hypothesis testing

Does smoking cause cancer?

(a) No; we don't know what causes cancer, but smokers are no more likely to get it than non-smokers

(b) Yes; a much greater % of smokers get it

Note: even in case (b), "cause" is a stretch, but for simplicity, "causes" and "correlates with" will be loosely interchangeable today

Programmers using the Eclipse IDE make fewer errors

(a) Hooey. Errors happen, IDE or not.

(b) Yes. On average, programmers using Eclipse produce code with fewer errors per thousand lines of code

Black Tie Linux has way better web-server throughput than Red Shirt.

(a) Ha! Linux is linux, throughput will be the same

(b) Yes. On average, Black Tie response time is 20% faster.

This coin is biased!

(a) "Don't be paranoid, dude. It's a fair coin, like any other, P(Heads) = 1/2"

(b) "Wake up, smell coffee: P(Heads) = 2/3, totally!"

How do we decide?

Design an experiment, *gather* data, *evaluate*:

In a sample of N smokers + non-smokers, does % with cancer differ? Age at onset? Severity?

In N programs, some written using IDE, some not, do error rates differ?

Measure response times to N individual web transactions on both.

In N flips, does putative biased coin show an unusual excess of heads? More runs? Longer runs?

A complex, multi-faceted problem. Here, emphasize evaluation:

What N? How large of a difference is convincing?

General framework:

1. Data
2. H_0 – the "null hypothesis"
3. H_1 – the "alternate hypothesis"
4. A decision rule for choosing between H_0/H_1 based on data
5. Analysis: What is the probability that we get the right answer?

Example:

100 coin flips

$P(H) = 1/2$

$P(H) = 2/3$

"if #H $\leq$ 60, accept null, else reject null"

$P(H \leq 60 \mid 1/2) = ?$
$P(H > 60 \mid 2/3) = ?$

By convention, the null hypothesis is usually the "simpler" hypothesis, or "prevailing wisdom." E.g., Occam's Razor says you should prefer that unless there is *strong* evidence to the contrary.

Is coin fair (1/2) or biased (2/3)? How to decide? Ideas:

1. Count: Flip 100 times; if number of heads observed is ≤ 60, accept H_0
   or ≤ 59, or ≤ 61 ... $\Rightarrow$ different error rates
2. Runs: Flip 100 times. Did I see a longer run of heads or of tails?
3. Runs: Flip until I see either 10 heads in a row (reject H_0) or 10 tails is a row (accept H_0)
4. Almost-Runs: As above, but 9 of 10 in a row
5. . . .

