Agenda

Transformations & assessing fit with predictive plots

1. Administrative
2. Interpretation with transformed variables
3. Prior predictive plots
4. Visualizing model predictions
5. ***Hands on:*** *Visualizing model predictions in R*

Transfor-
mations

MONTREAL BIKE TRAFFIC

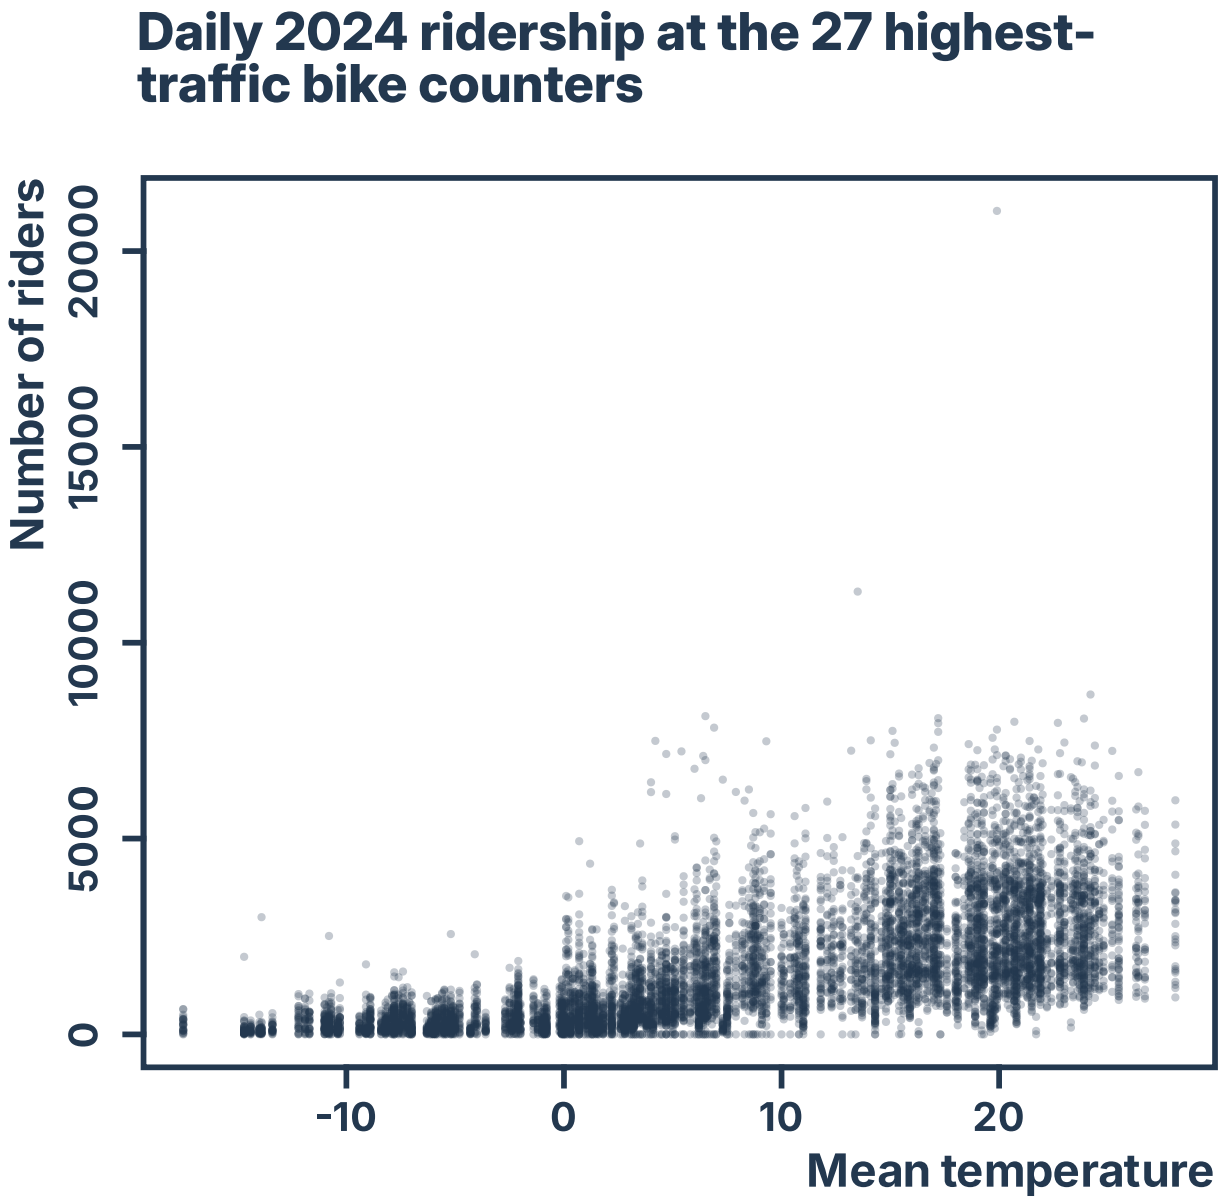

INTERPRETING COEFFICIENTS

Predicting ridership using temperature

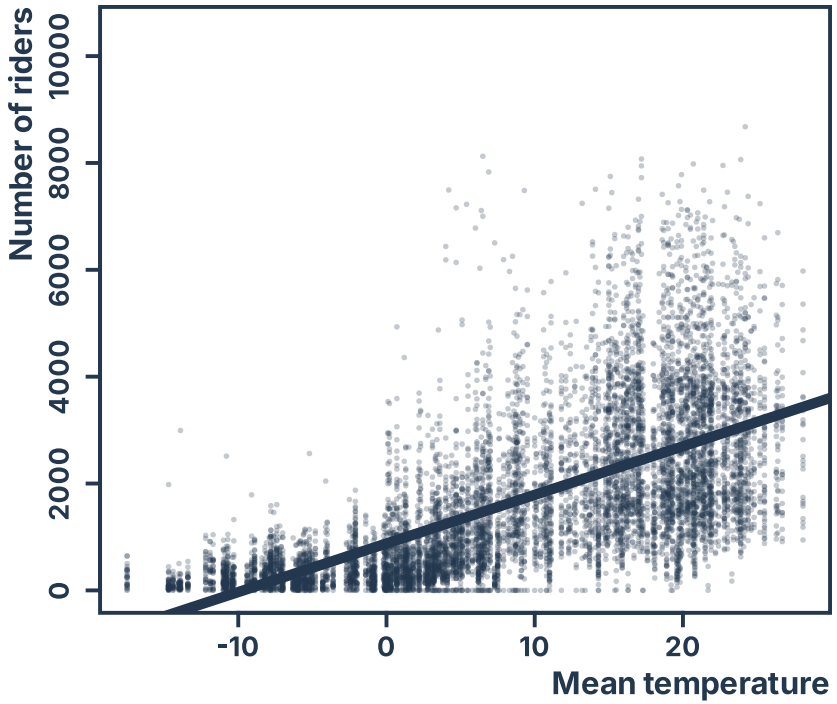

$$R_i \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta T_i$$

	Post. mean
α	878.28
β	90.96

If the temperature changes from t_1 to t_2, how much do we expect ridership to change?

