 -3.24 -0.894 -2.35
13 ## 2         1         1         1     -0.961 0.0671 1.33      2 -2.00 -0.826 -1.18
14 ## 3         1         1         1     -0.961 0.0671 1.33      3 -2.59 -0.759 -1.84
15 ## 4         1         1         1     -0.961 0.0671 1.33      4 -3.07 -0.692 -2.38
16 ## 5         1         1         1     -0.961 0.0671 1.33      5 -1.88 -0.625 -1.26
17 ## 6         1         1         1     -0.961 0.0671 1.33      6 -0.807 -0.558 -0.24
18 ## 7         1         1         1     -0.961 0.0671 1.33      7 -2.09 -0.491 -1.60
```

Bayesian R^2

When we compute any statistic for our model we want to do so at each iteration so that we can obtain the posterior distribution of that particular statistic (e.g. the posterior distribution of R^2 in this case).

```
1 ( b_R2 = b_post_full %>%
2   group_by(.iteration) %>%
3   summarize(
4     R2_classic = var(y_hat) / var(y),
5     R2_bayes   = var(y_hat) / (var(y_hat) + var(resid))
6   )
7 )
8 ## # A tibble: 1,000 × 3
9 ##   .iteration R2_classic R2_bayes
10 ##      <int>      <dbl>   <dbl>
11 ## 1         1      0.676     0.642
12 ## 2         2      0.757     0.663
13 ## 3         3      0.593     0.610
14 ## 4         4      0.667     0.637
15 ## 5         5      0.630     0.625
16 ## 6         6      0.627     0.624
17 ## 7         7      0.557     0.593
18 ## 8         8      0.640     0.625
19 ## 9         9      0.730     0.657
```

Uh oh ...



Sanity check

```
1 ( l_pred = broom::augment(l) )
2 ## # A tibble: 100 × 8
3 ##       y       x .fitted .resid  .hat .sig
4 ##   <dbl> <int>   <dbl> <dbl>  <dbl> <dbl>
5 ## 1 -3.24     1 -0.962 -2.28  0.0394 1.0
6 ## 2 -2.00     2 -0.893 -1.11  0.0382 1.0
7 ## 3 -2.59     3 -0.825 -1.77  0.0371 1.0
8 ## 4 -3.07     4 -0.757 -2.31  0.0359 1.0
9 ## 5 -1.88     5 -0.688 -1.20  0.0348 1.0
10 ## 6 -0.807    6 -0.620 -0.187 0.0338 1.0
11 ## 7 -2.09     7 -0.551 -1.54  0.0327 1.0
12 ## 8 -0.227    8 -0.483  0.256 0.0317 1.0
13 ## 9  0.333    9 -0.415  0.747 0.0307 1.0
14 ## 10  1.98    10 -0.346  2.32  0.0297 1.0
15 ## # ... with 90 more rows
```

```
1 broom::glance(l)$r.squared
2 ## [1] 0.6265565
3
4 var(l_pred$.fitted) / var(l_pred$y)
5 ## [1] 0.6265565
6
7 var(l_pred$.fitted) / (var(l_pred$.fitted) +
8 ## [1] 0.6265565
```


What if we collapsed first?

Here we calculate the posterior mean of \hat{y} and use that to estimate R^2 ,