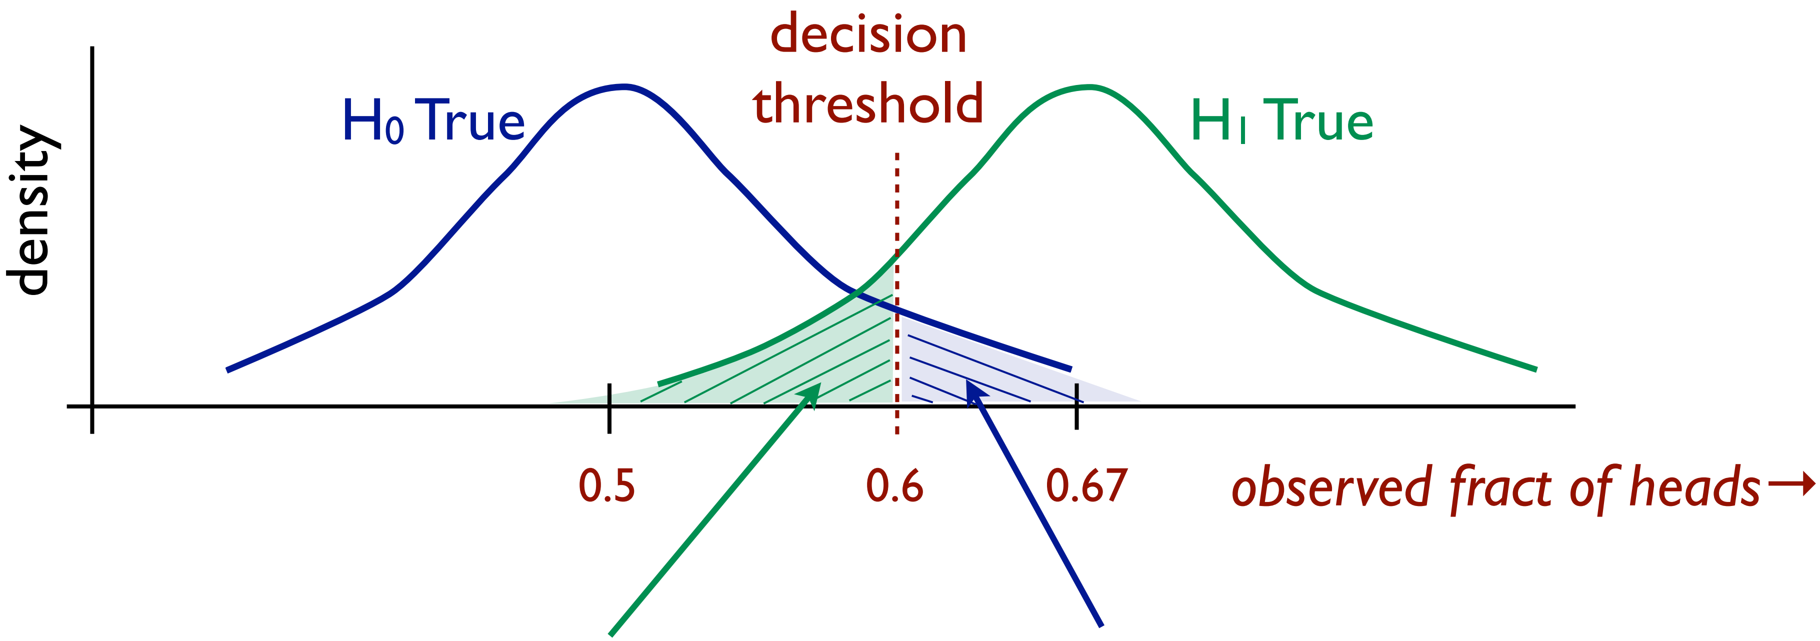

Type II error: false accept;
accept H_0 when it is false.
β = P(type II error)

Type I error: false reject;
reject H_0 when it is true.
α = P(type I error)

Goal: make both α, β small (but it's a tradeoff; they are interdependent).
$\alpha \le 0.05$ common in scientific literature.

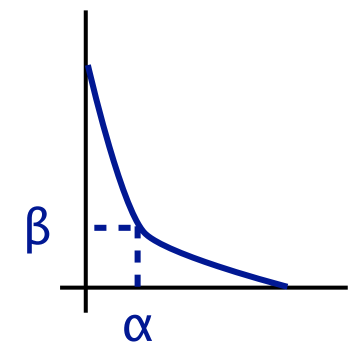

One general approach: a "Likelihood Ratio Test"

$$\frac{L(x_1, x_2, \ldots, x_n \mid H_1)}{L(x_1, x_2, \ldots, x_n \mid H_0)} :: c \quad \begin{cases} < c & \text{accept } H_0 \\ = c & \text{arbitrary} \\ > c & \text{reject } H_0 \end{cases}$$

E.g.:

c = 1: accept H_0 if observed data is *more* likely under that hypothesis than it is under the alternate

c = 5: accept H_0 unless there is *strong* evidence that the alternate is more likely (i.e. 5 x)

Changing the threshold c shifts α, β, of course.

Given: A coin, either fair (p(H)=1/2) or biased (p(H)=2/3)

Decide: which

How? Flip it 5 times. Suppose outcome D = HHHTH

Null Model/Null Hypothesis M_0: p(H) = 1/2

Alternative Model/Alt Hypothesis M_1: p(H) = 2/3

Likelihoods:

$P(D \mid M_0) = (1/2)\ (1/2)\ (1/2)\ (1/2)\ (1/2) = 1/32$

$P(D \mid M_1) = (2/3)\ (2/3)\ (2/3)\ (1/3)\ (2/3) = 16/243$

Likelihood Ratio: $$\frac{p(D \mid M_1)}{p(D \mid M_0)} = \frac{16/243}{1/32} = \frac{512}{243} \approx 2.1$$

I.e., alt model is ≈ 2.1x more likely than null model, given data

The Neyman-Pearson Lemma

If an LRT for some simple hypotheses H_0 versus H_1 has error probabilities α, β, then any test with type I error $\alpha' \leq \alpha$ must have type II error $\beta' \geq \beta$

In other words, to compare a simple hypothesis to a simple alternative, a likelihood ratio test will be as good as any for a given error bound.

example (cont.)

