

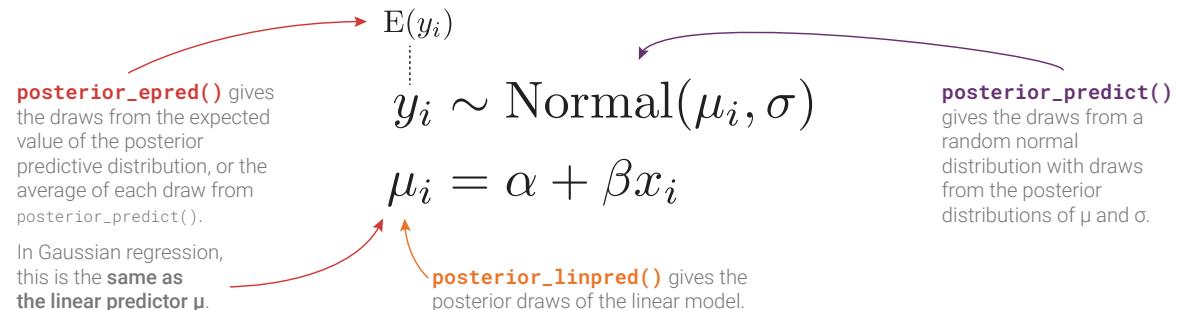
Posterior predictions, linear predictions, and expectation of posterior predictions cheat sheet

v2.0 (2022-10-23) Andrew Heiss (<https://www.andrewheiss.com>; @andrewheiss)

Normal Gaussian models

Normal or Gaussian models are roughly equivalent to standard frequentist ordinary least squares (OLS) models and are generally interpreted the same way.

All the Bayesian prediction functions return values on the original scale of the outcome. In Gaussian models, the results from `posterior_epred()` and `posterior_linpred()` are identical. The draws from `posterior_predict()` will generally have the same mean as the linear predictor, but with more variance since the overall model σ is incorporated into the predictions.

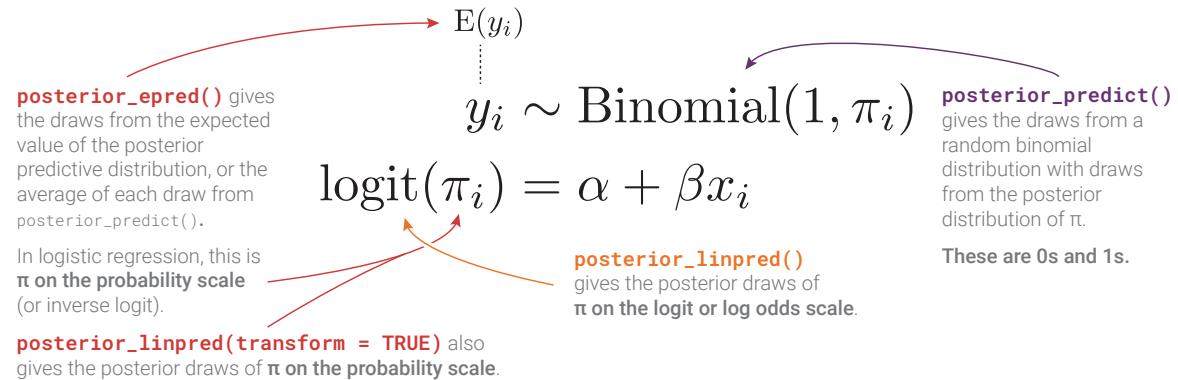


Generalized linear models with link transformations

Generalized linear models (e.g., logistic, probit, ordered logistic, exponential, Poisson, negative binomial, etc.) use special link functions (e.g. logit, log, etc.) to transform the outcome variable into a scale that is more amenable to linear regression.

Estimates from these models can be used in their transformed scales (e.g., log odds in logistic regression) or can be back-transformed into their original scale (e.g., probabilities in logistic regression).

When working with links, the various Bayesian prediction functions return values on different scales, each corresponding to different parts of the model.



