Stat310

Confidence intervals

Hadley Wickham

1. Test
2. Roadmap
3. Properties of ML estimators
4. Confidence intervals
5. New distribution (the t-distribution)

Test 2

Grading in progress

Bivariate change of variables much better :) But I accidentally made it too hard :(So bonus points available :)

Roadmap

What we want to do

Given data:

- Estimate true value of parameter
  (last week)
- **Quantify uncertainty of estimate**
  (this week)
- Test whether true value is a certain value
  (next week)

Last week

I repeated an experiment defined by Poisson(λ) 10 times, and recorded the following results: 6 11 10 6 12 7 8 5 7 10

What is a good estimator of λ?

This week

I repeated an experiment defined by Poisson(λ) 10 times, and recorded the following results: 6 11 10 6 12 7 8 5 7 10

What is a 95% confidence interval for λ?

Next week

I repeated an experiment defined by Poisson(λ) 10 times, and recorded the following results: 6 11 10 6 12 7 8 5 7 10

Is $\lambda > 10$?

MLE properties

MLE's are

Invariant

Consistent

Asymptotically normally distributed

$$Var(\hat{\theta}_{ML}) = \frac{-1}{E\frac{\delta^2}{\delta\theta^2}l(X|\theta)}$$

But

That math is too hard for this course :(

So we need some other ways to work out how much error our estimators have.

Your turn

What is the variance of $\hat{\lambda}_{ML}$?

Recall what this estimator is. What do we know about it?

Other estimators

We will derive a few other estimators and their distributions. Most of the time we will use the CLT. Sometime we will have to do more work, and may need **new distributions** (t, chi-square).

Confidence intervals

CI

Point estimator has probability 0 of being correct. Need to do better!

Create a range of possible values so that we have a X% of including the true value.

A 95% confidence has a 95% chance of containing the true value.

Set up

I repeated an experiment defined by Poisson(λ) 10 times, and recorded the following results: 6 11 10 6 12 7 8 5 7 10

The MLE of λ is 8.2, and its variance is 0.82.

What is the distribution of the estimate? Can you construct an interval that will contain λ 95% of the time?

Steps

1. **Identify distribution that connects estimator and true value.**

2. Form confidence interval for known (sampling) distribution, and work out bounds. Use plug-in principle for anything we don't know.

3. Back transform.

4. Write as interval.

5. Plug in sample estimates (actual numbers).

X_i iid, and n large:

$$\frac{\bar{X}_n - E(\bar{X})}{sd(\bar{X})} \dot{\sim} Z$$

and even more so

$$\frac{\bar{X}_n - E(\bar{X})}{\hat{sd}(\bar{X})} \dot{\sim} Z$$

Your turn

Work through the steps on the handout.

There are many ways of coming up with a confidence interval

http://www.jstor.org/pss/3087376

How do we choose?

A confidence interval is a simple numerical summary of the uncertainty of an estimate.

A 95% confidence interval will contain the true value 95% of the time.

An additional constraint is that we want the confidence interval to be a short as possible.

