STAT 226 Lecture 3
Small Sample Binomial Inference
Section 1.4.3

Yibi Huang
Department of Statistics
University of Chicago

Example: Medical Consultants for Organ Donors

- People providing an organ for donation sometimes seek the help of a special “medical consultant.” These consultants assist the patient in all aspects of the surgery, with the goal of reducing the possibility of complications during the medical procedure and recovery.
- One consultant tried to attract patients by noting the average complication rate for liver donor surgeries in the US is about 10%, but her clients have only had 3 complications in the 62 liver donor surgeries she has facilitated.
- Is this strong evidence that her work meaningfully contributes to reducing complications (and therefore she should be hired!)?

Example: Medical Consultants for Organ Donors (Cont'd)

- H_0: $\pi = 0.1$ vs. H_a: $\pi < 0.1$
- estimate of π is $\hat{\pi} = 3/62 \approx 0.048$

Example: Medical Consultants for Organ Donors (Cont'd)

- H_0: $\pi = 0.1$ vs. H_a: $\pi < 0.1$
- estimate of π is $\hat{\pi} = 3/62 \approx 0.048$
- Wald, score, likelihood ratio tests are based on *large samples*: only appropriate when *numbers of successes and failures are both at least 10* (or 15), but there were only 3 successes (having complications) in this example

Example: Medical Consultants for Organ Donors (Cont'd)

- H_0: $\pi = 0.1$ vs. H_a: $\pi < 0.1$
- estimate of π is $\hat{\pi} = 3/62 \approx 0.048$
- Wald, score, likelihood ratio tests are based on *large samples*: only appropriate when *numbers of successes and failures are both at least 10* (or 15), but there were only 3 successes (having complications) in this example
- For small sample, one can use the exact distribution of the data — **Binomial**, instead of its normal approximation.

Example: Medical Consultants for Organ Donors (Cont'd)

- H_0: $\pi = 0.1$ vs. H_a: $\pi < 0.1$
- estimate of π is $\hat{\pi} = 3/62 \approx 0.048$
- Wald, score, likelihood ratio tests are based on *large samples*: only appropriate when *numbers of successes and failures are both at least 10* (or 15), but there were only 3 successes (having complications) in this example
- For small sample, one can use the exact distribution of the data — **Binomial**, instead of its normal approximation.
- Under H_0: number of complications $\sim Bin(n = 62, \pi = 0.1)$

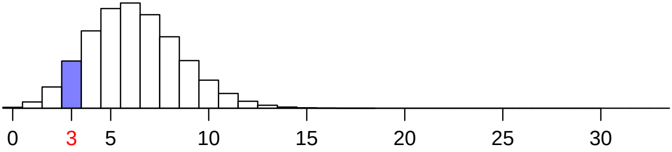

Exact Binomial Tests

For conventional large sample tests based on normal approximation, the lower one sided P-value is the area under the normal curve below 3

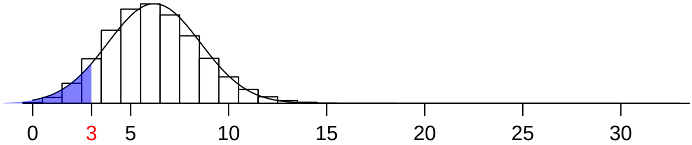

Exact Binomial Tests

For conventional large sample tests based on normal approximation, the lower one sided P-value is the area under the normal curve below 3

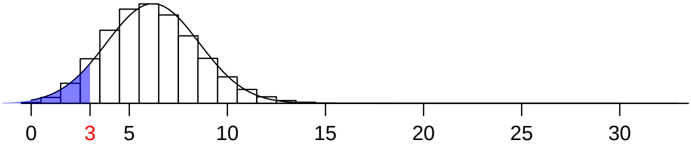

For the exact binomial test, the lower one-sided P-value is the area under the probability histogram below 3.

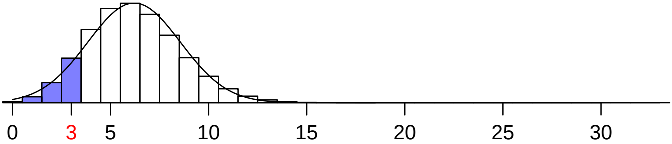

Exact Binomial Tests

Let Y = number of complications among the 62 liver donors.

$Y \sim Binomial(n = 62, \pi = 0.1)$ under H_0.

$$P(Y = k) = \binom{62}{k}(0.1)^k(0.9)^{62-k}$$

The **lower one-sided *P*-value** for exact binomial test of $\pi = 0.1$ is

$$\begin{aligned}
P(Y \leq 3) &= P(Y = 0) + P(Y = 1) + P(Y = 2) + P(Y = 3) \\
&= \binom{62}{0}(0.1)^0(0.9)^{62} + \binom{62}{1}(0.1)^1(0.9)^{61} + \binom{62}{2}(0.1)^2(0.9)^{60} + \binom{62}{3}(0.1)^3(0.9)^{59} \\
&= 0.1210
\end{aligned}$$

```r
dbinom(0:3, size=62, p=0.1)
[1] 0.001456 0.010027 0.033981 0.075514
sum(dbinom(0:3, size=62, p=0.1))
[1] 0.121
```

Not enough evidence to support the consultant's claim.

