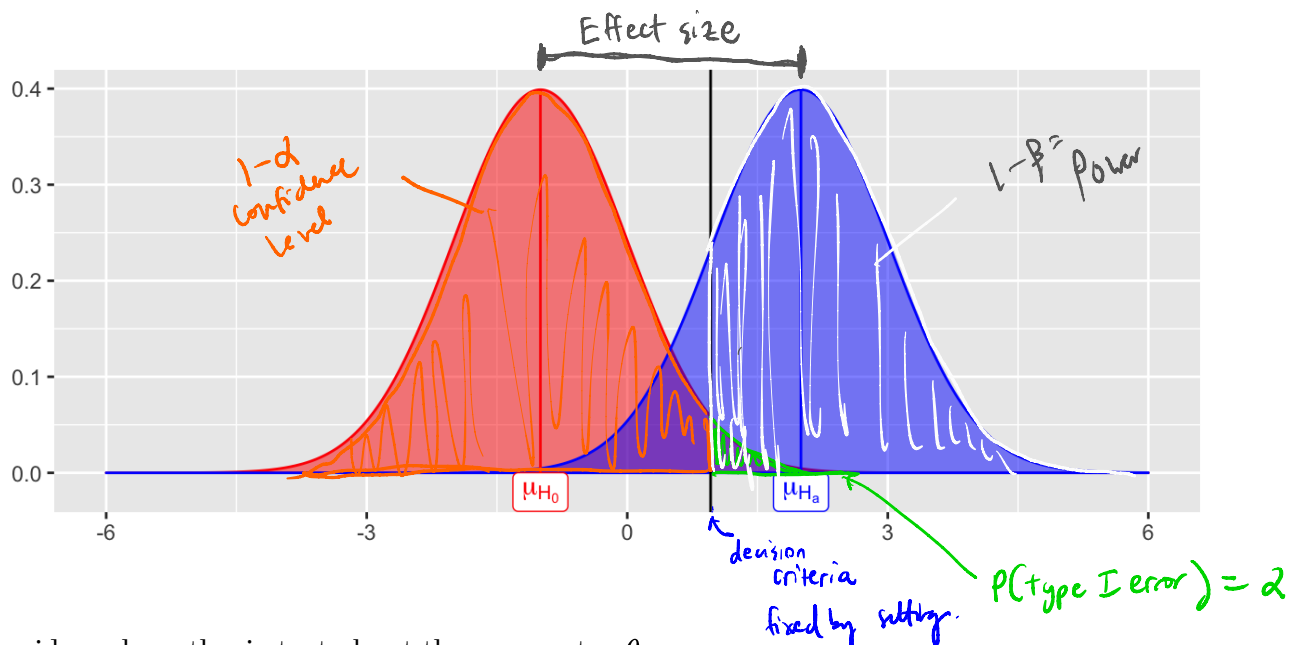


2.3 Power



Consider a hypothesis test about the parameter θ :

$$H_0 : \theta = \theta_0$$

$$H_a : \theta > \theta_0$$

We let $\beta = P(\text{fail to reject } H_0 | H_0 \text{ is false}) = P(\text{Type II error})$, then Power = $P(\text{reject } H_0 | H_0 \text{ is false}) = 1 - \beta$.

Power depends on the distance between the hypothesized value of the parameter θ_0 and the actual value θ_1 , so we can write $1 - \beta(\theta_1)$.

Why is power important?
 \hookrightarrow i.e. effect size.

1. If we have multiple statistical testing method for the same hypothesis, choose test w/ most power.
2. If you're going to spend time/money to do an experiment, need to check beforehand that your study will be powerful enough to detect an effect.

For a few simple cases, you can derive a closed form expression of power.

Example 2.4 Consider a one-sample z -test. Sample $X_1, \dots, X_n \stackrel{iid}{\sim} N(\mu, \sigma^2)$.

So power is a function of

- 1.
- 2.
- 3.
- 4.

2.4 MC Estimator of $1 - \beta$

Assume $X_1, \dots, X_n \sim F(\theta_0)$ (i.e., assume H_0 is true).

Then, we have the following hypothesis test –

$$\begin{aligned} H_0 : \theta &= \theta_0 \\ H_a : \theta &> \theta_0 \end{aligned}$$

and the statistics T^* , which is a test statistic computed from data. Then we **reject** H_0 if $T^* >$ the critical value from the distribution of the test statistic.

This leads to the following algorithm to estimate the power of the test $(1 - \beta)$

Your Turn


Consider data generated from the following mixture distribution:

$$f(x) = (1 - \epsilon)f_1(x) + \epsilon f_2(x), \quad x \in \mathbb{R}$$

where f_1 is the pdf of a $N(0, 1)$ distribution, f_2 is the pdf of a $N(0, 100)$ distribution, and $\epsilon \in [0, 1]$.

```
r_noisy_normal <- function(n, epsilon) {
  z <- rbinom(n, 1, 1 - epsilon)
  z*rnorm(n, 0, 1) + (1 - z)*rnorm(n, 0, 10)
}

n <- 100
data.frame(e = 0, sample = r_noisy_normal(n, 0)) %>%
  rbind(data.frame(e = 0.1, sample = r_noisy_normal(n, 0.1))) %>%
  rbind(data.frame(e = 0.6, sample = r_noisy_normal(n, 0.6))) %>%
  rbind(data.frame(e = 0.9, sample = r_noisy_normal(n, 0.9))) %>%
  ggplot() +
  geom_histogram(aes(sample)) +
  facet_wrap(~e, scales = "free")
```



We will compare the power of various tests of normality. Let F_X be the distribution of a random variable X . We will consider the following hypothesis test,

$$H_0 : F_x \in N \quad \text{vs.} \quad H_a : F_x \notin N,$$

where N denotes the family of univariate Normal distributions.

Recall Pearson's moment coefficient of skewness (See Example 2.2).

We will compare Monte Carlo estimates of power for different levels of contamination ($0 \leq \epsilon \leq 1$). We will use $\alpha = 0.1$, $n = 100$, and $m = 100$.

```
# skewness statistic function
skew <- function(x) {
  xbar <- mean(x)
  num <- mean((x - xbar)^3)
  denom <- mean((x - xbar)^2)
  num/denom^1.5
}

# setup for MC
alpha <- .1
n <- 100
m <- 100
epsilon <- seq(0, 1, length.out = 200)
var_sqrt_b1 <- 6*(n - 2)/((n + 1)*(n + 3)) # adjusted variance for
  skewness test
crit_val <- qnorm(1 - alpha/2, 0, sqrt(var_sqrt_b1)) #crit value for
  the test
empirical_pwr <- rep(NA, length(epsilon)) #storage

# estimate power for each value of epsilon
for(j in 1:length(epsilon)) {
  # perform MC to estimate empirical power
  ## Your turn

}

## store empirical se
empirical_se <- "Your Turn: fill this in"

## plot results --
## x axis = epsilon values
## y axis = empirical power
## use lines + add band of estimate +/- se
```

Compare the power with $n = 100$ to the power with $n = 10$. Make a plot to compare the two for many values of ϵ .