Note: $\hat{\alpha} = \text{E}(\alpha)$ *and* $\hat{\beta} = \text{E}(\beta)$

$$\begin{aligned}\text{E}(R_i|T_i = t_2, \mathbf{M}) - \text{E}(R_i|T_i = t_1, \mathbf{M}) &= (\hat{\alpha} + \hat{\beta}t_2) - (\hat{\alpha} + \hat{\beta}t_1) \\ &= \hat{\beta}(t_2 - t_1) \\ &= 90.96(t_2 - t_1)\end{aligned}$$

INTERPRETING COEFFICIENTS

Predicting ridership using temperature

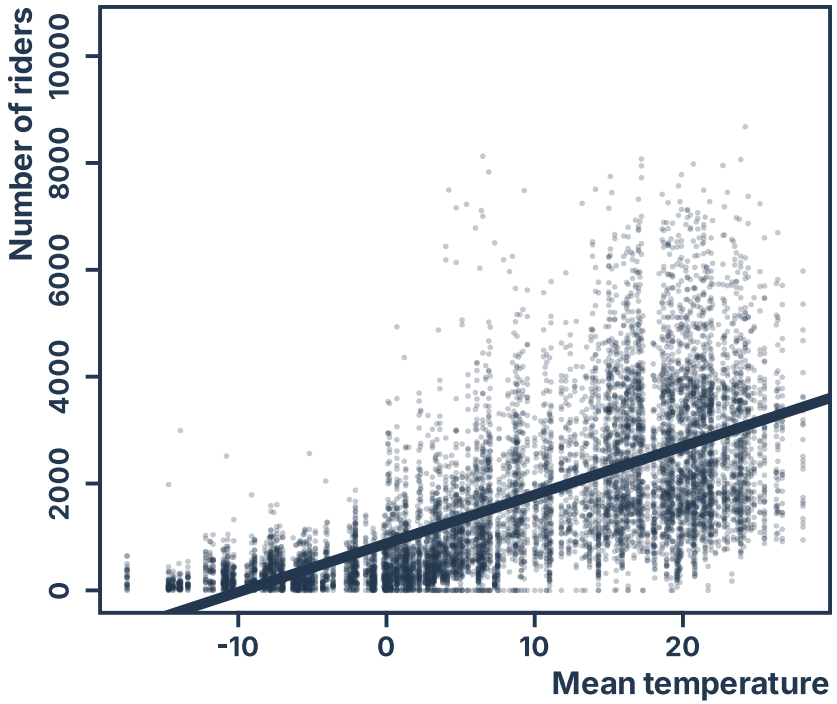

$$R_i \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta T_i$$

	Post. mean
α	878.28
β	90.96

Units of ridership: Degrees celsius

Units of ridership: Number of riders

Interpretation of β: For every increase of ***one degree celsius***, the model predicts an average of ***90.96 more riders*** per day at each bike counter

STANDARDIZED VARIABLES

Predicting St(ridership) using St(temperature)

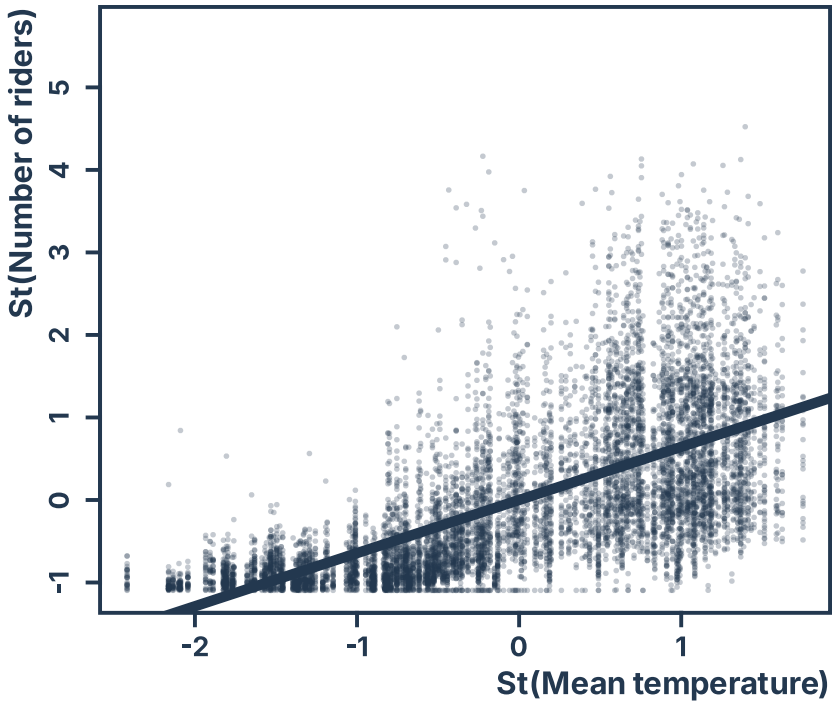

$$\text{St}(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

	Post. mean
α	0.0
β	0.644

Units of temperature: Standard deviations of temperature

Units of ridership: Standard deviations of ridership

Interpretation of β: For every increase of ***one standard deviation of temperature***, the model predicts an average of ***0.644* more standard deviations of ridership** per day at each bike counter

LOG OF OUTCOME VARIABLE

Predicting Log(ridership) using temperature

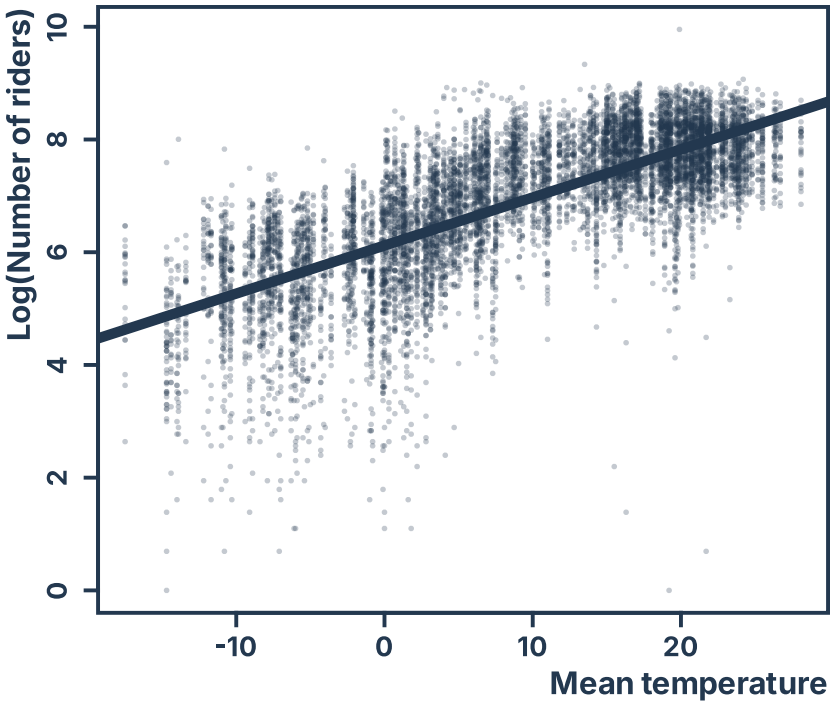

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta T_i$$

	Post. mean	Exp(mean)
α	6.12	454.47
β	0.085	1.089

If the temperature changes from t_1 to t_2, how much do we expect ridership to change?