```
1  b_post_full %>%
2    group_by(x) %>%
3    summarize(
4      y_hat = mean(y_hat),
5      y = mean(y),
6      resid = mean(y - y_hat),
7      .groups = "drop"
8    ) %>%
9    summarize(
10     R2_classic = var(y_hat) / var(y),
11     R2_bayes   = var(y_hat) / (var(y_hat) + var(resid))
12   )
13 ## # A tibble: 1 × 2
14 ##   R2_classic R2_bayes
15 ##       <dbl>   <dbl>
16 ## 1       0.627   0.627
```

Some problems with R^2

Some new issues,

- R^2 doesn't really make sense in the Bayesian context
 - multiple possible definitions with different properties
 - fundamental equality doesn't hold anymore
 - Possible to have $R^2 > 1$

Some old issues,

- R^2 always increases (or stays the same) when adding a predictor
- R^2 is highly susceptible to over fitting
- R^2 is sensitive to outliers
- R^2 depends heavily on values of y (can differ for two equivalent models)

Some Other Metrics

Root Mean Square Error

The traditional definition of rmse is as follows

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$$

In the bayesian context, we have posterior samples from each parameter / prediction of interest so we can express this as

$$\frac{1}{m} \sum_{s=1}^m \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_{i,s})^2}$$