Each line = 95% confidence interval from one experiment

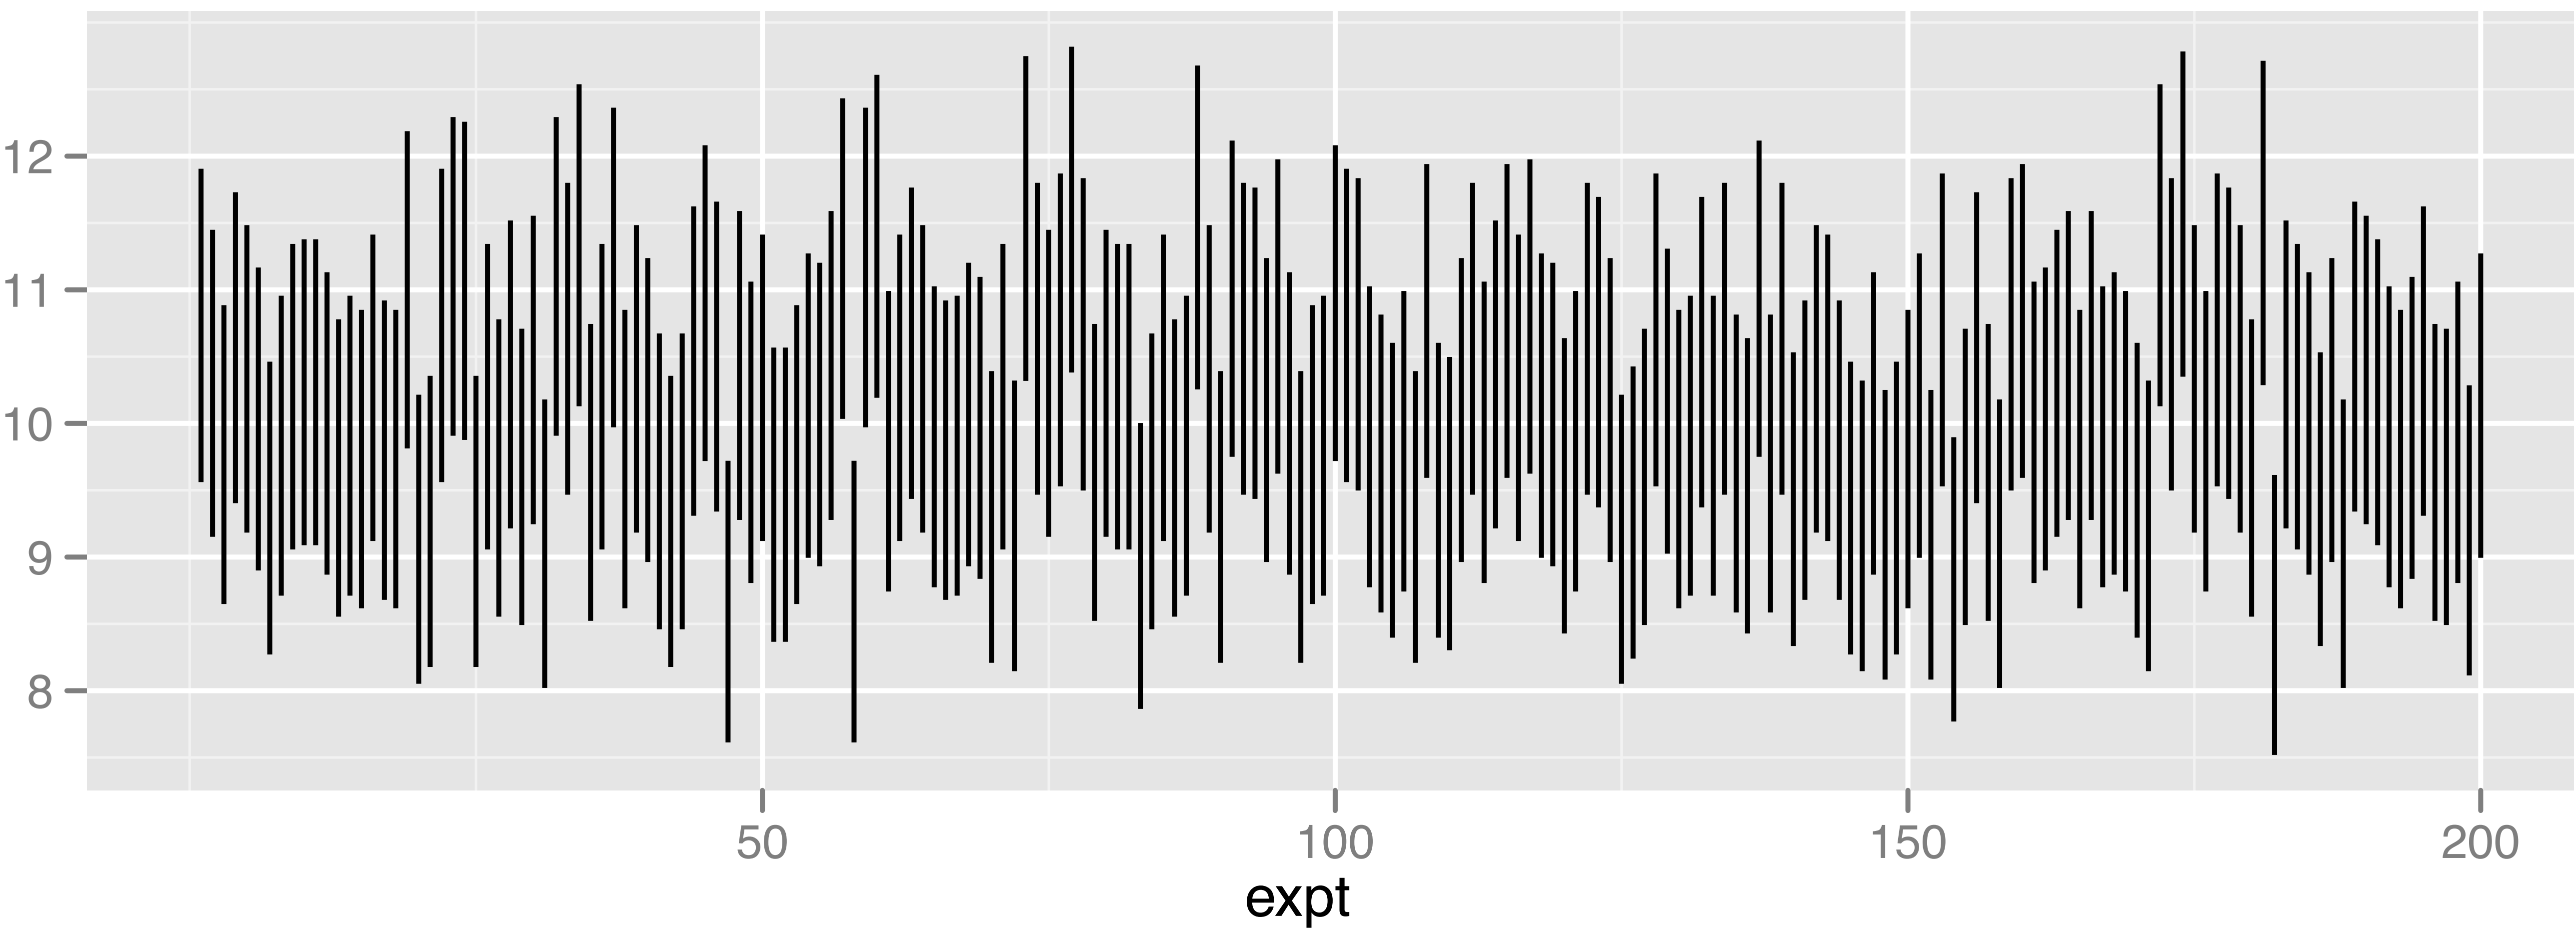

One experiment = 10 repeats

Horizontal line = true value