$$\begin{aligned}\frac{\text{E}(R_i|T_i = t_2, \mathbf{M})}{\text{E}(R_i|T_i = t_1, \mathbf{M})} &= \exp\left(\log\left(\frac{\text{E}(R_i|T_i = t_2, \mathbf{M})}{\text{E}(R_i|T_i = t_1, \mathbf{M})}\right)\right) \\ &= \exp\Big((\hat{\alpha} + \hat{\beta}t_2) - (\hat{\alpha} + \hat{\beta}t_1)\Big) \\ &= \exp(\hat{\beta}(t_2 - t_1)) \\ &= \exp(0.085(t_2 - t_1))\end{aligned}$$

LOG OF OUTCOME VARIABLE

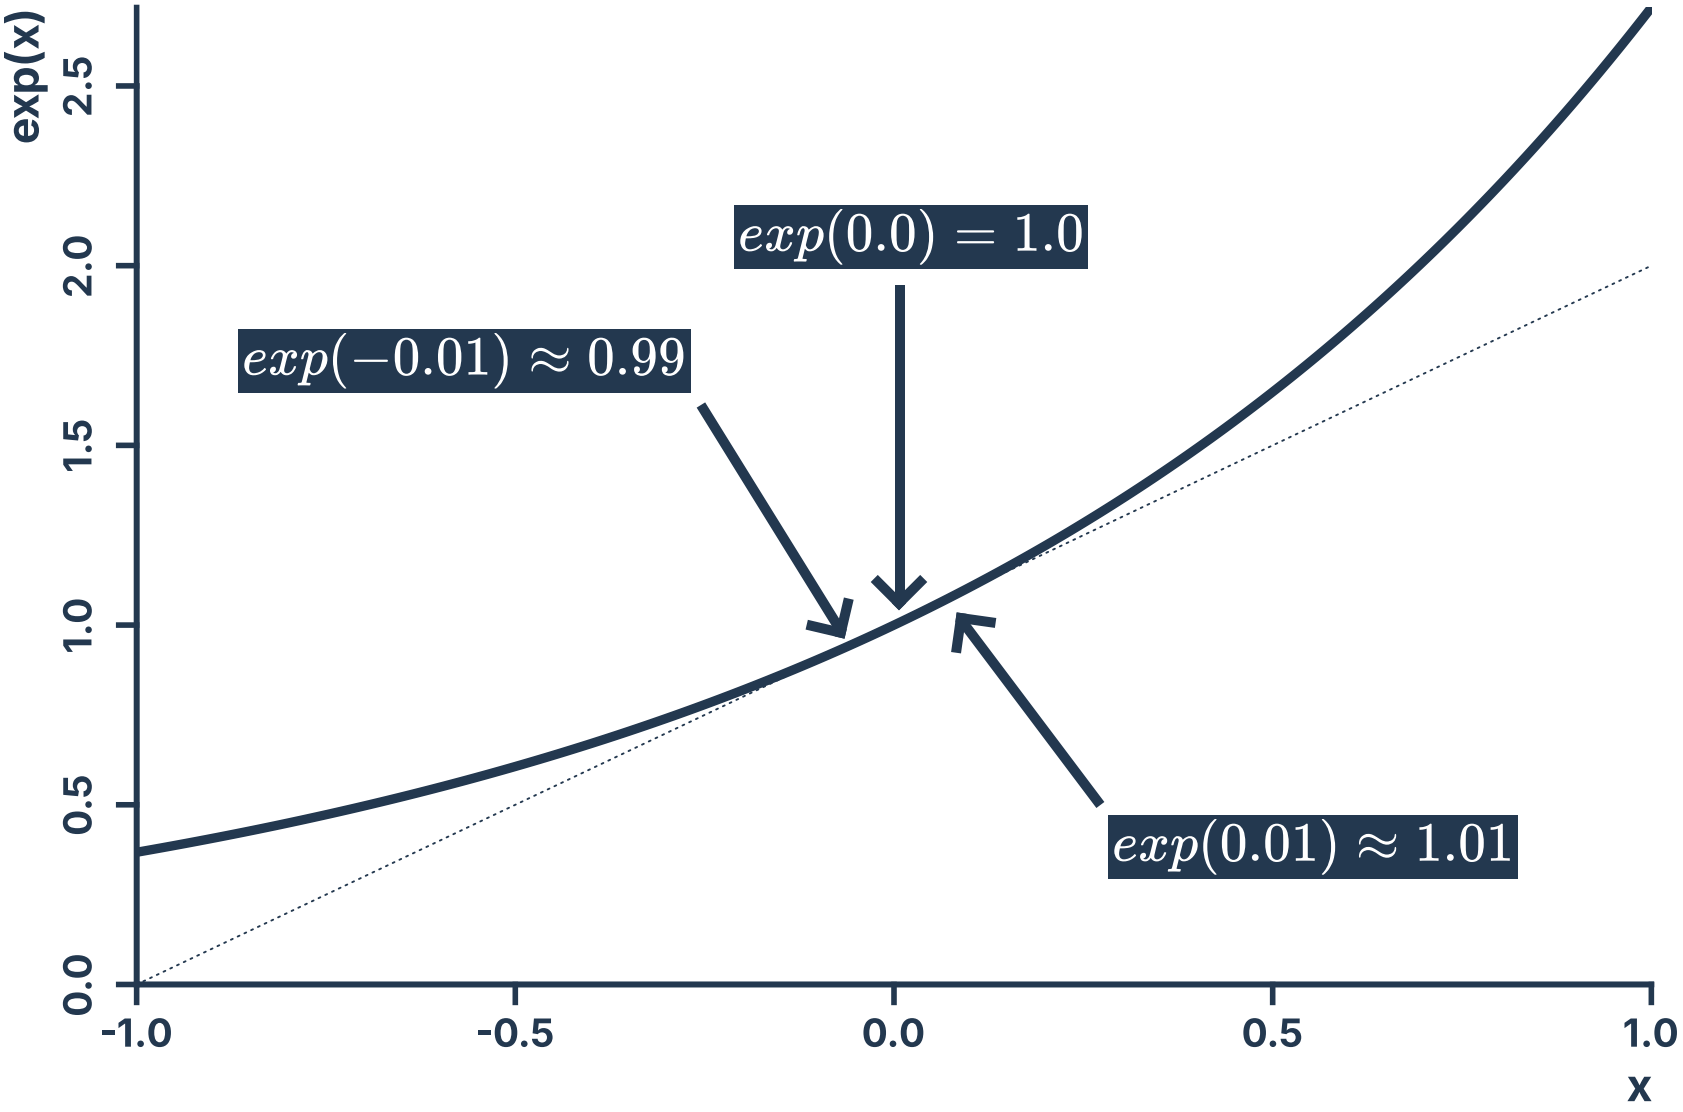

Predicting Log(ridership) using temperature

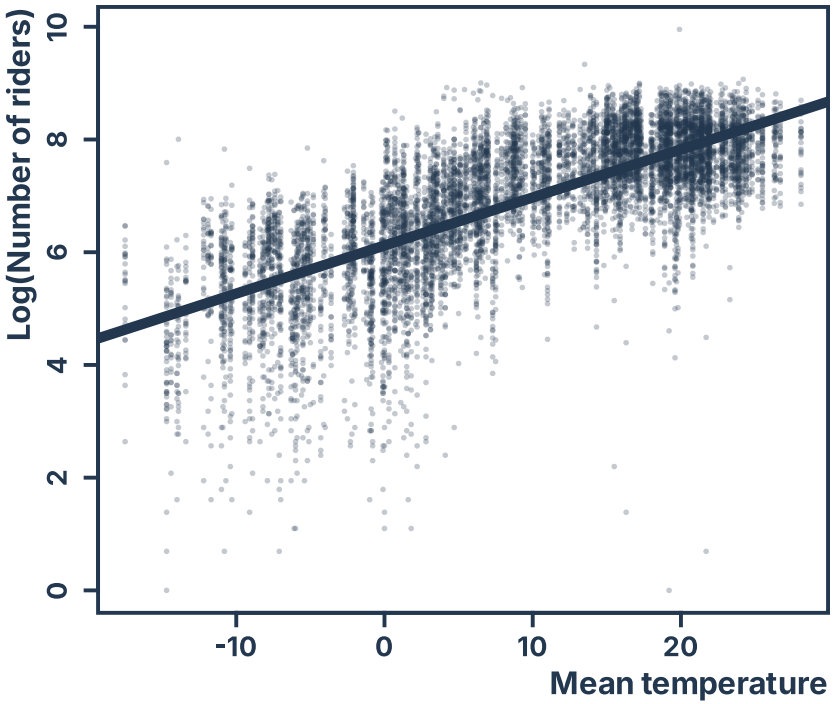

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta T_i$$

	Post. mean	Exp(mean)
α	6.12	454.47
β	0.085	1.089

Units of temperature: Degrees celsius

Units of ridership: Log ridership

Interpretation of β: For every increase of ***one degree celsius***, the model predicts an average increase of **8.9% in ridership** per day at each bike counter

Prior and posterior predictive plots

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

Prior predictive plots allow you to visualize the implications of a set of priors