where m is the number of iterations and \hat{Y}_i^s is the prediction for Y_i at iteration s .

Continuous Rank Probability Score

Another approach is the continuous rank probability score which comes from the probabilistic forecasting literature, it compares the full posterior predictive distribution to the observation / truth.


$$\text{CRPS} = \int_{-\infty}^{\infty} (F_{\hat{y}}(z) - 1_{z \geq y})^2 dz$$

where $F_{\hat{y}}$ is the CDF of \hat{y} (the posterior predictive distribution for y) and $1_{z \geq Y}$ is an indicator function which equals 1 when $z \geq y$, the true/observed value of y .


Since this calculates a score for a single probabilistic prediction we can naturally extend it to multiple predictions by calculating an average CRPS

$$\frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} (F_{\hat{y}_i}(z) - 1_{z \geq y_i})^2 dz$$

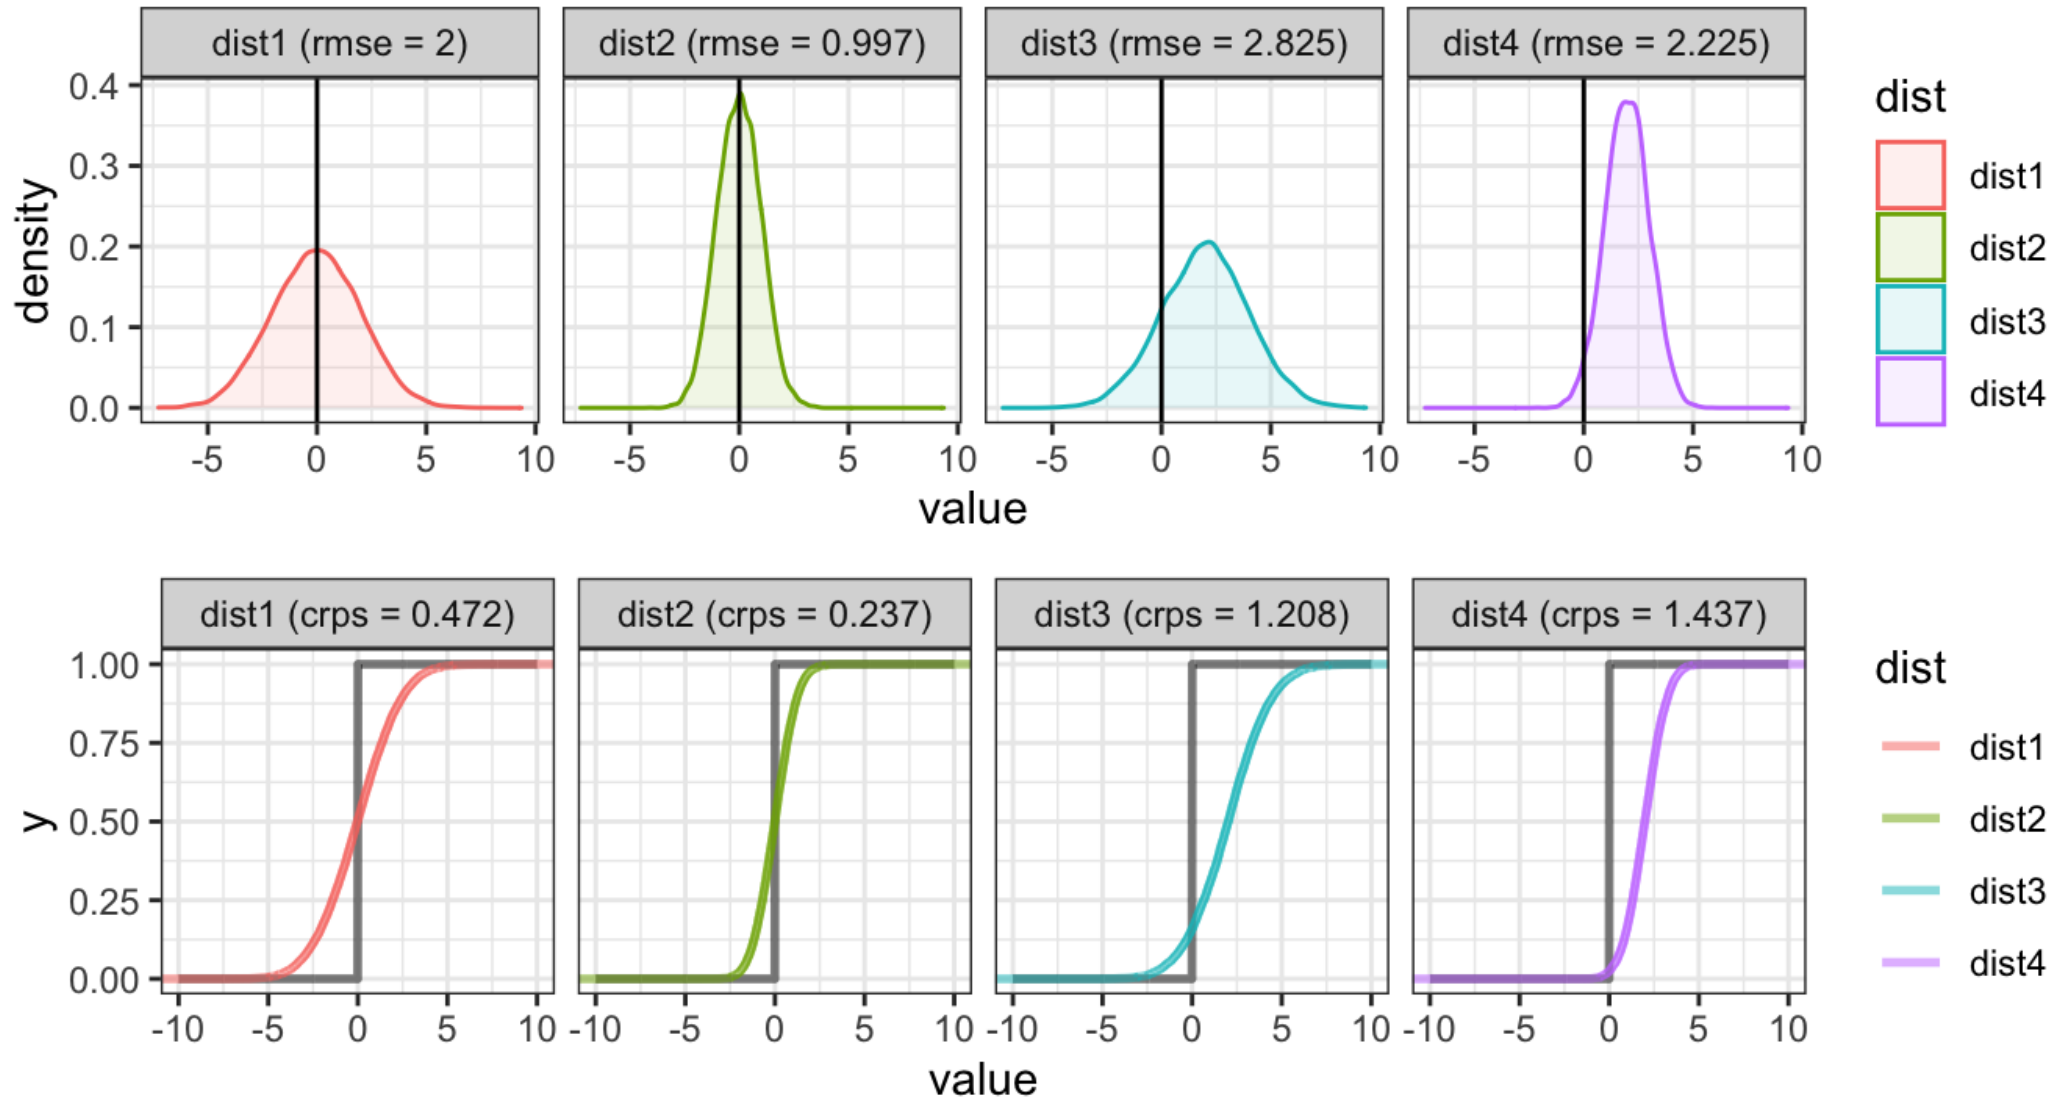
CDF vs Indicator



Empirical CDF vs Indicator




Accuracy vs. Precision



Empirical Coverage

One final method, which assesses model calibration is to examine how well credible intervals, derived from the posterior predictive distributions of the y s, capture the true/observed values.



Back to our example


RMSE - \hat{y}

