PRIOR PREDICTIVE DISTRIBUTION

- tells us how prior belief about θ, $p(\theta)$ translates into beliefs about $\tilde{y}$, according to our sampling model.

PICTURE:

$$\begin{tikzcd} \theta \arrow[d] & p(\theta) \arrow[l] \\ \tilde{y} & \end{tikzcd}$$

MATH:

$$p(\tilde{y}) = \int p(\tilde{y} \mid \theta)\, p(\theta)\, d\theta$$

pseudo-CODE:

```
Let S be large.
for (s in 1:S) {
draw θ^(s) ~ p(θ)
draw ỹ^(s) ~ p(ỹ | θ^(s))
}
```

POSTERIOR PREDICTIVE DISTRIBUTION

- tells us how our <u>posterior</u> beliefs about θ, $p(\theta|\vec{y})$, translate into beliefs about $\tilde{y}$.
- useful for <u>evaluating</u> our sampling model, $p(y|\theta)$.
- Does not "rule model in" but may help expose a flawed model. Note: this depends on many factors, e.g. the statistic of interest.

<u>PICTURE</u>:

$$
\begin{array}{c}
y_1, \dots, y_n := \vec{y} \\
\downarrow \\
\theta \\
\downarrow \\
\tilde{y}
\end{array}
$$

<u>MATH</u>:

$$p(\tilde{y}|y_1,\dots,y_n) = \int p(\tilde{y}|\theta)\, p(\theta|y_1,\dots,y_n)\, d\theta$$

<u>PSEUDO-CODE</u>: To approximate $p(\tilde{y}|y_1,\dots,y_n)$ w/ Monte Carlo:

Let S large.

for $(s$ in $1:S)$ {

  $\theta^{(s)} \sim p(\theta|\vec{y})$

  $\tilde{y}^{(s)} \sim p(\tilde{y}|\theta)$

}

Posterior Predictive Check

- Given the posterior predictive distr., compute the tail probability of some statistic.

Recall: definition of statistic $t(Y)$: a function of the data

random

ex: $t(Y) = \frac{1}{n} \sum Y_i$

ex: $t(Y) = \max\{Y_1, \ldots, Y_n\}$

PICTURE:

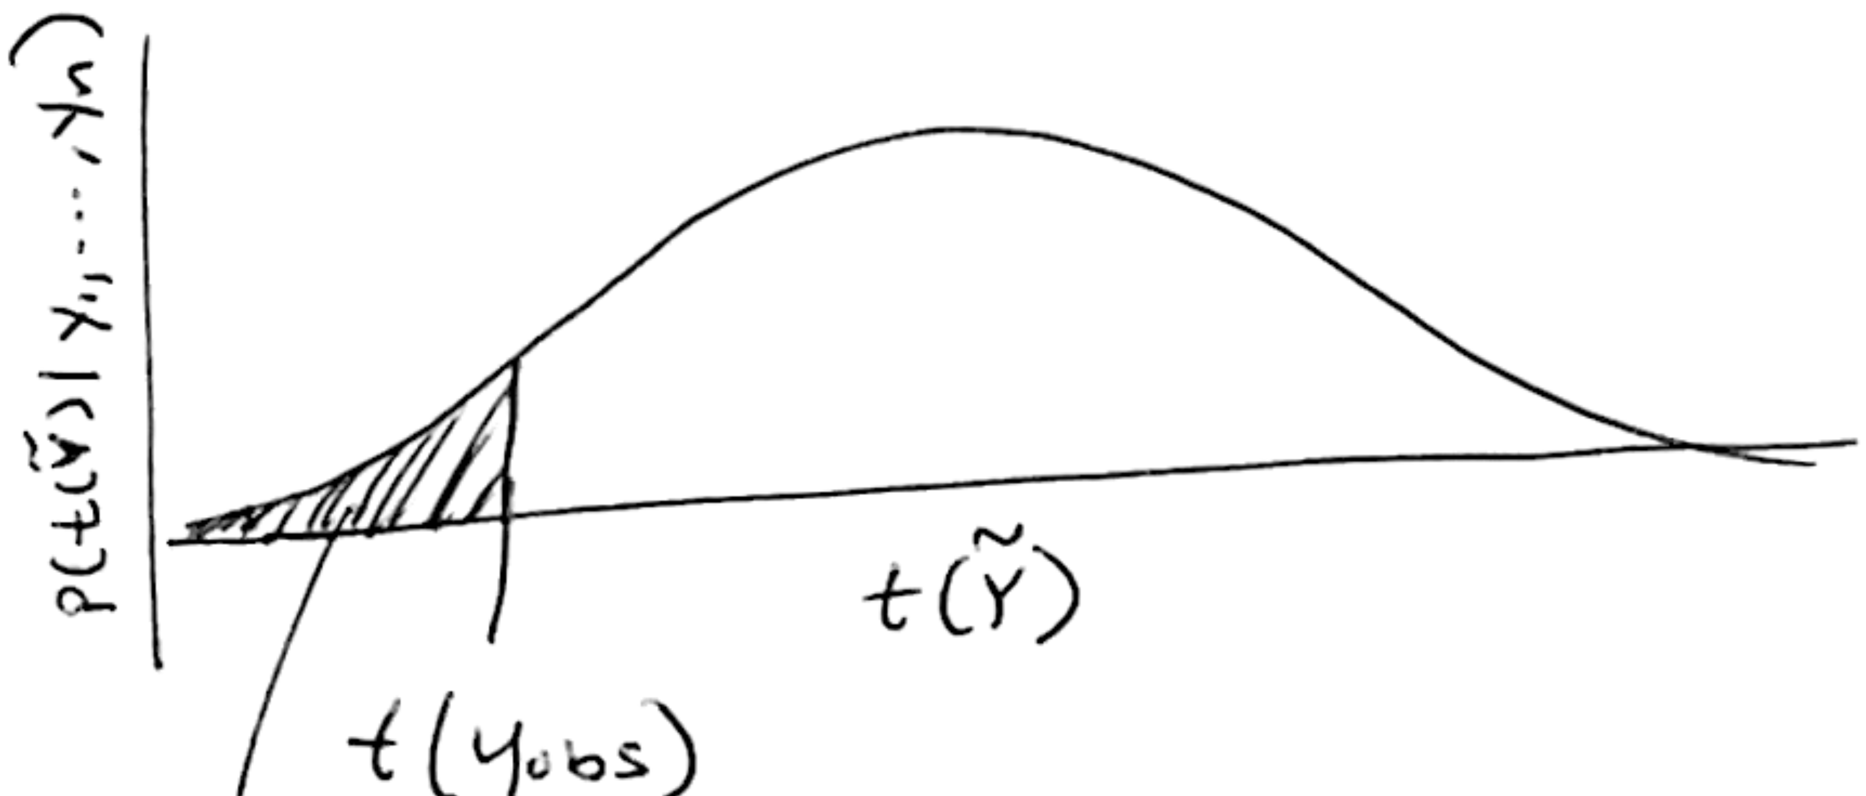

p-value: tail probability.

ex: $H_0: \theta = \theta_0$ (null hypothesis)

More generally: p-value $= p = \text{prob}\left(t(Y) \leq t(Y_{obs}) \,\middle|\, \text{something}\right)$

ex: $\theta = \hat{\theta}_{MLE}$ "parametric bootstrap"

$\{y_1, \ldots, y_n\}$ "posterior predictive"

* NOT distributed unif(0,1) like a traditional p-value!

posterior predictive p-value:*

$$p = \int_{-\infty}^{t(y_{obs})} \underbrace{p(t(\tilde{y})|y_1,\ldots,y_n)}_{\int_\Theta p(t(\tilde{y})|\theta)\,p(\theta|y_1,\ldots,y_n)\,d\theta} \, d\,t(\tilde{y})$$

POSTERIOR PREDICTIVE P-VALUE PSEUDOCODE:

To APPROXIMATE $\text{prob}(t(y) \leq t(y_{obs}) \mid y_1, \dots, y_n)$ using Monte Carlo:

```
Let S large.
for (s in 1:S) {
    sample θ^(s) ~ p(θ | y_1, ..., y_n)
    sample n ỹ ~ p(ỹ | θ^(s))
    compute & save t(ỹ^(s)_1, ..., ỹ^(s)_n)
}
Report mean(t(ỹ) < t(y_obs)).
```