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

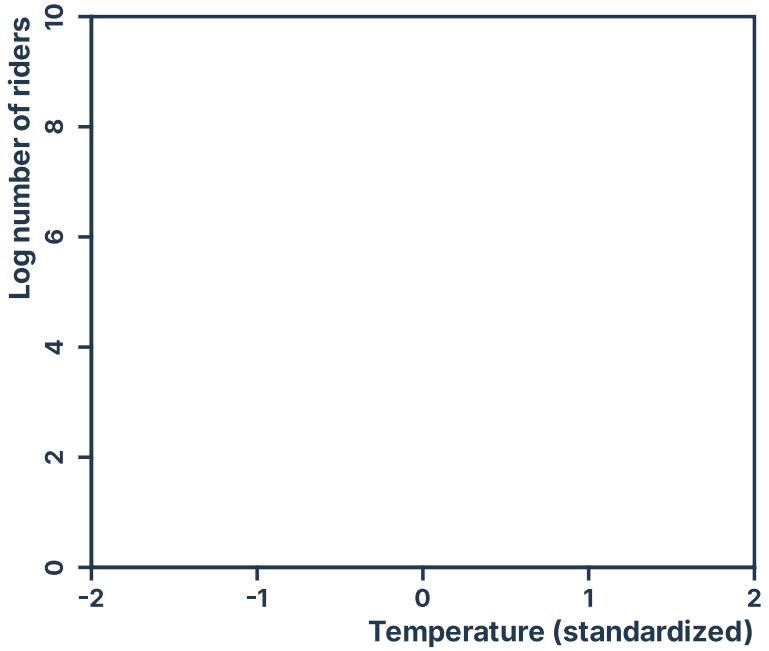

α	β
12.45	2.58

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

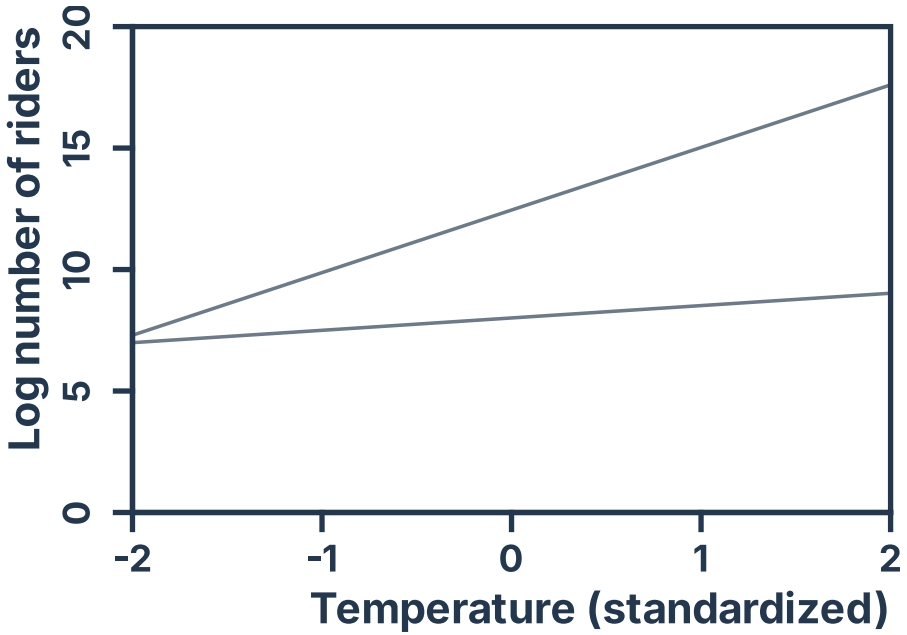

α	β
12.45	2.58
8.01	0.51

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

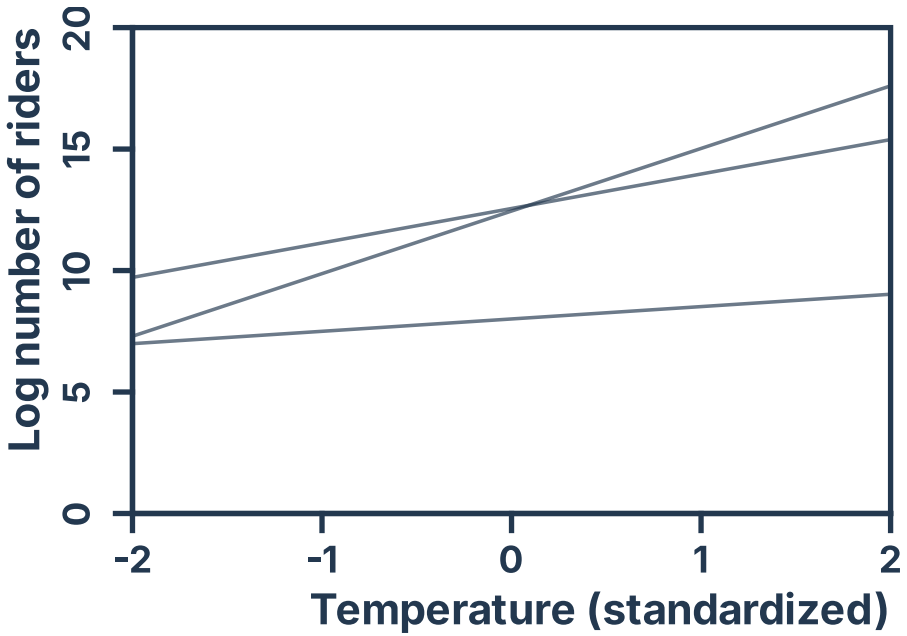

α	β
12.45	2.58
8.01	0.51
12.55	1.42

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

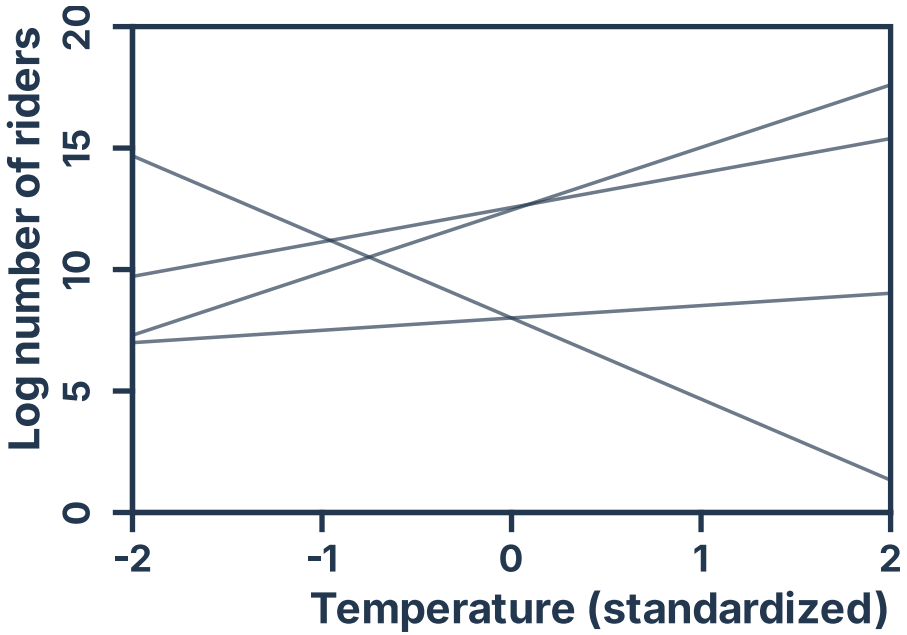