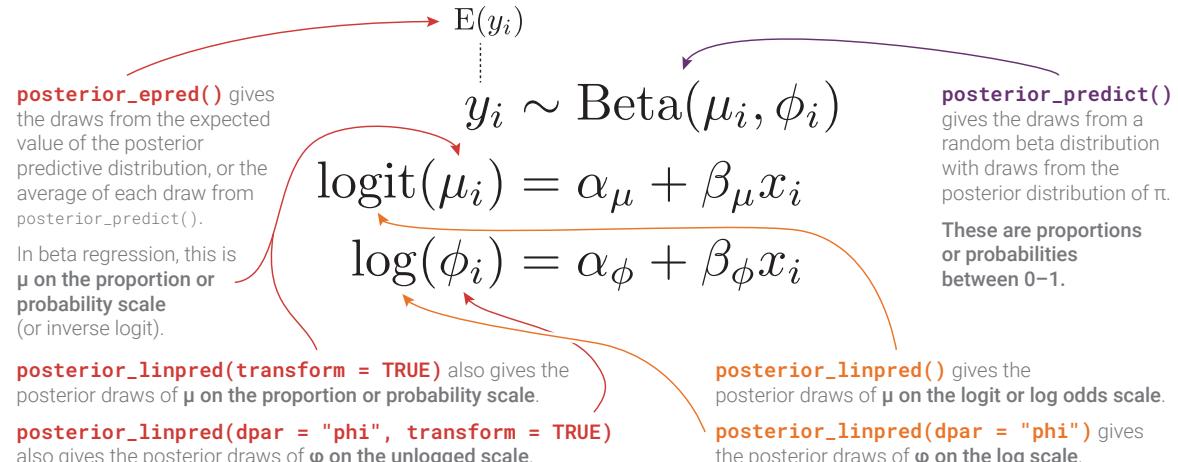
Distributional models with link transformations

Regression models often focus solely on the location parameter of the model (e.g., μ in $\text{Normal}(\mu, \sigma)$; π in $\text{Binomial}(n, \pi)$). However, it is also possible to specify separate predictors for the scale or shape parameters of models (e.g., σ in $\text{Normal}(\mu, \sigma)$, ϕ in $\text{Beta}(\mu, \phi)$).

More complex models use a collection of distributional parameters (e.g., zero-inflated beta models estimate a mean μ , precision ϕ , and a zero-inflated parameter z_i ; hurdle lognormal models estimate a mean μ , scale σ , and a hurdle parameter h_u).

Even plain old Gaussian models become distributional models when a set of predictors is specified for σ (e.g. `brm(y ~ x1 + x2, sigma ~ x2 + x3)`)

When working with extra distributional parameters (dpar), the various Bayesian prediction functions return values on different scales for different components of the model.