```
1 broom::augment(l) |>
2   yardstick::rmse(y, .fitted)
3 ## # A tibble: 1 × 3
4 ##   .metric .estimator .estimate
5 ##   <chr>   <chr>         <dbl>
6 ## 1 rmse     standard         1.52
7
8 ( b_rmse = b_post_full %>%
9   group_by(.chain, .iteration) %>%
10  summarize(
11    rmse = sqrt( sum( (y - y_hat)^2 ) / n() )
12  )
13 )
14 ## # A tibble: 4,000 × 3
15 ## # Groups:   .chain [4]
16 ##   .chain .iteration rmse
```

RMSE - \hat{y} - Results

```
1 b_rmse |>
2   group_by(.chain) |>
3   summarize(post_mean = mean(rmse))
4 ## # A tibble: 4 × 2
5 ##   .chain post_mean
6 ##   <int>     <dbl>
7 ## 1         1     1.54
8 ## 2         2     1.54
9 ## 3         3     1.54
10 ## 4         4     1.54
```

```
1 ggplot(b_rmse) +
2   geom_density(
3     aes(x=rmse, fill=as.factor(.chain))
4     alpha=0.33
5   ) +
6   guides(fill="none")
```



CRPS - \hat{y}

```
1 ( b_crps = b_post_full |>
2   group_by(.chain, x) |>
3   summarise(
4     crps = calc_crps(y_hat, y)
5   )
6 )
7 ## # A tibble: 400 × 3
8 ## # Groups:   .chain [4]
9 ##   .chain     x crps
10 ##   <int> <int> <dbl>
11 ## 1       1     1  2.10
12 ## 2       1     2  0.934
13 ## 3       1     3  1.60
14 ## 4       1     4  2.14
15 ## 5       1     5  1.03
16 ## 6       1     6  0.113
17 ## 7       1     7  1.38
18 ## 8       1     8  0.156
```

```
1 b_crps |>
2   group_by(.chain) |>
3   summarize(
4     avg_crps = mean(crps)
5   )
6 ## # A tibble: 4 × 2
7 ##   .chain avg_crps
8 ##   <int>   <dbl>
9 ## 1       1     1.19
10 ## 2       2     1.19
11 ## 3       3     1.19
12 ## 4       4     1.19
```

Empirical Coverage - \hat{y}

```
1 ( b_cover = b_post_full %>%
2   group_by(x, y) %>%
3   tidybayes::mean_hdi(
4     y_hat, .prob = c(0.5, 0.9, 0.95)
5   )
6 )
7 ## # A tibble: 300 × 8
8 ##       x     y y_hat .lower .upper .width .point .interval
9 ##   <int> <dbl> <dbl> <dbl> <dbl> <dbl> <chr> <chr>
10 ## 1     1 -3.24 -0.963 -1.14  -0.735  0.5 mean  hdi
11 ## 2     2 -2.00 -0.894 -1.07  -0.669  0.5 mean  hdi
12 ## 3     3 -2.59 -0.826 -1.02  -0.630  0.5 mean  hdi
13 ## 4     4 -3.07 -0.758 -0.947 -0.562  0.5 mean  hdi
14 ## 5     5 -1.88 -0.689 -0.847 -0.467  0.5 mean  hdi
15 ## 6     6 -0.807 -0.621 -0.771 -0.398  0.5 mean  hdi
16 ## 7     7 -2.09 -0.552 -0.745 -0.378  0.5 mean  hdi
17 ## 8     8 -0.227 -0.484 -0.667 -0.306  0.5 mean  hdi
18 ## 9     9  0.333 -0.415 -0.594 -0.240  0.5 mean  hdi
```

Empirical Coverage - \hat{y} - Results

```
1 b_cover %>%
2   mutate(contains = y >= .lower & y <= .upper) %>%
3   group_by(prob = .width) %>%
4   summarize(
5     emp_cov = sum(contains)/n()
6   )
7 ## # A tibble: 3 × 2
8 ##   prob emp_cov
9 ##   <dbl> <dbl>
10 ## 1  0.5    0.02
11 ## 2  0.9    0.11
12 ## 3  0.95   0.14
```

What went wrong now?