α	β
12.45	2.58
8.01	0.51
12.55	1.42
8.01	-3.34

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

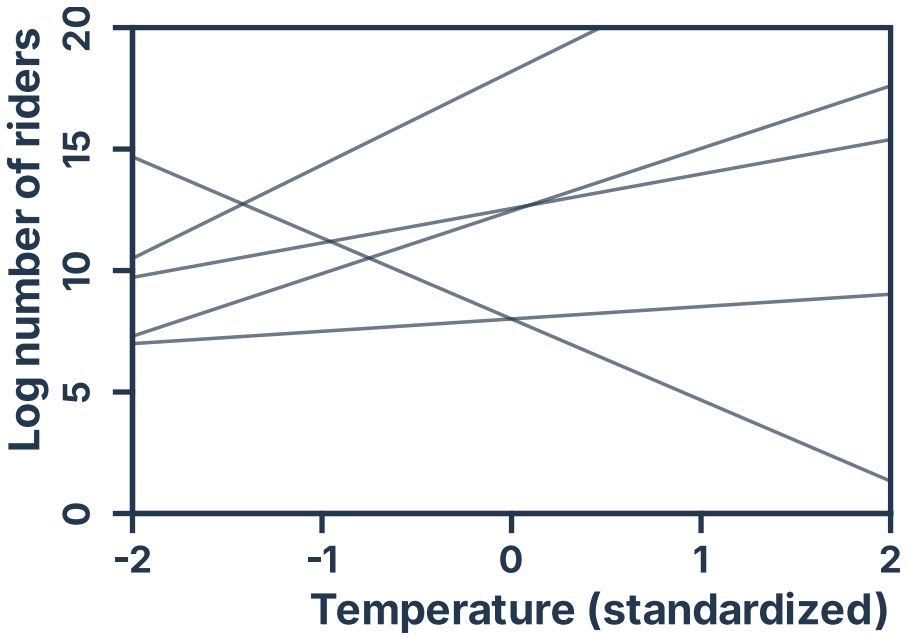

α	β
12.45	2.58
8.01	0.51
12.55	1.42
8.01	-3.34
18.19	3.85

PRIOR PREDICTIVE PLOT

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

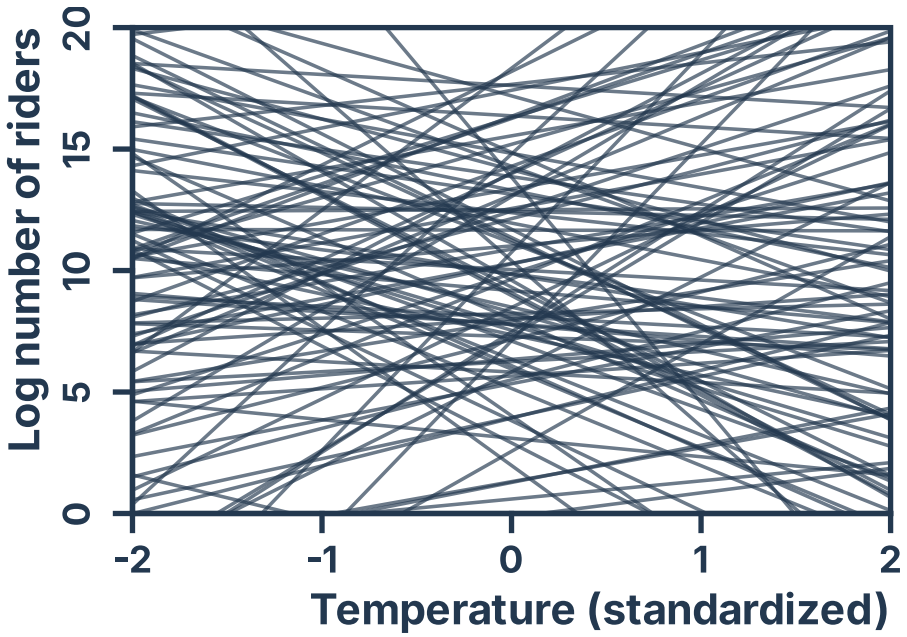

α	β
12.45	2.58
8.01	0.51
12.55	1.42
8.01	-3.34
18.19	3.85
13.11	3.24
11.01	0.66
15.54	1.36
8.97	-4.10
8.11	-0.42
8.48	3.76
⋮	⋮