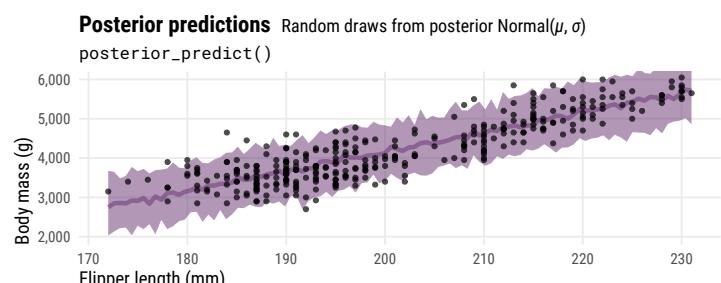
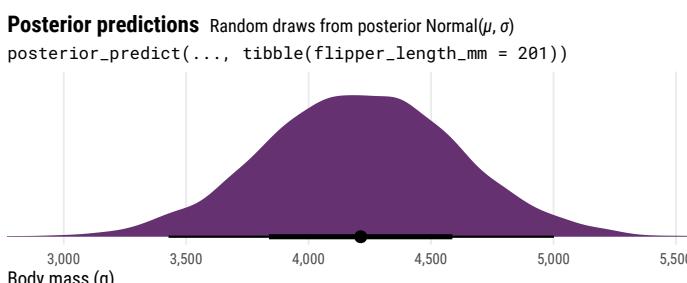
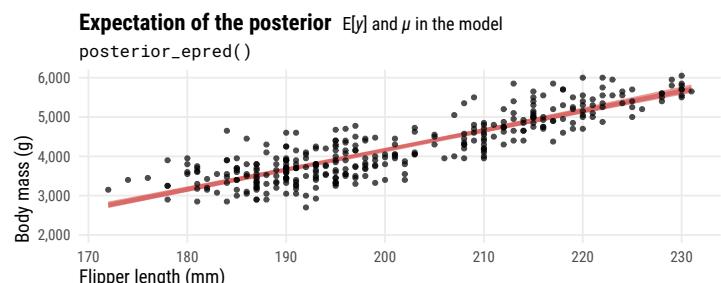
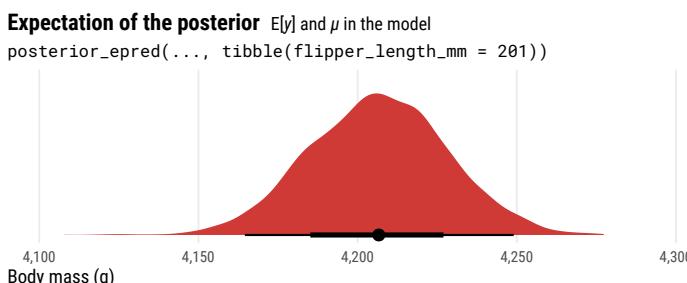
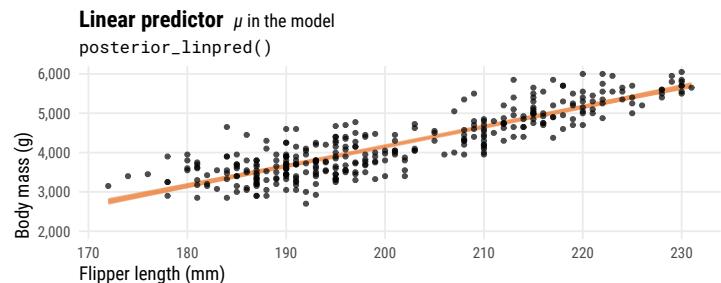
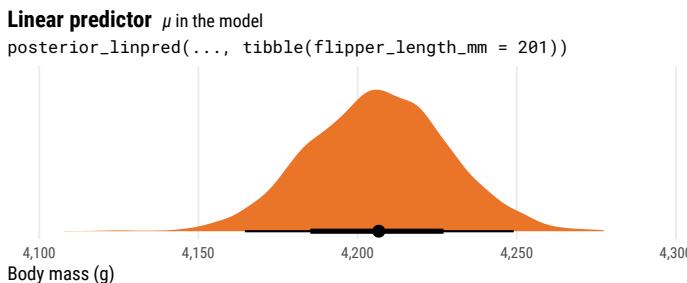
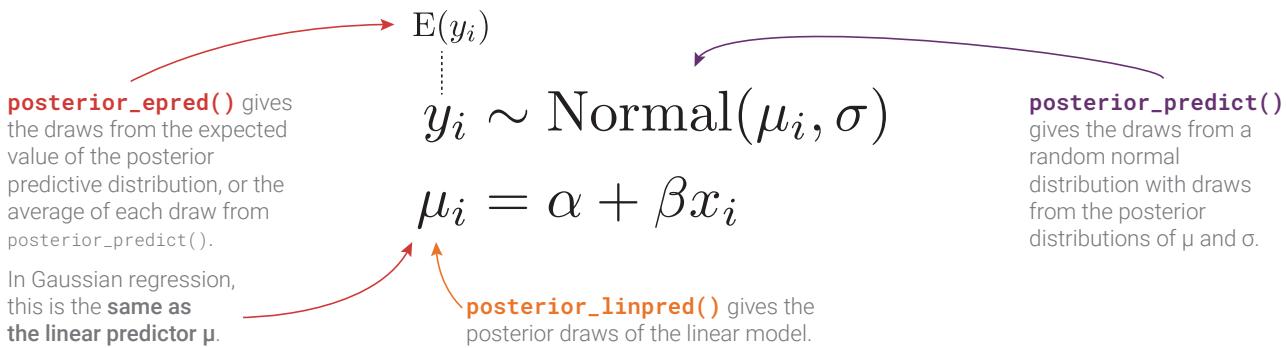


Gaussian regression example

Normal or Gaussian models are roughly equivalent to standard frequentist ordinary least squares (OLS) models and are generally interpreted the same way.