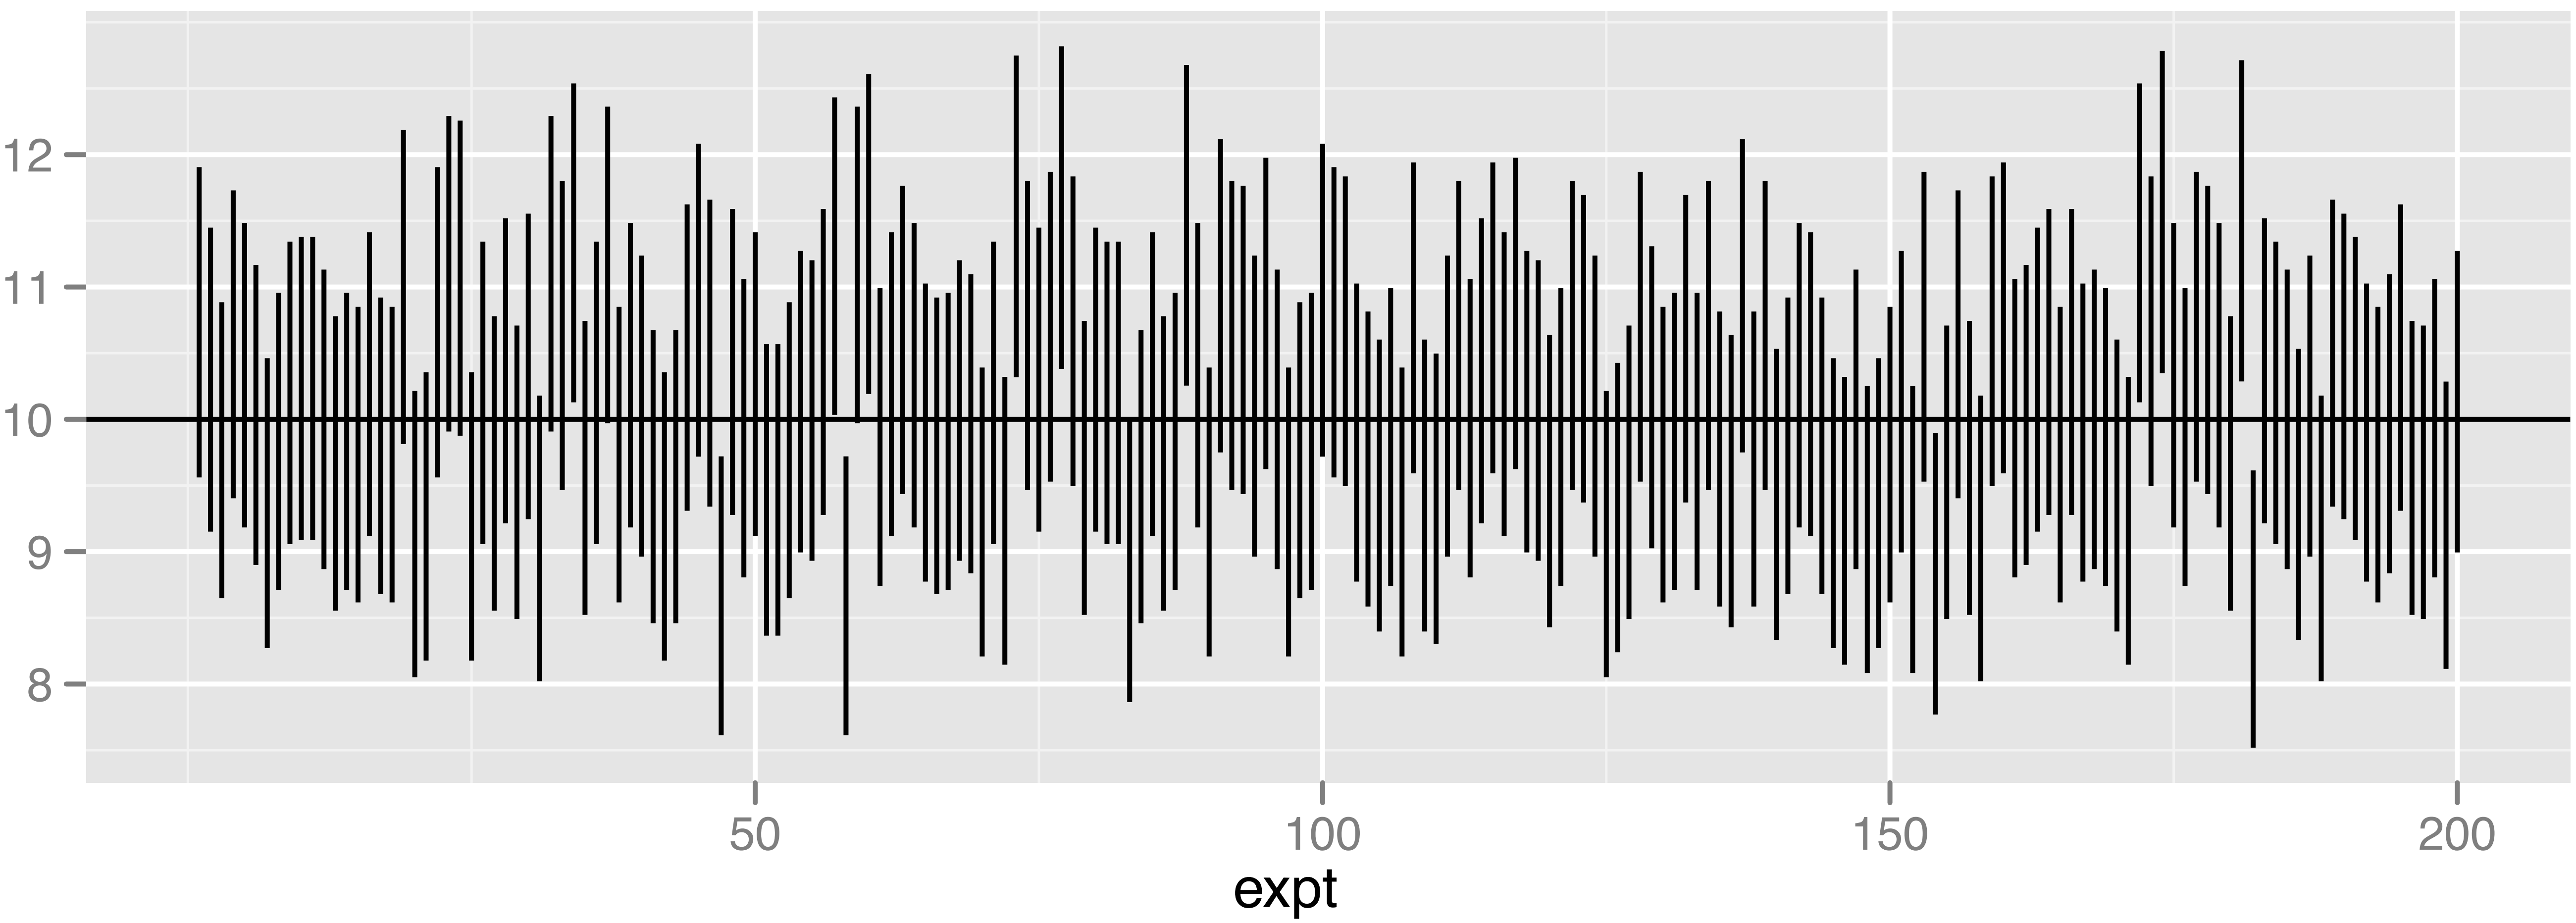

Red intervals don't include true value

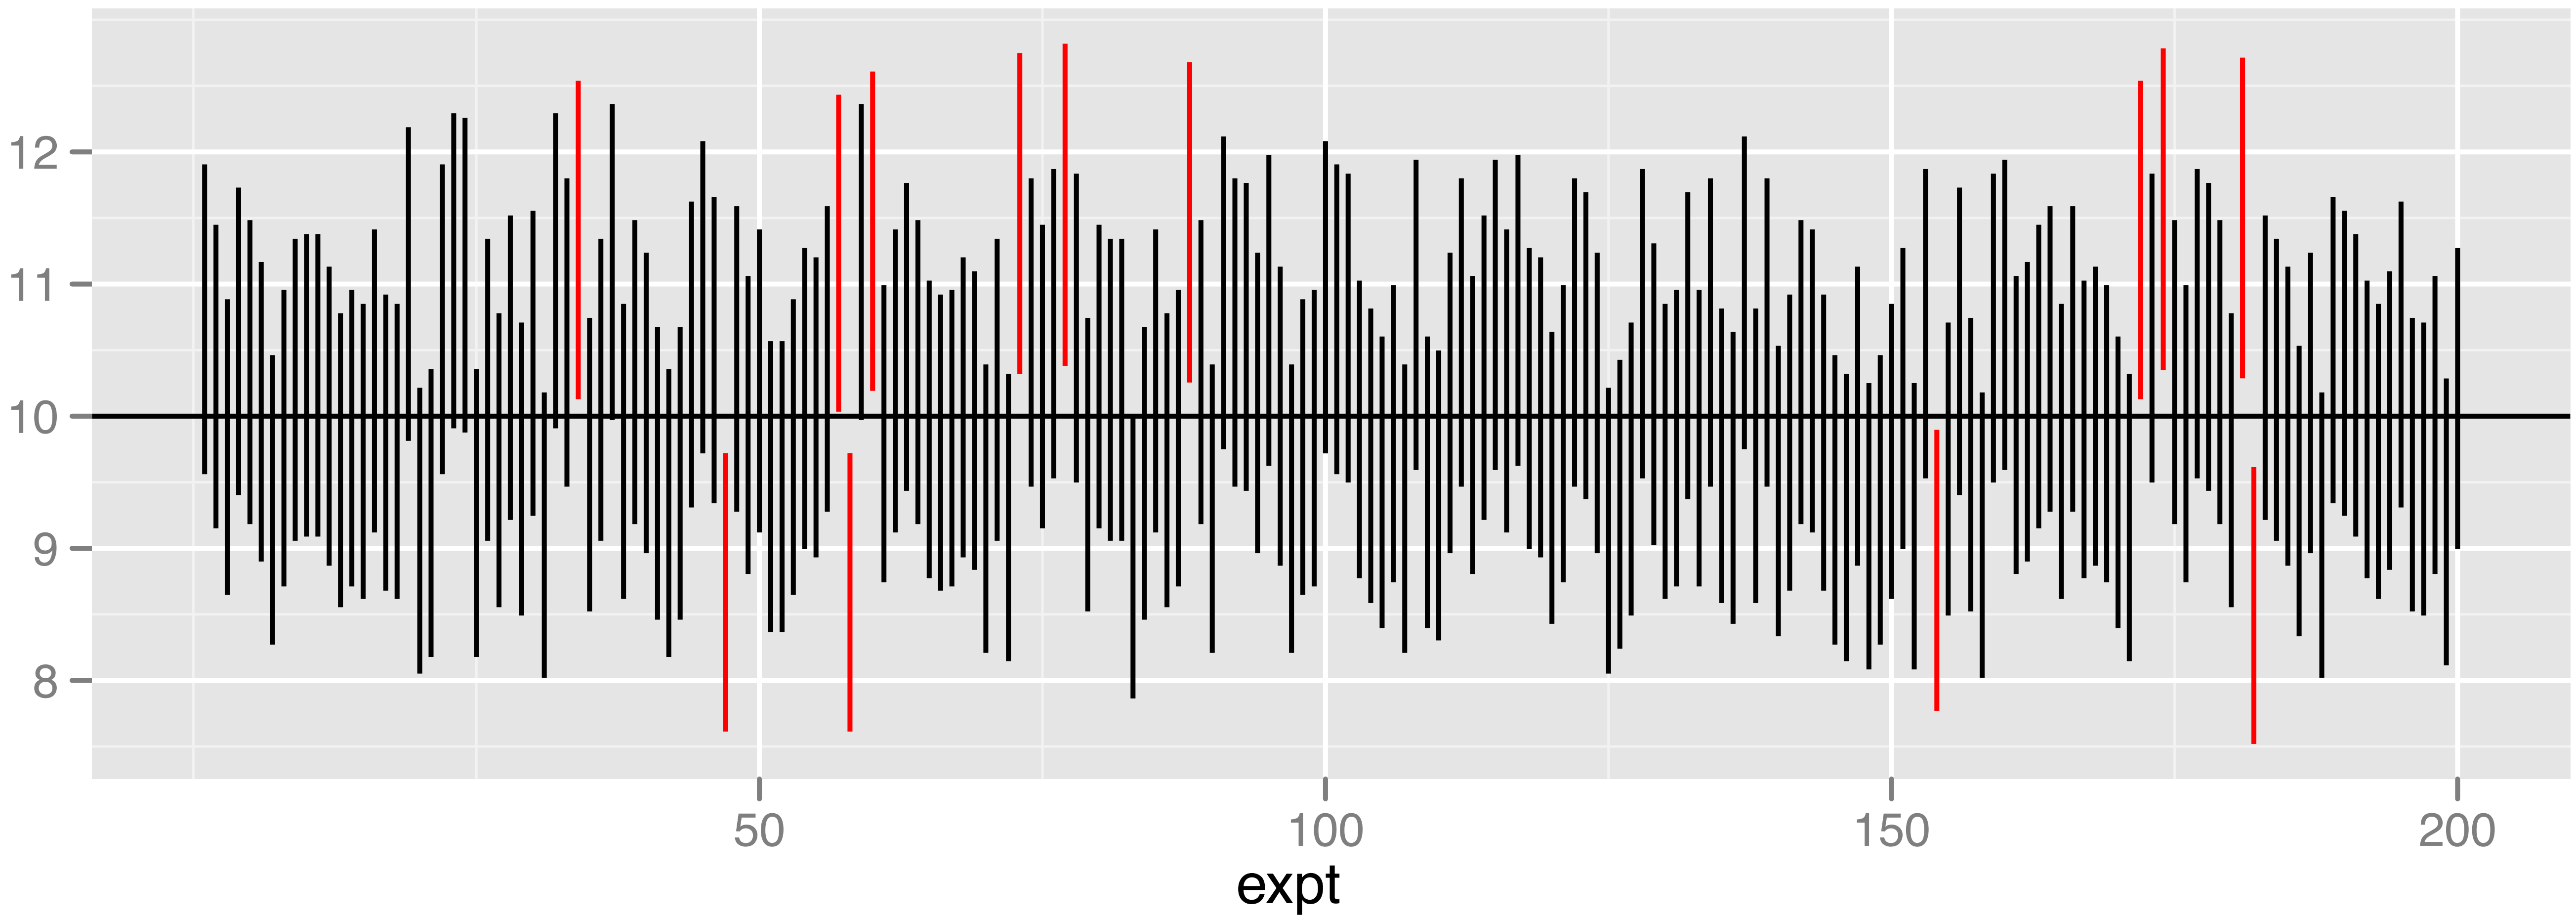

There are 13 red lines and 200 experiments. Is this an ok interval?

Your turn

What if the experimenter only did two replicates? What do you expect will happen to the confidence intervals?