VISUALIZING PREDICTIONS

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

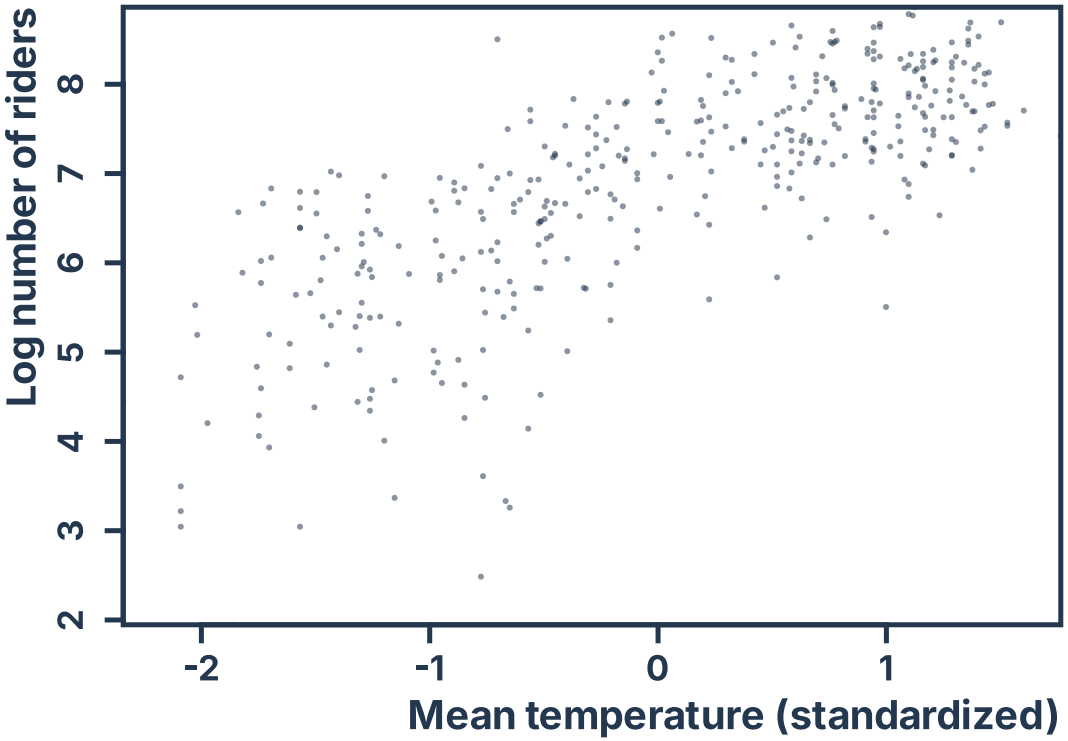

To better illustrate uncertainty in posterior estimates, we will use a random subsample of 400 counter–days.

VISUALIZING PREDICTIONS

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

	Post. median	80% post. interval
α	6.877	(6.823, 6.931)
β	0.940	(0.887, 0.996)
σ	0.841	(0.804, 0.981)

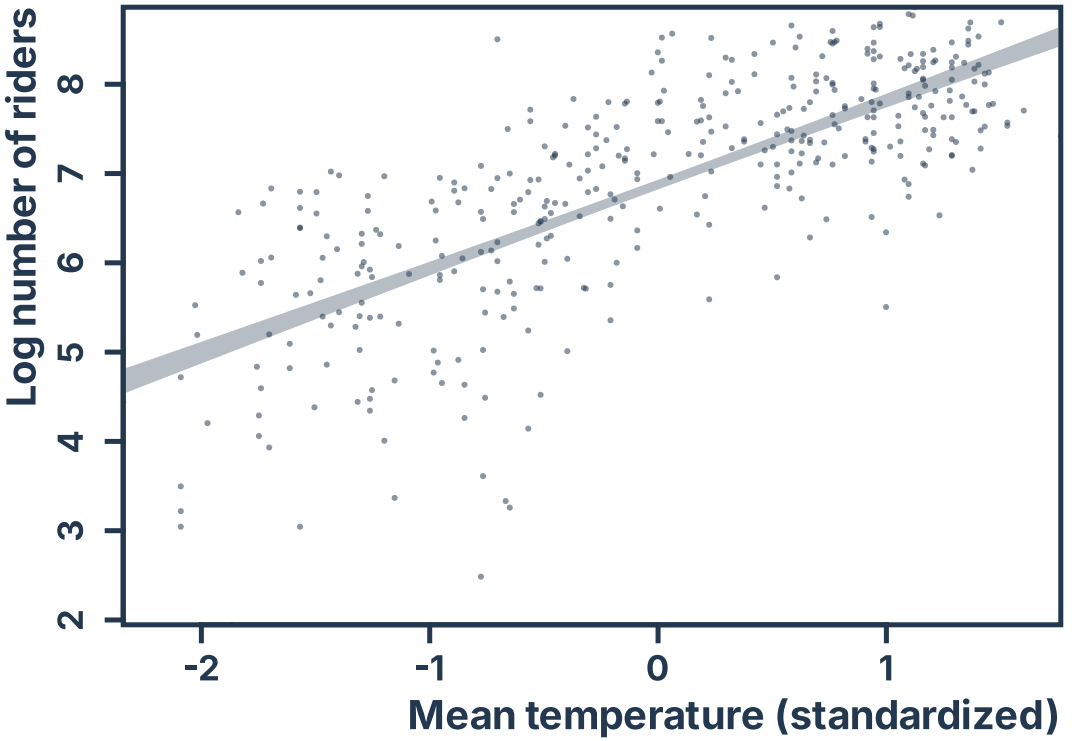

Posterior distribution of mean:

$$\Pr(\mu | T = t)$$

1. Take a sample of size N from posterior $\Pr(\alpha, \beta, \sigma | D)$.
2. For each value of standardized temperature t, calculate N values $\mu_n = \alpha_n + \beta_n t$.
3. Calculate quantiles (say, 10% and 90%) of the posterior draws $\{\mu_n\}$ at each value of t.

VISUALIZING PREDICTIONS

$$\log(R_i) \sim \text{Norm}(\mu_i, \sigma)$$
$$\mu_i = \alpha + \beta \text{St}(T_i)$$

$$\alpha \sim \text{Norm}(10, 5)$$
$$\beta \sim \text{Norm}(0, 3)$$
$$\sigma \sim \text{Unif}(0, 5)$$

	Post. median	80% post. interval
α	6.877	(6.823, 6.931)
β	0.940	(0.887, 0.996)
σ	0.841	(0.804, 0.981)