All the Bayesian prediction functions return values on the original scale of the outcome. In Gaussian models, the results from `posterior_epred()` and `posterior_linpred()` are identical. The draws from `posterior_predict()` will generally have the same mean as the linear predictor, but with more variance since the overall model σ is incorporated into the predictions.

```
model_normal <- brm(
  bf(body_mass_g ~ flipper_length_mm),
  family = gaussian(),
  data = penguins
)
```



Logistic regression example (generalized linear model with link transformation)

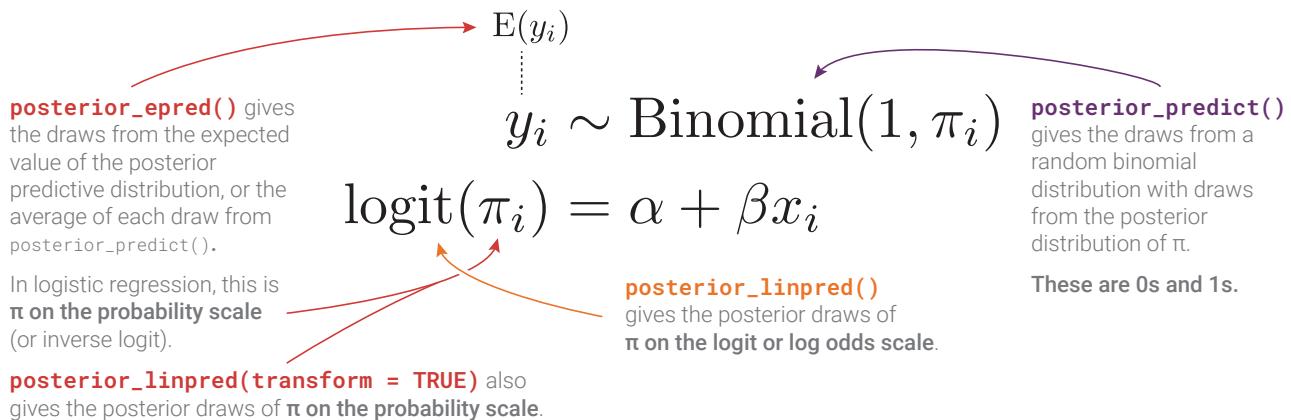
Generalized linear models (e.g., logistic, probit, ordered logistic, exponential, Poisson, negative binomial, etc.) use special link functions (e.g. logit, log, etc.) to transform the outcome variable into a scale that is more amenable to linear regression.

Estimates from these models can be used in their transformed scales (e.g., log odds in logistic regression) or can be back-transformed into their original scale (e.g., probabilities in logistic regression).

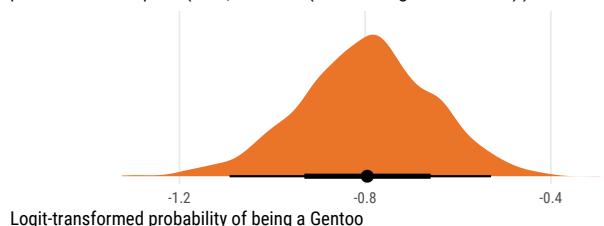
When working with links, the various Bayesian prediction functions return values on different scales, each corresponding to different parts of the model.

```
penguins <- penguins |>
  mutate(is_gentoo = species == "Gentoo")

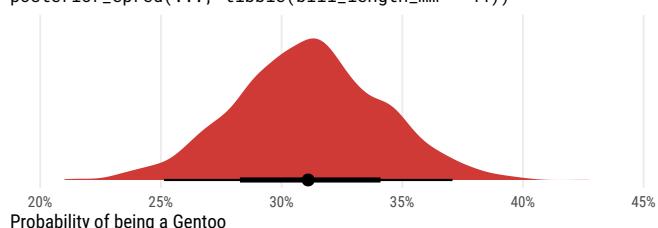
model_logit <- brm(
  bf(is_gentoo ~ bill_length_mm),
  family = bernoulli(link = "logit"),
  data = penguins
)
```