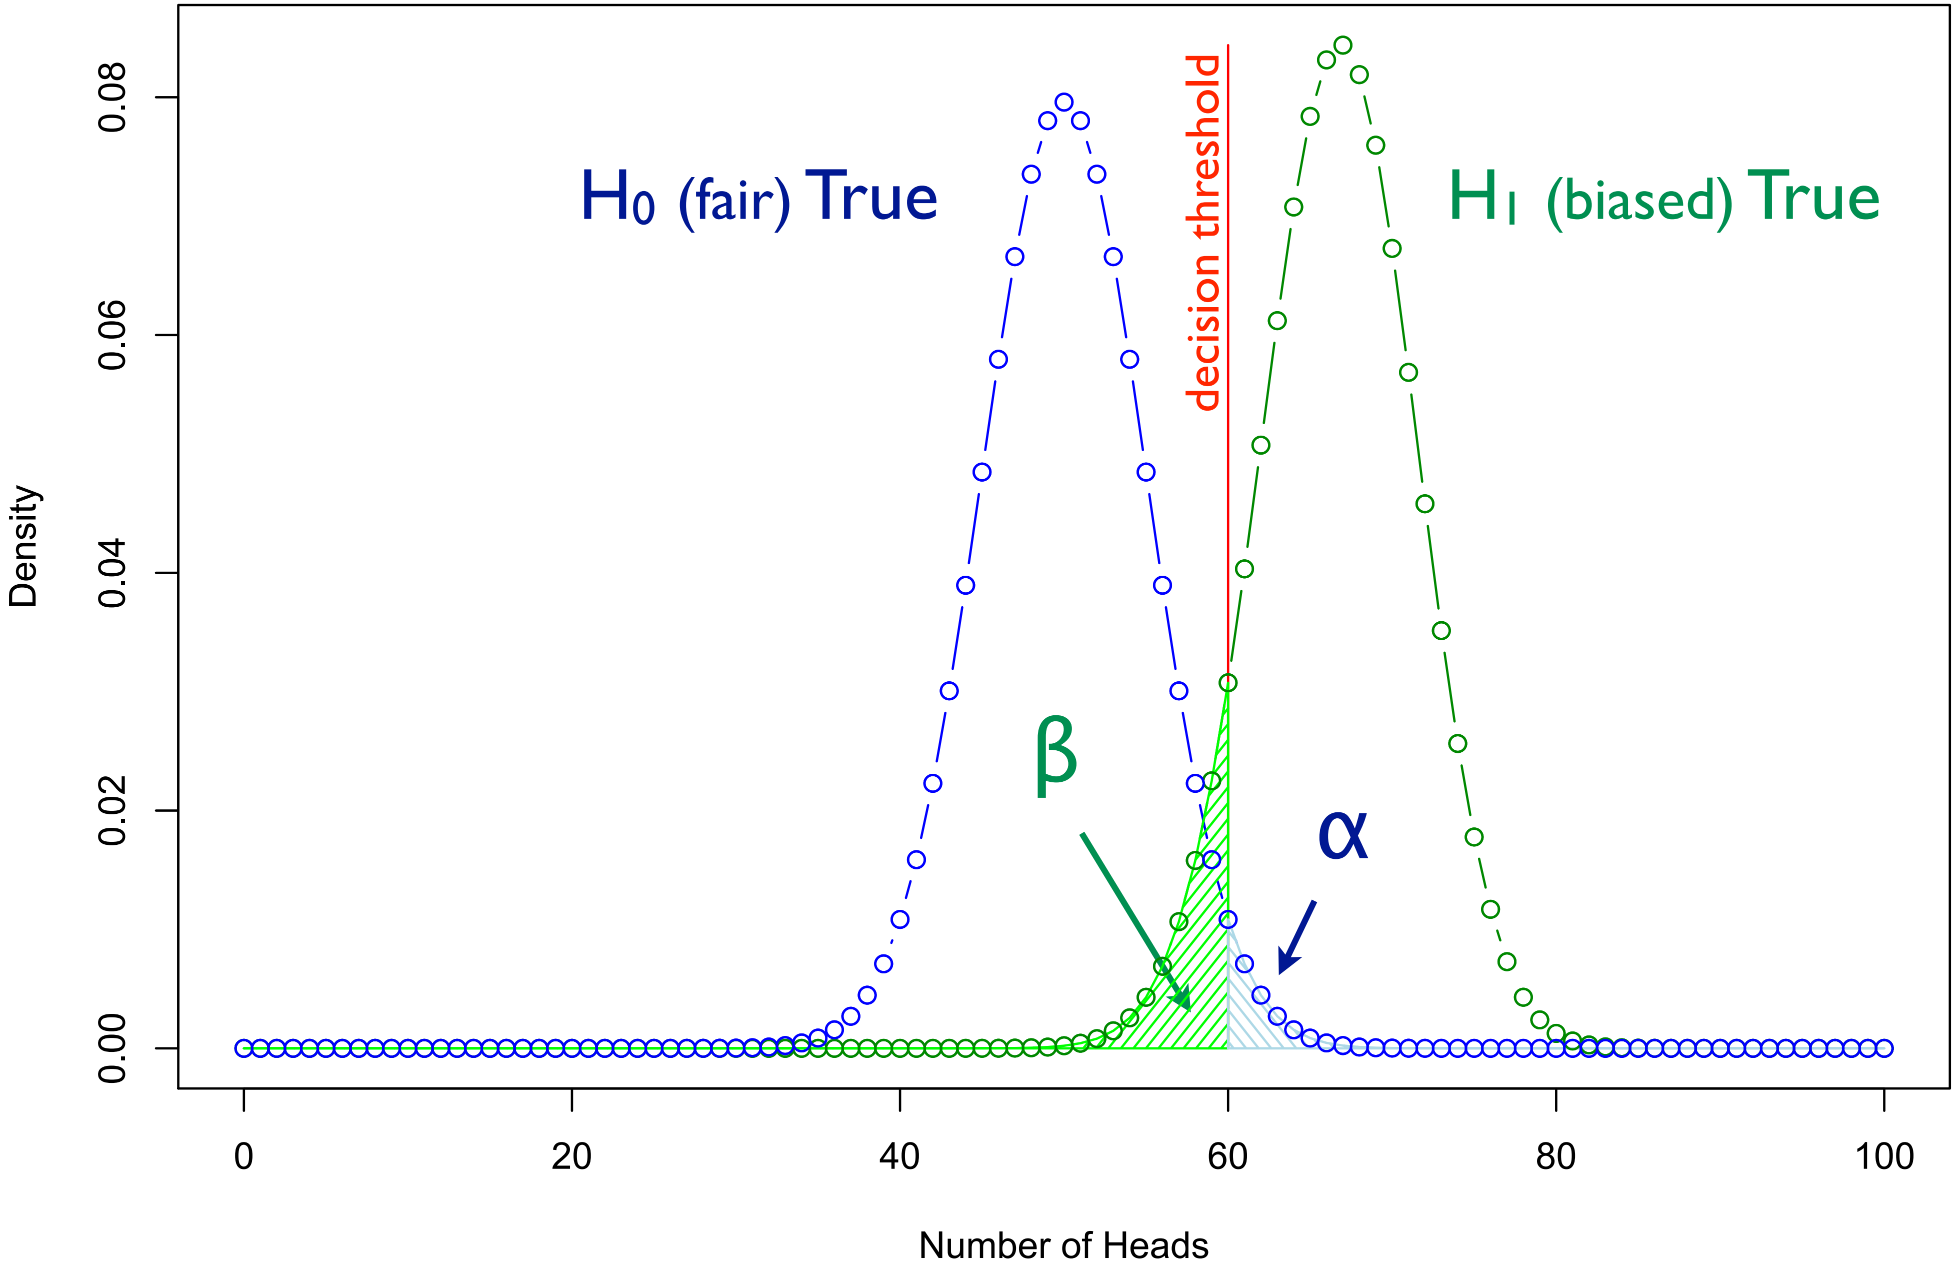

Null/Alternative hypotheses - specify distributions from which data are assumed to have been sampled

Simple hypothesis - one distribution

E.g., "Normal, mean = 42, variance = 12"

Decision rule; "accept/reject null if sample data..."; *many* possible

Type 1 error: false reject/reject null when it is true

Type 2 error: false accept/accept null when it is false

α = P(type 1 error), β = P(type 2 error)

Likelihood ratio tests: for simple null vs simple alt, compare ratio of likelihoods under the 2 competing models to a fixed threshold.

Neyman-Pearson: LRT is best possible in this scenario.