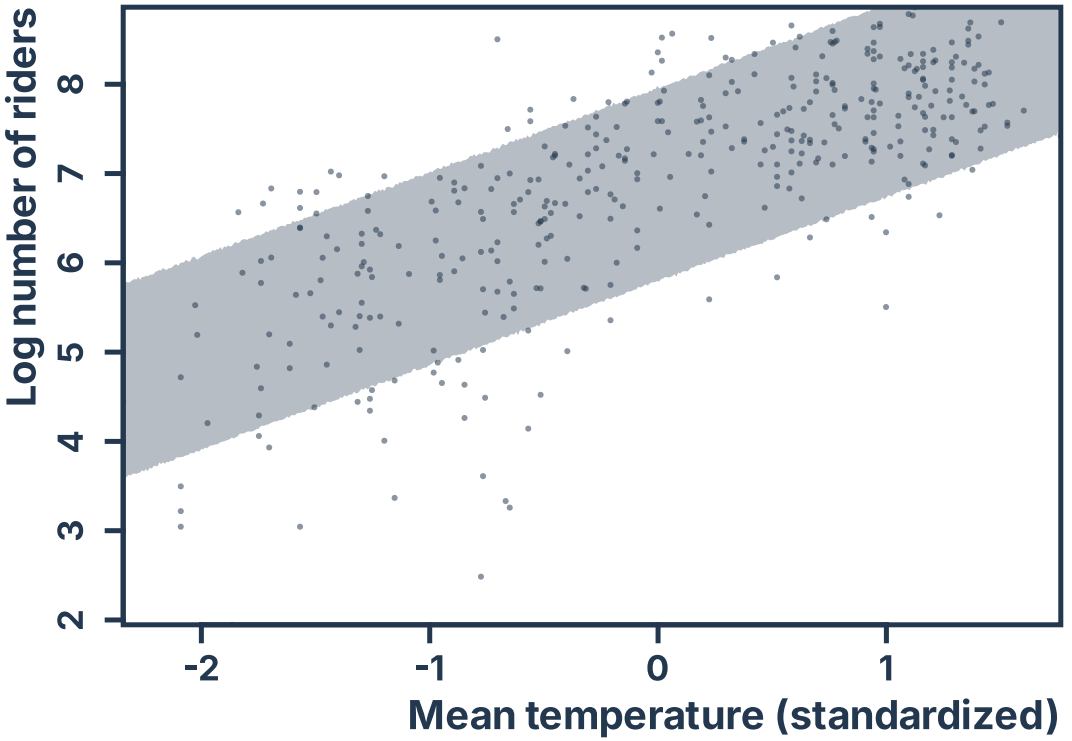

Posterior predictive distribution:

$$\Pr(\log(R_i)|T = t)$$

1. Take a sample of size N from posterior $\Pr(\alpha, \beta, \sigma|D)$.
2. For each value of standardized temperature t, calculate N values $\mu_n = \alpha_n + \beta_n t$.
3. For each of the N posterior draws for temperature t, draw $p_n \sim \text{Norm}(\mu_n, \sigma_n)$.
4. Calculate quantiles (say, 10% and 90%) of the posterior draws $\{p_n\}$ at each value of t.

POSTERIOR MEAN VS. PREDICTED

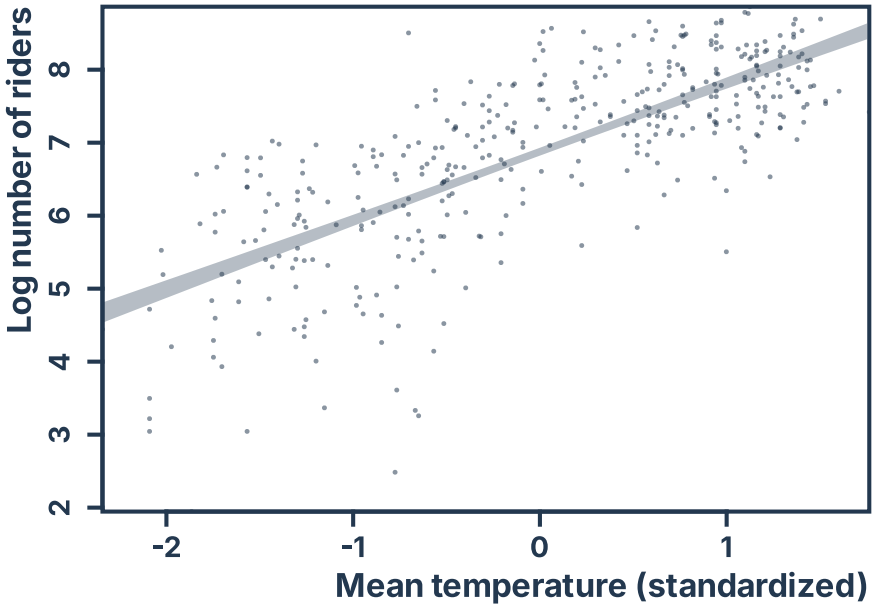

For any given mean temperature, μ is the value of log ridership that is "expected" by the model on a day with that temperature.

The *posterior distribution* of μ describes our uncertainty about the value of μ after seeing the data.

This distribution takes into account the model coefficients α and β, but *not* the model standard deviation σ.

The 80% posterior interval of μ should get narrower as more data is added.

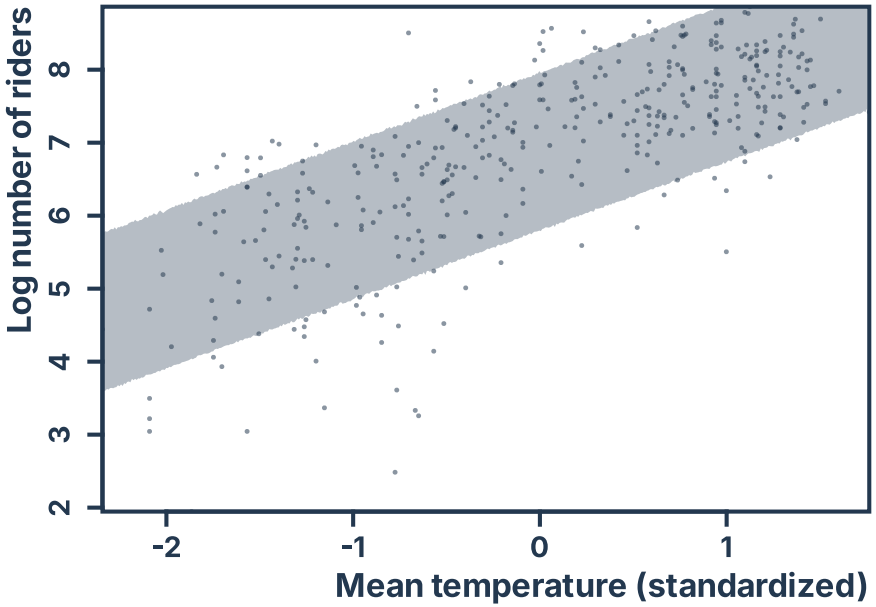

For any given mean temperature, the *posterior predictive distribution* describes the range of riderships we would expect for any day with that temperature.

The *posterior predictive distribution* describes our uncertainty about the value of $\log(R_i)$ after seeing the data.

This distribution takes into account all the model parameters: α, β, and σ.

The 80% posterior *predictive* interval should contain about 80% of the data, and should not get appreciably narrower as more data is added.