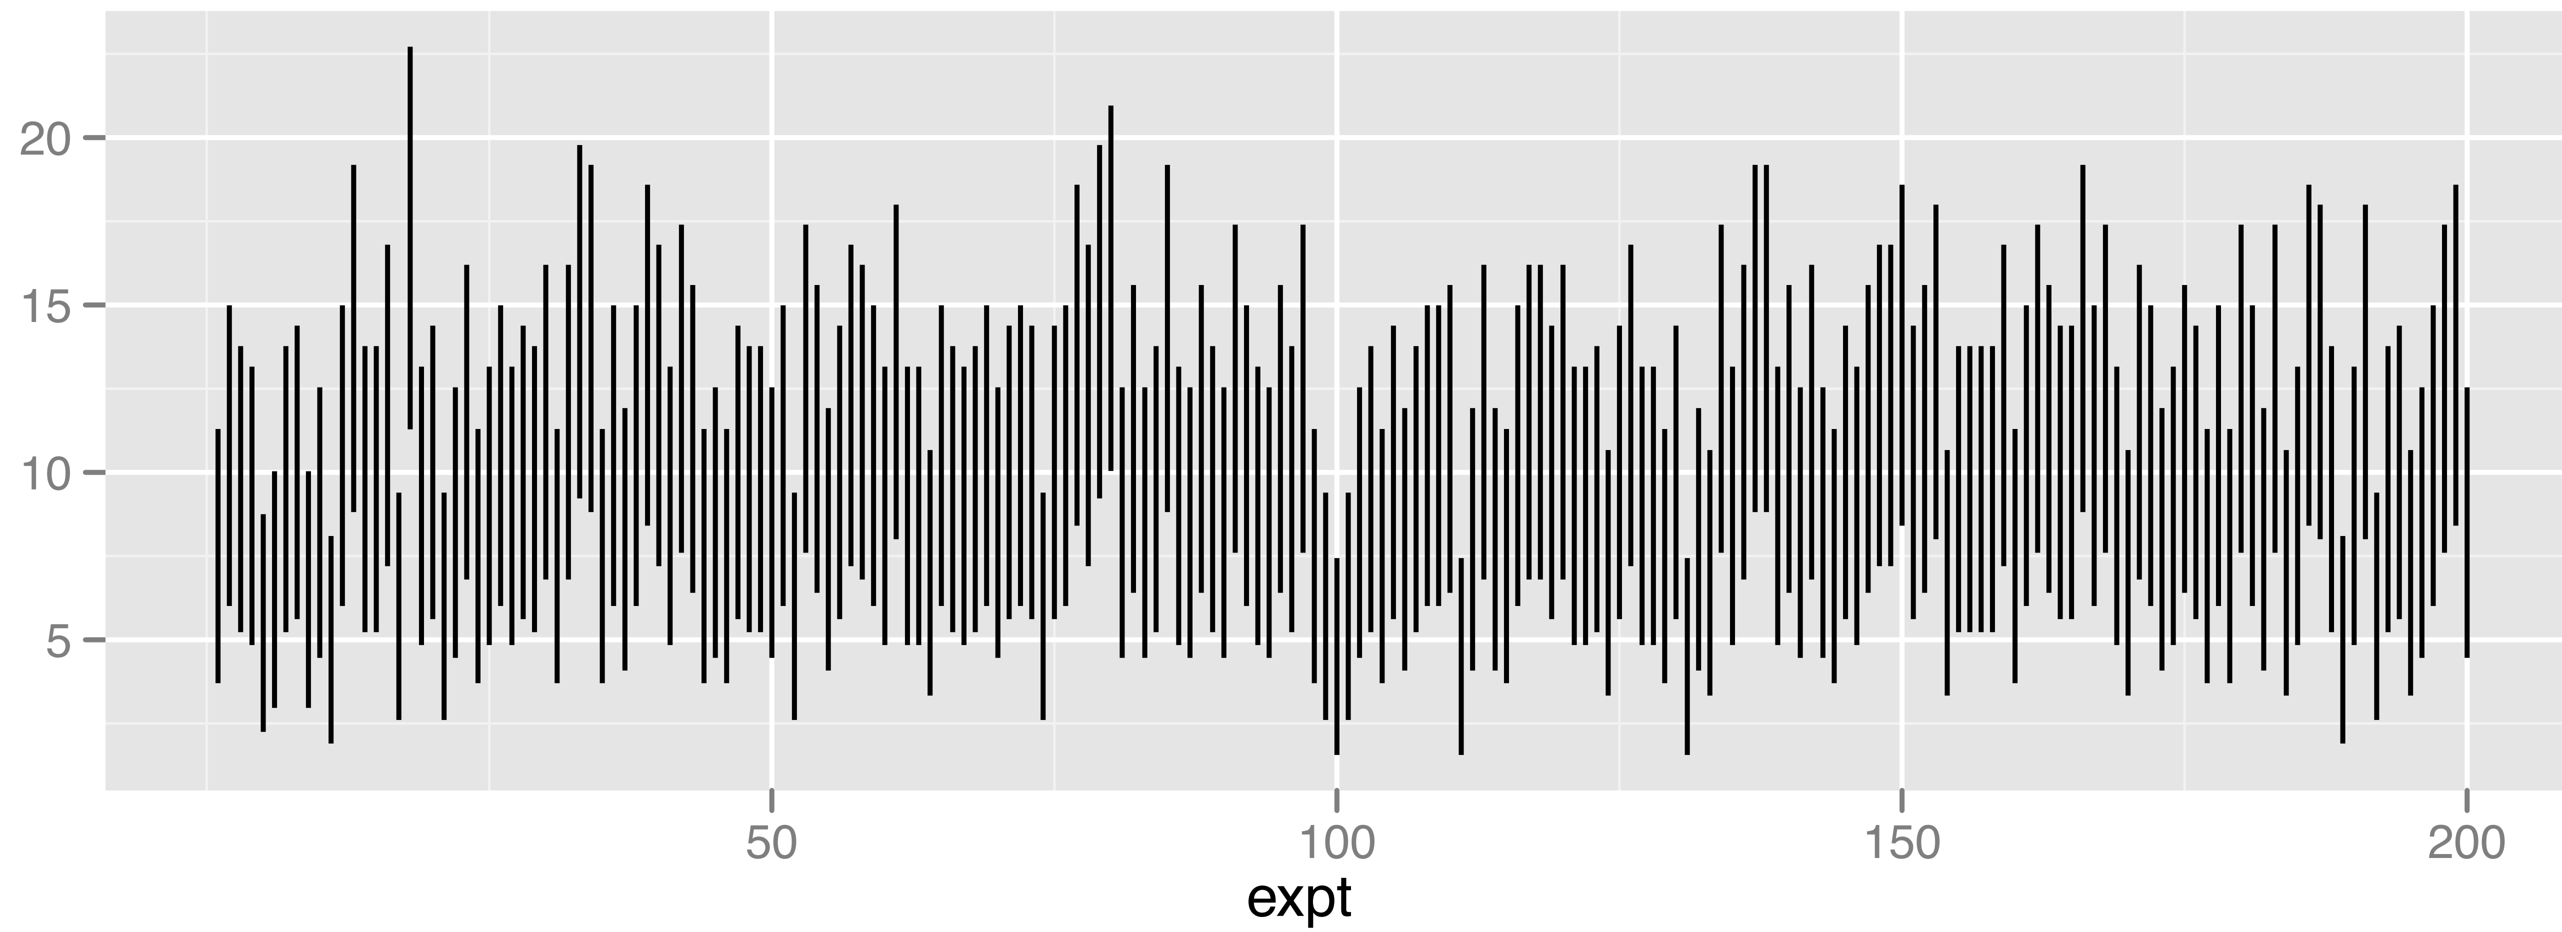

20
15
10
5
50
100
150
200
expt

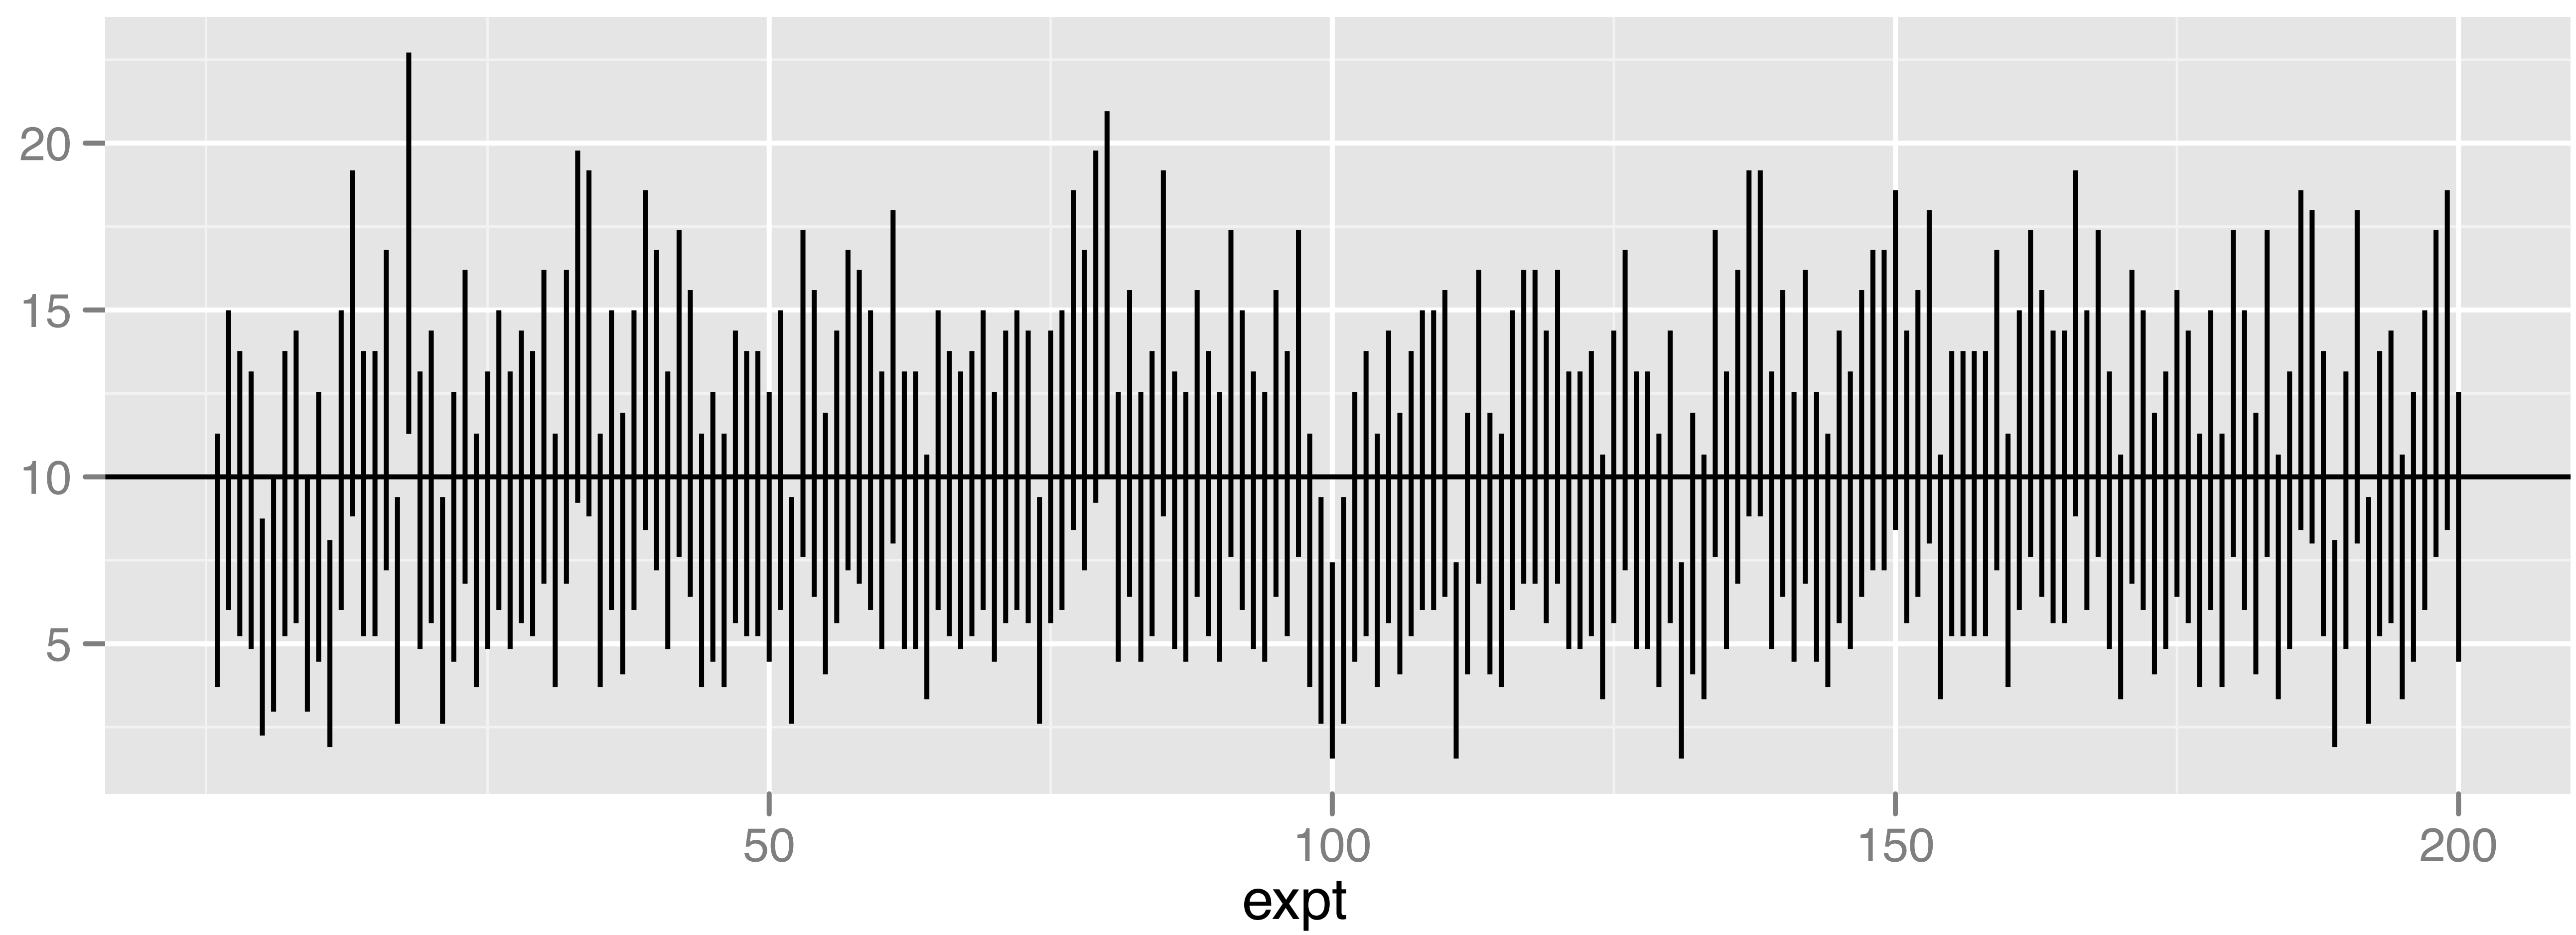

20
15
10
5
50
100
150
200
expt

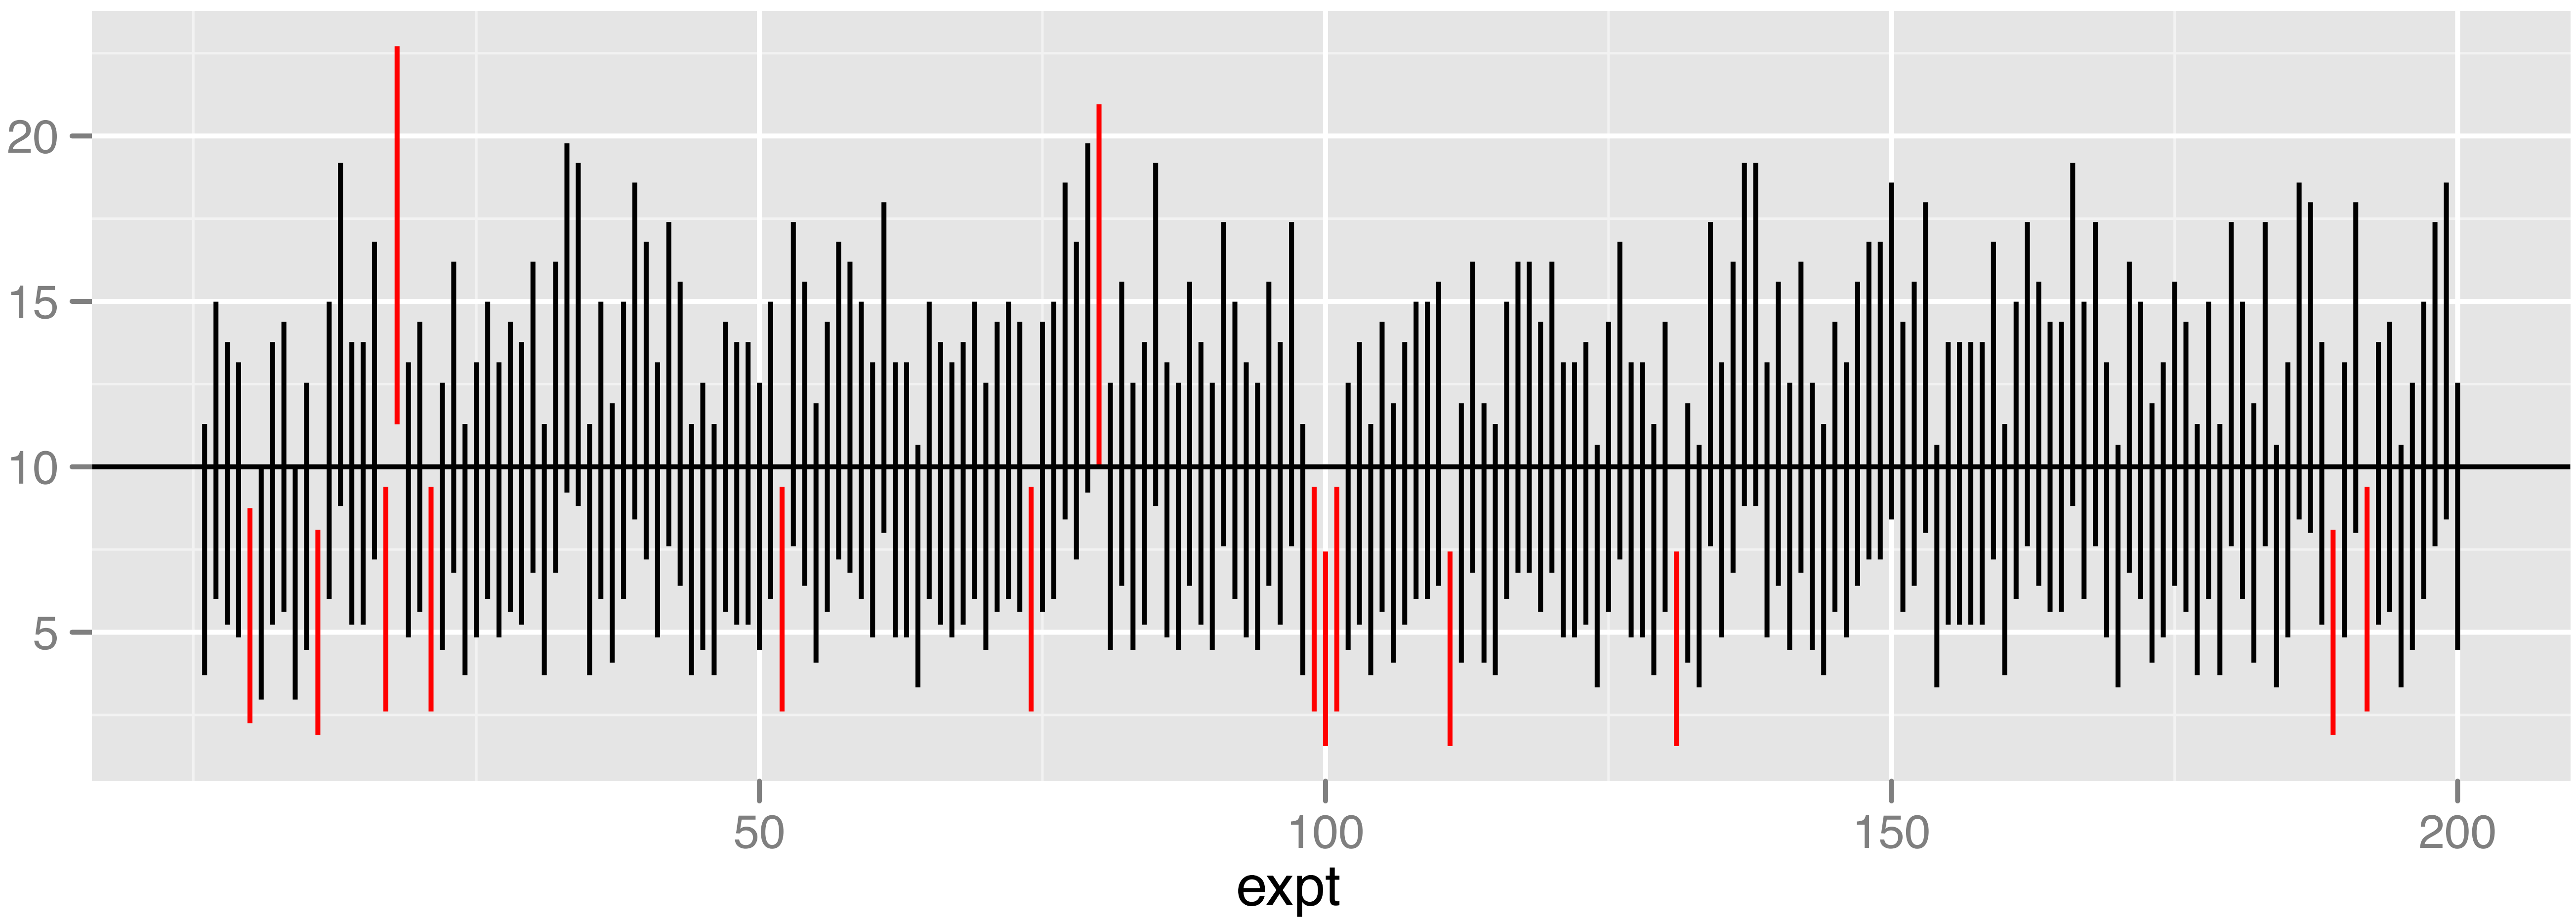

There are 15 red lines and 200 experiments. Is this an ok interval?

Your turn

What's wrong with a statement like this:

$P(8 < \lambda < 10) = 0.95$?

t-distribution

Steps

Identify distribution that connects estimator and true value.

Form confidence interval for known (sampling) distribution.

Write as probability statement.

Back transform.

Write as interval.

What other distributions can we use?

$$X_i \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$$

$$\frac{\bar{X}_n - E(\bar{X})}{sd(\bar{X})} \sim Z$$

but

$$\frac{\bar{X}_n - E(\bar{X})}{\hat{sd}(\bar{X})} \sim t_{n-1}$$

$$X_i \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2)$$

$$\frac{\bar{X}_n - E(\bar{X})}{sd(\bar{X})} \sim Z$$

but

$$\frac{\bar{X}_n - E(\bar{X})}{\hat{sd}(\bar{X})} \sim \boxed{t_{n-1}}$$

New distribution!

Parameter: v (nu, degrees of freedom)

don't need to recall

$$f(x, \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

$$\mathrm{E}(X) = 0$$

If v > 1

$$\mathrm{Var}(X) = \frac{\nu}{\nu - 2}$$

If v > 2

No mgf!