Exact Binomial Tests in R

The R function to do exact binomial test is `binom.test()`.

```
binom.test(3, 62, p=0.1, alternative="less")


    Exact binomial test


data:  3 and 62
number of successes = 3, number of trials = 62, p-value = 0.121
alternative hypothesis: true probability of success is less than 0.1
95 percent confidence interval:
 0.0000000 0.1203362
sample estimates:
probability of success
             0.0483871
```

The p-value given by R is 0.121, which agrees with our calculation.

P-values of Exact Binomial Tests

For testing H_0: $\pi = \pi_0$, suppose the observed binomial count is y_{obs}.

- P-value $= P(Y \leq y_{obs}) = \sum_{k \leq y_{obs}} \binom{n}{k} \pi_0^k (1-\pi_0)^{n-k}$ for a lower one-sided alternative H_a: $\pi < \pi_0$
- P-value $= P(X \geq y_{obs}) = \sum_{k \geq y_{obs}} \binom{n}{k} \pi_0^k (1-\pi_0)^{n-k}$ for a upper one-sided alternative H_a: $\pi > \pi_0$
- For a two-sided alternative H_a: $\pi \neq \pi_0$, the P-value is the sum of all the $P(Y = k)$ such that $P(Y = k) \leq P(Y = y_{obs})$

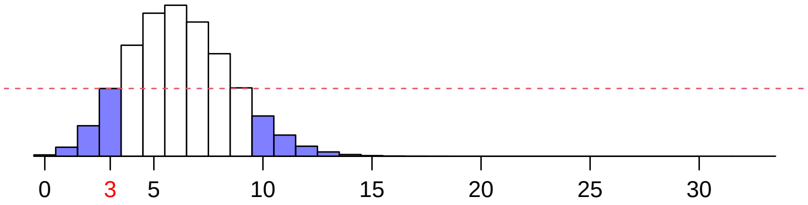

Example: Medical Consultants for Organ Donors (Cont'd)

In this example, the observed count y_{obs} is 3.

As $P(Y = 9) > P(Y = 3)$ and $P(Y = k) < P(Y = 3)$ for all $k \geq 10$, the two-sided P-value is

$$P(Y \leq 3) + P(Y \geq 10) \approx 0.1210 + 0.0872 = 0.2082$$

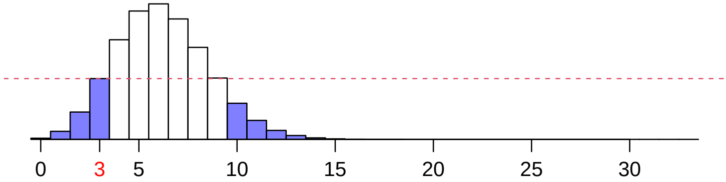

Note that the two-sided P-value for an exact binomial test may not be twice of the one-sided P-value since a binomial distribution may not be symmetric

```
k = 0:12
prob = dbinom(k, 62, 0.1)  # P(Y=k) for k=0,1,2,...,11
data.frame(k,prob)
    k     prob
1   0 0.001456
2   1 0.010027
3   2 0.033981
4   3 0.075514
5   4 0.123760
6   5 0.159512
7   6 0.168374
8   7 0.149666
9   8 0.114328
10  9 0.076219
11 10 0.044884
12 11 0.023576
13 12 0.011133
```

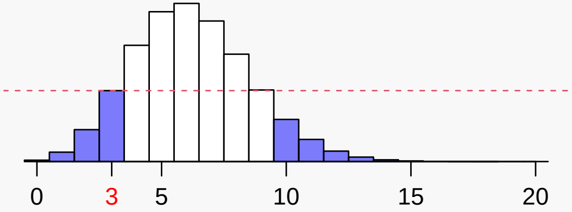

Two-Sided Exact Binomial Tests in R

```
binom.test(3, 62, p=0.1, alternative="two.sided")


    Exact binomial test


data:  3 and 62
number of successes = 3, number of trials = 62, p-value = 0.2081
alternative hypothesis: true probability of success is not equal to 0.1
95 percent confidence interval:
 0.01009195 0.13496195
sample estimates:
probability of success
             0.0483871
```

The P-value given by R 0.2081 agrees with our calculation.

Exact Binomial Confidence Intervals

- Just like Wald, score, or LRT confidence intervals, one can invert the two-sided exact binomial test to construct confidence intervals for π.
- The $100(1-\alpha)\%$ exact binomial confidence interval for π is the collection of those π_0 such that the two-sided P-value for testing H_0: $\pi = \pi_0$ using the exact binomial test is at least α.
- The computation of the exact binomial confidence interval is tedious to do by hand, but easier for a computer.
- For the medical consultant example, the R command `binom.test()` gives the 95% exact confidence interval $(0.01009195, 0.13496195)$ for π from the R output in the previous slide However, this interval is not obtained by inverting a two-sided exact Binomial test.

```r
binom.test(3, 62, p=0.01009195, conf.level=0.95, alternative="two.sided
[1] 0.02500002
binom.test(3, 62, p=0.13496195, conf.level=0.95, alternative="two.sided
[1] 0.04121624
```

Neither P-values equal to 0.05