POSTERIOR MEAN VS. PREDICTED

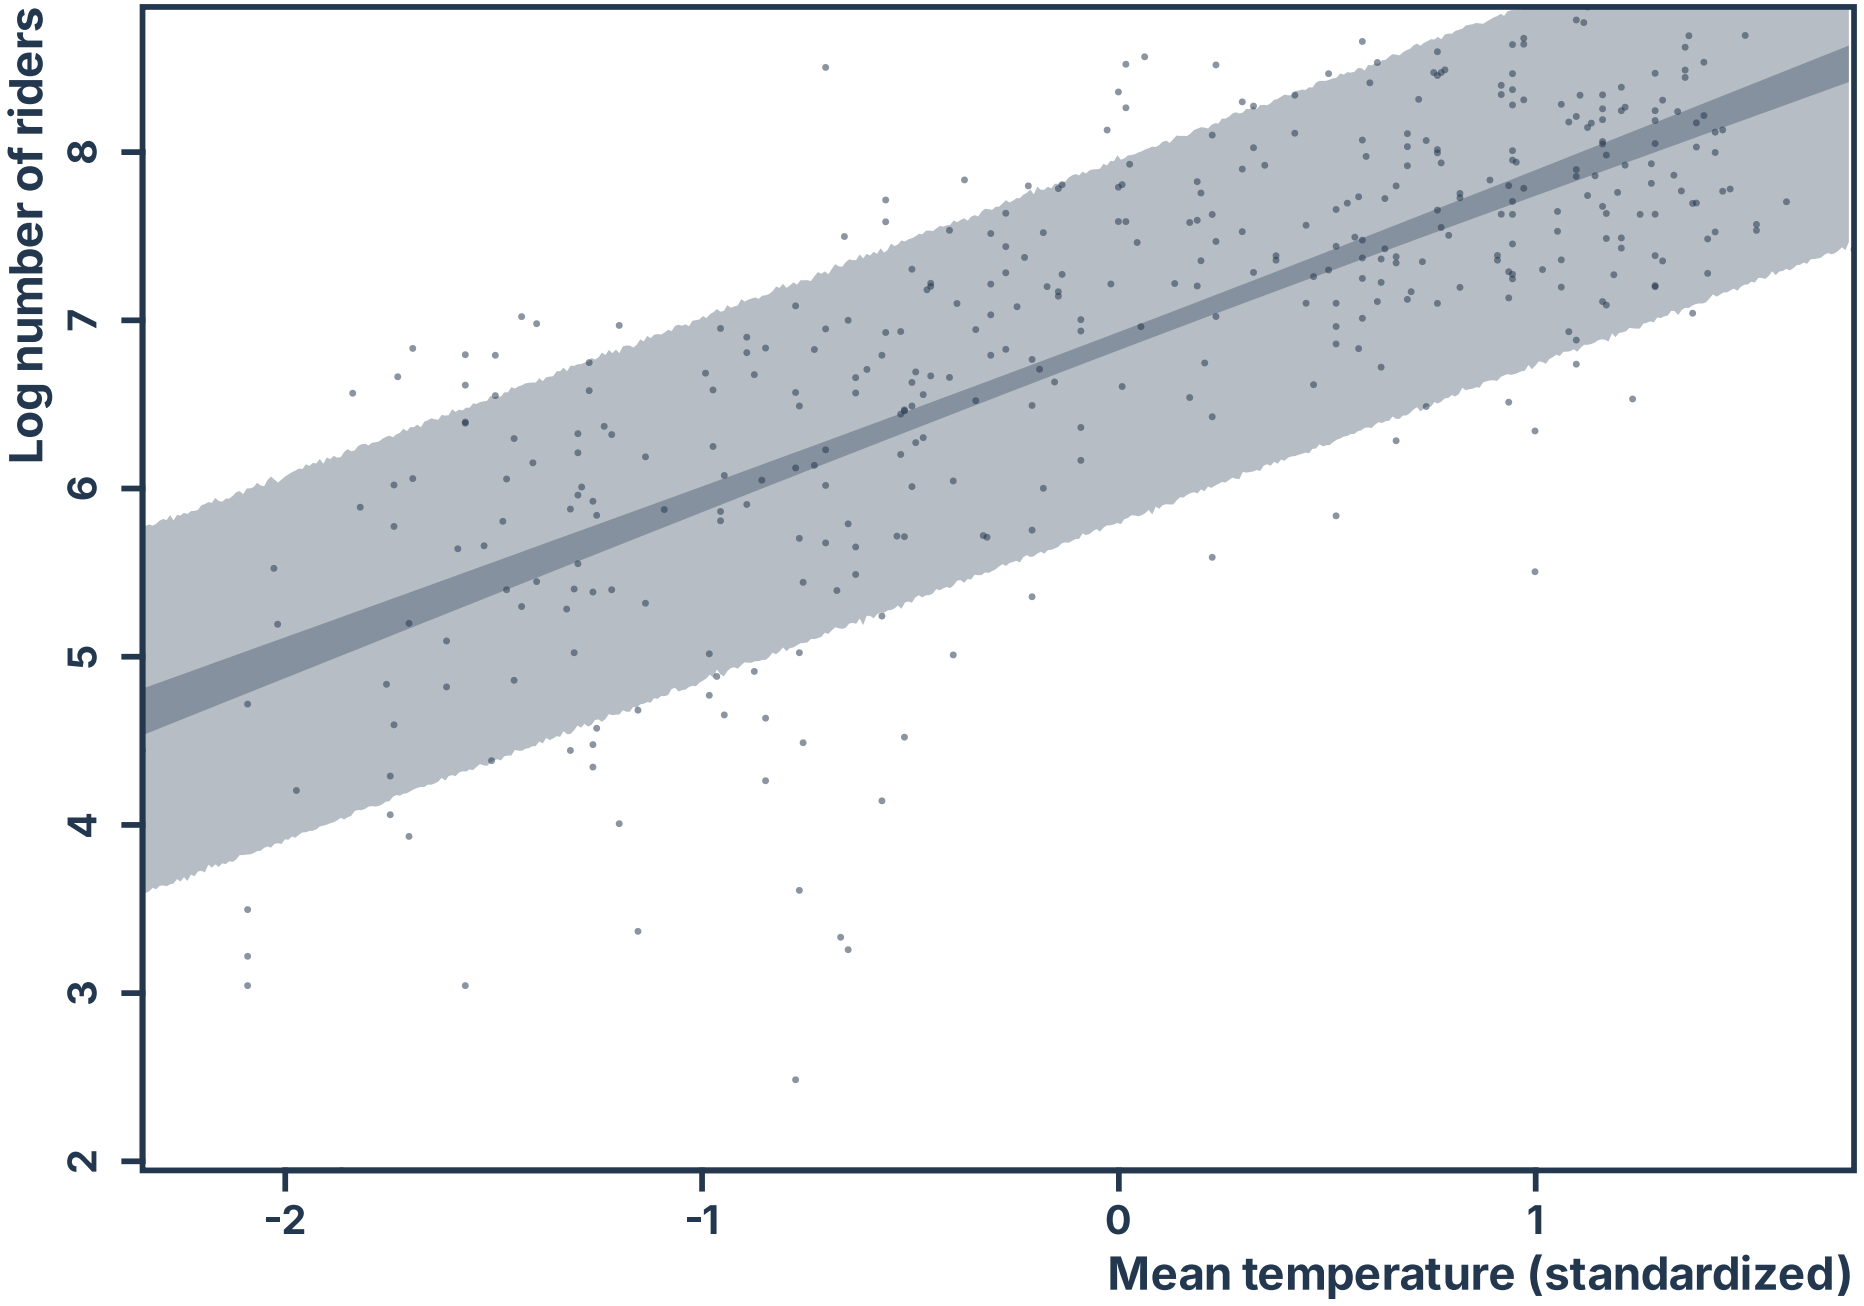

Image credit

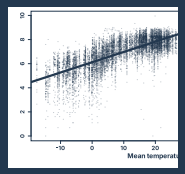

Figures by Peter McMahan (source code)

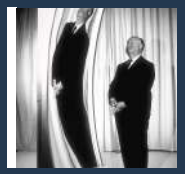

Alfred Hitchcock in Funhouse Mirror by Globe Photos

Photo by Peter McMahan

Photo via Flickr user Midnight Believer