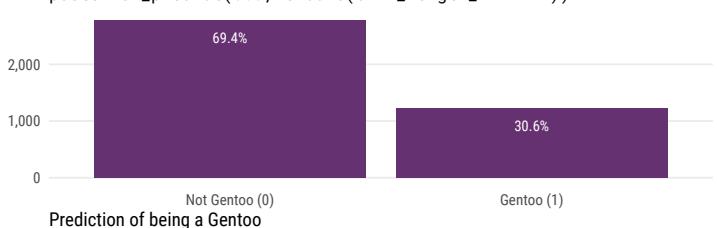
Linear predictor $\text{logit}(\pi)$ in the model
`posterior_linpred(..., tibble(bill_length_mm = 44))`



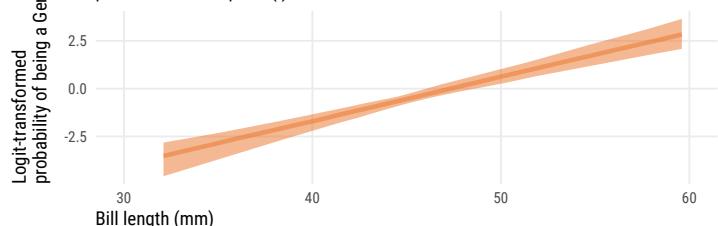
Expectation of the posterior $E[y]$ and π in the model
`posterior_epred(..., tibble(bill_length_mm = 44))`



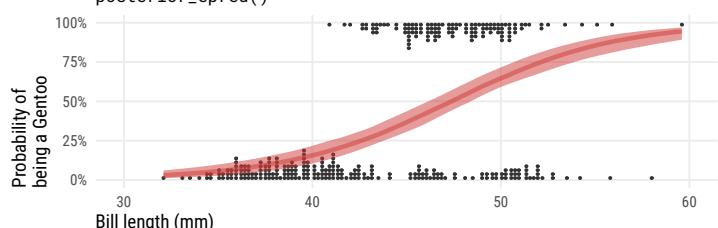
Posterior predictions Random draws from posterior $\text{Binomial}(1, \pi)$
`posterior_predict(..., tibble(bill_length_mm = 44))`



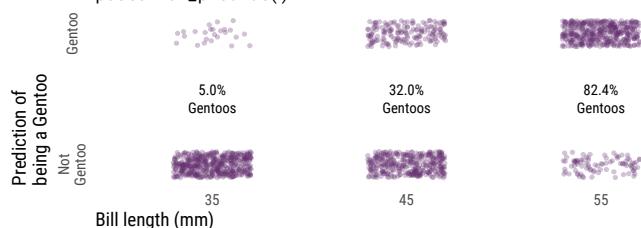
Linear predictor posterior $\text{logit}(\pi)$ in the model
`posterior_linpred()`



Expectation of the posterior $E[y]$ and π in the model
`posterior_epred()`



Posterior predictions Random draws from posterior $\text{Binomial}(1, \pi)$
`posterior_predict()`



Beta regression example (distributional model with multiple link transformations)

Regression models often focus solely on the location parameter of the model (e.g., μ in $\text{Normal}(\mu, \sigma)$; π in $\text{Binomial}(n, \pi)$). However, it is also possible to specify separate predictors for the scale or shape parameters of models (e.g., σ in $\text{Normal}(\mu, \sigma)$, ϕ in $\text{Beta}(\mu, \phi)$).

More complex models use a collection of distributional parameters (e.g., zero-inflated beta models estimate a mean μ , precision ϕ , and a zero-inflated parameter zi ; hurdle lognormal models estimate a mean μ , scale σ , and a hurdle parameter hu).

Even plain old Gaussian models become distributional models when a set of predictors is specified for σ (e.g. $\text{brm}(y \sim x_1 + x_2, \text{sigma} \sim x_2 + x_3)$)

When working with extra distributional parameters (dpar), the various Bayesian prediction functions return values on different scales for different components of the model.