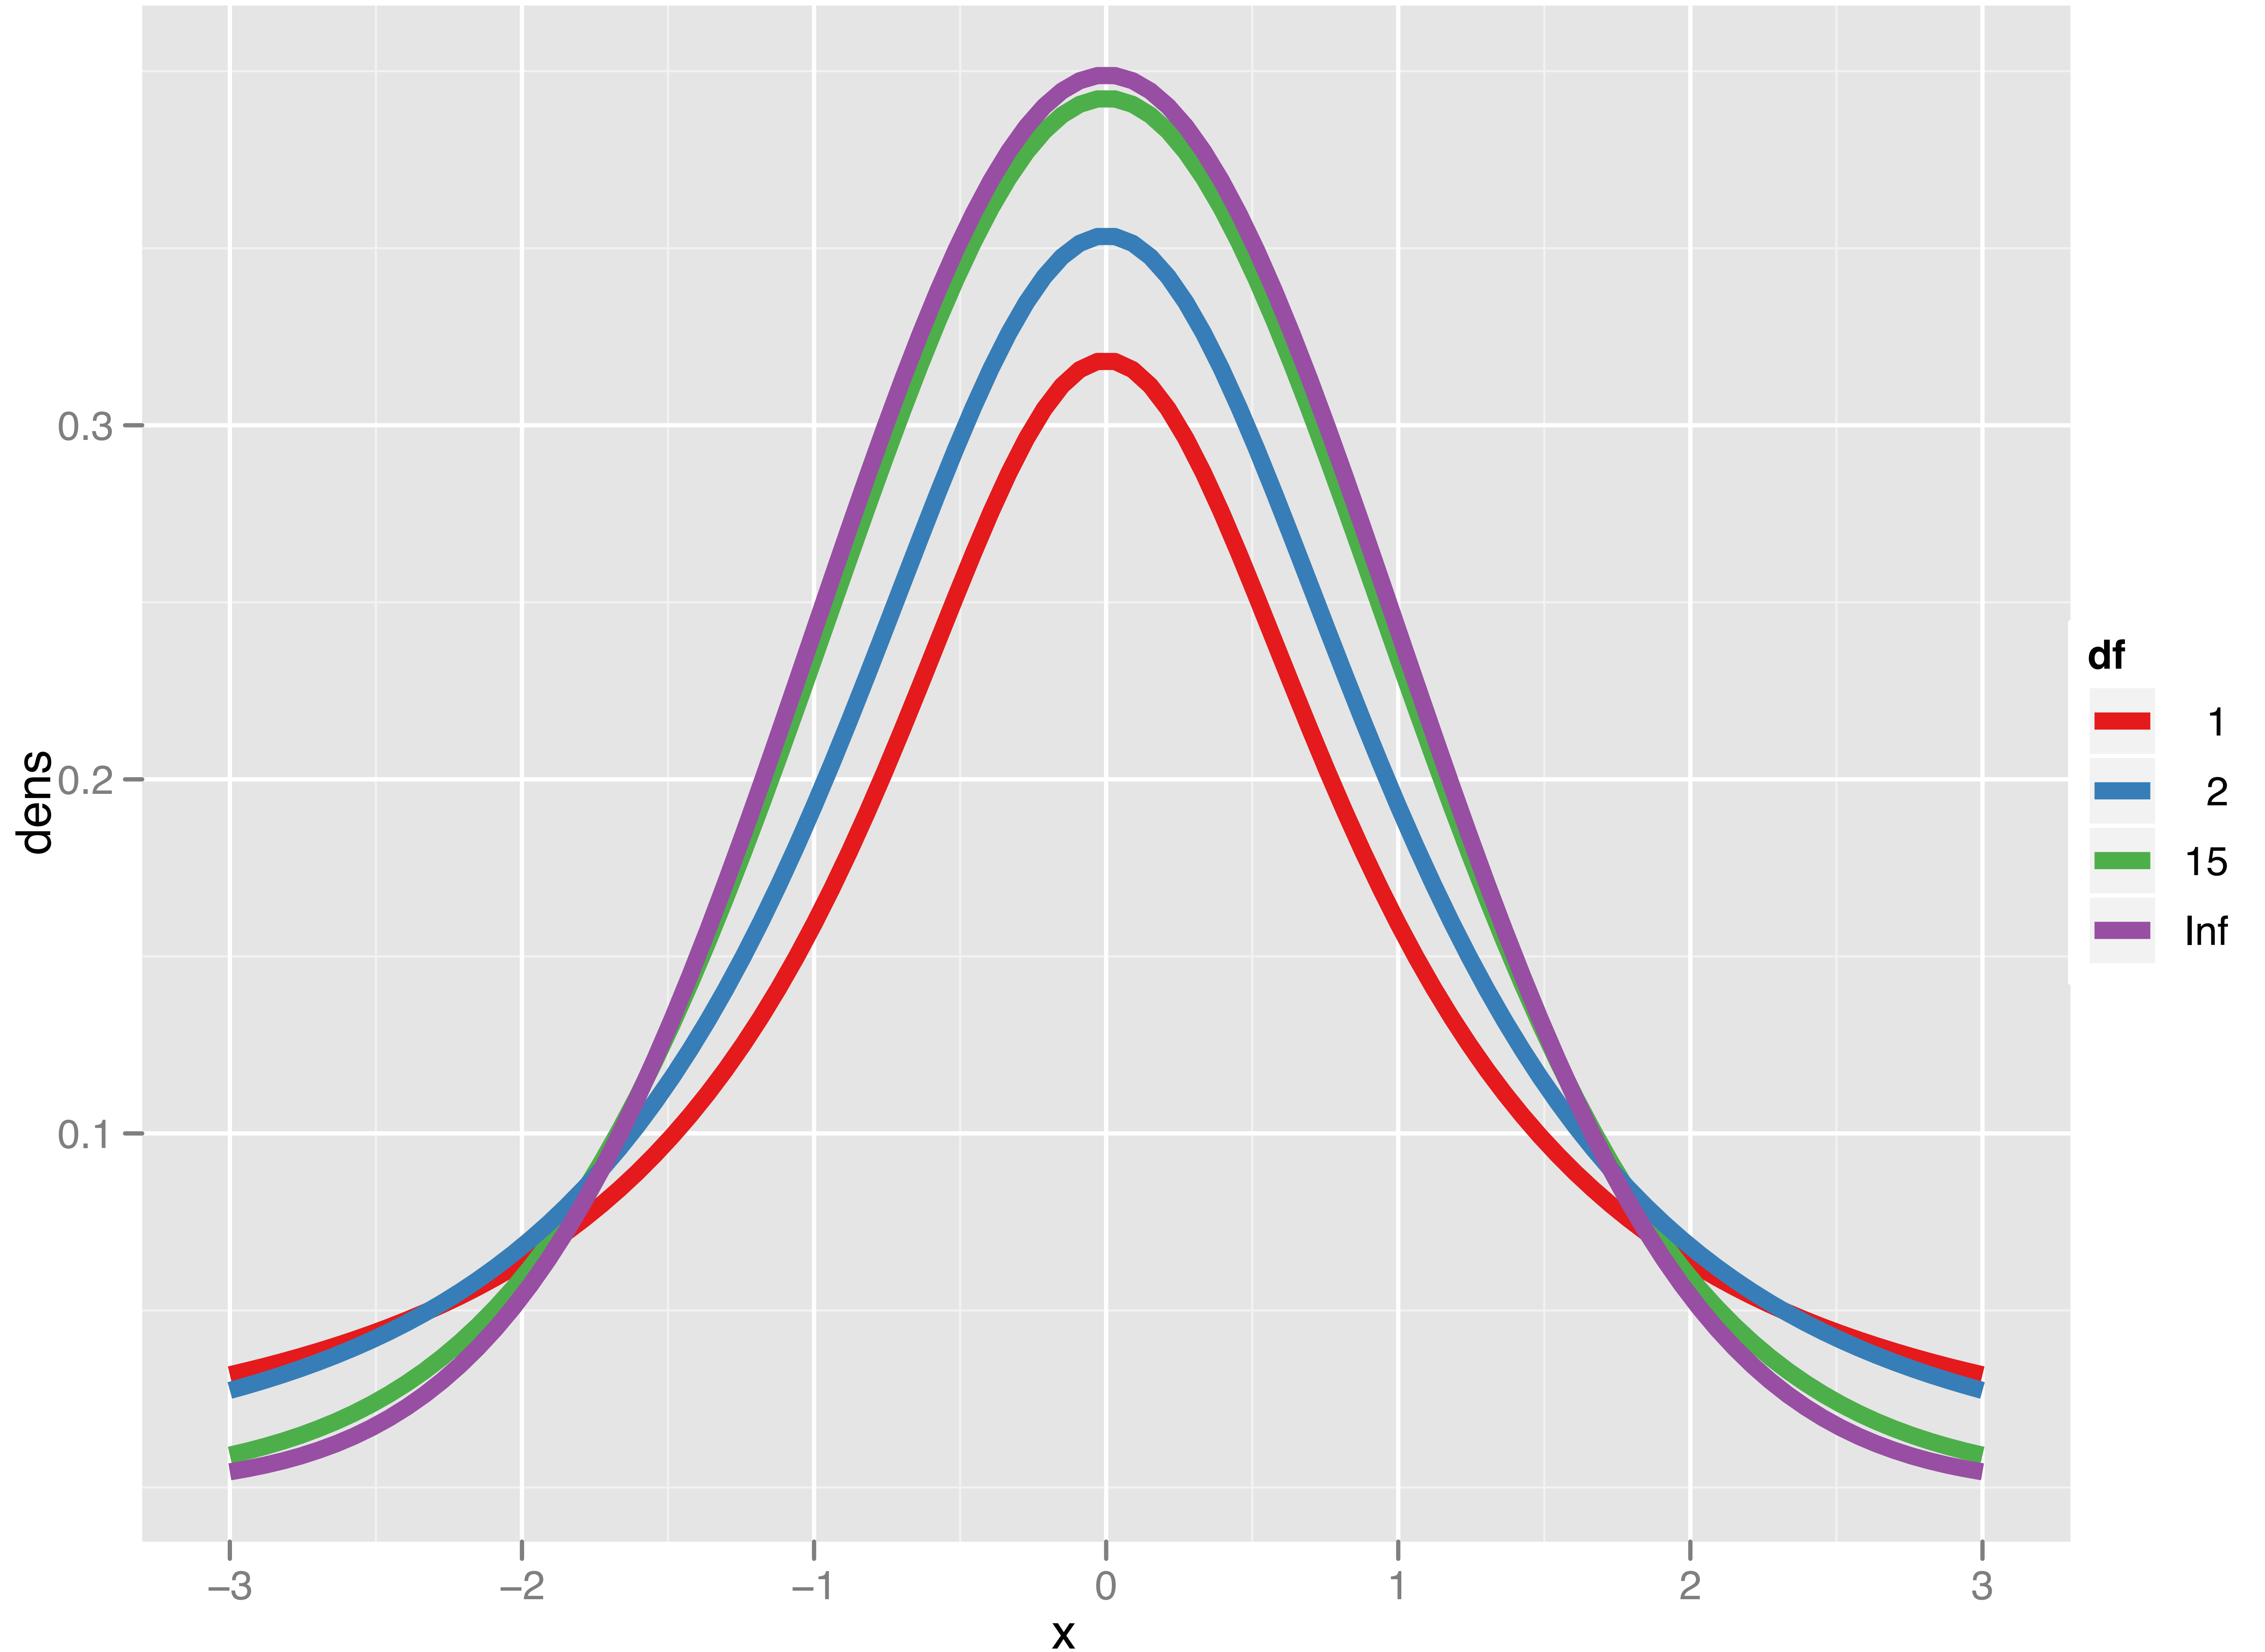
0.3
0.2
0.1
dens
−3
−2
−1
0
1
2
3
x
df
1
2
15
Inf

Properties of the t-dist

Heavier tails compared to the normal distribution.

$$\lim_{n \to \infty} t_n = Z$$

Practically, if n > 30, the t distribution is essentially equivalent to the normal.

```
limit_(nu -> oo) (nu/(x^2+nu))^((1+nu)/2)/(sqrt(nu) beta(nu/2, 1/2))
```

t-tables

Basically the same as the standard normal. But one table for each value of degrees of freedom.

Easiest to use calculator or computer: http://www.stat.tamu.edu/~west/applets/tdemo.html

(For homework, use this applet, for final, I'll give you a small table, if necessary)

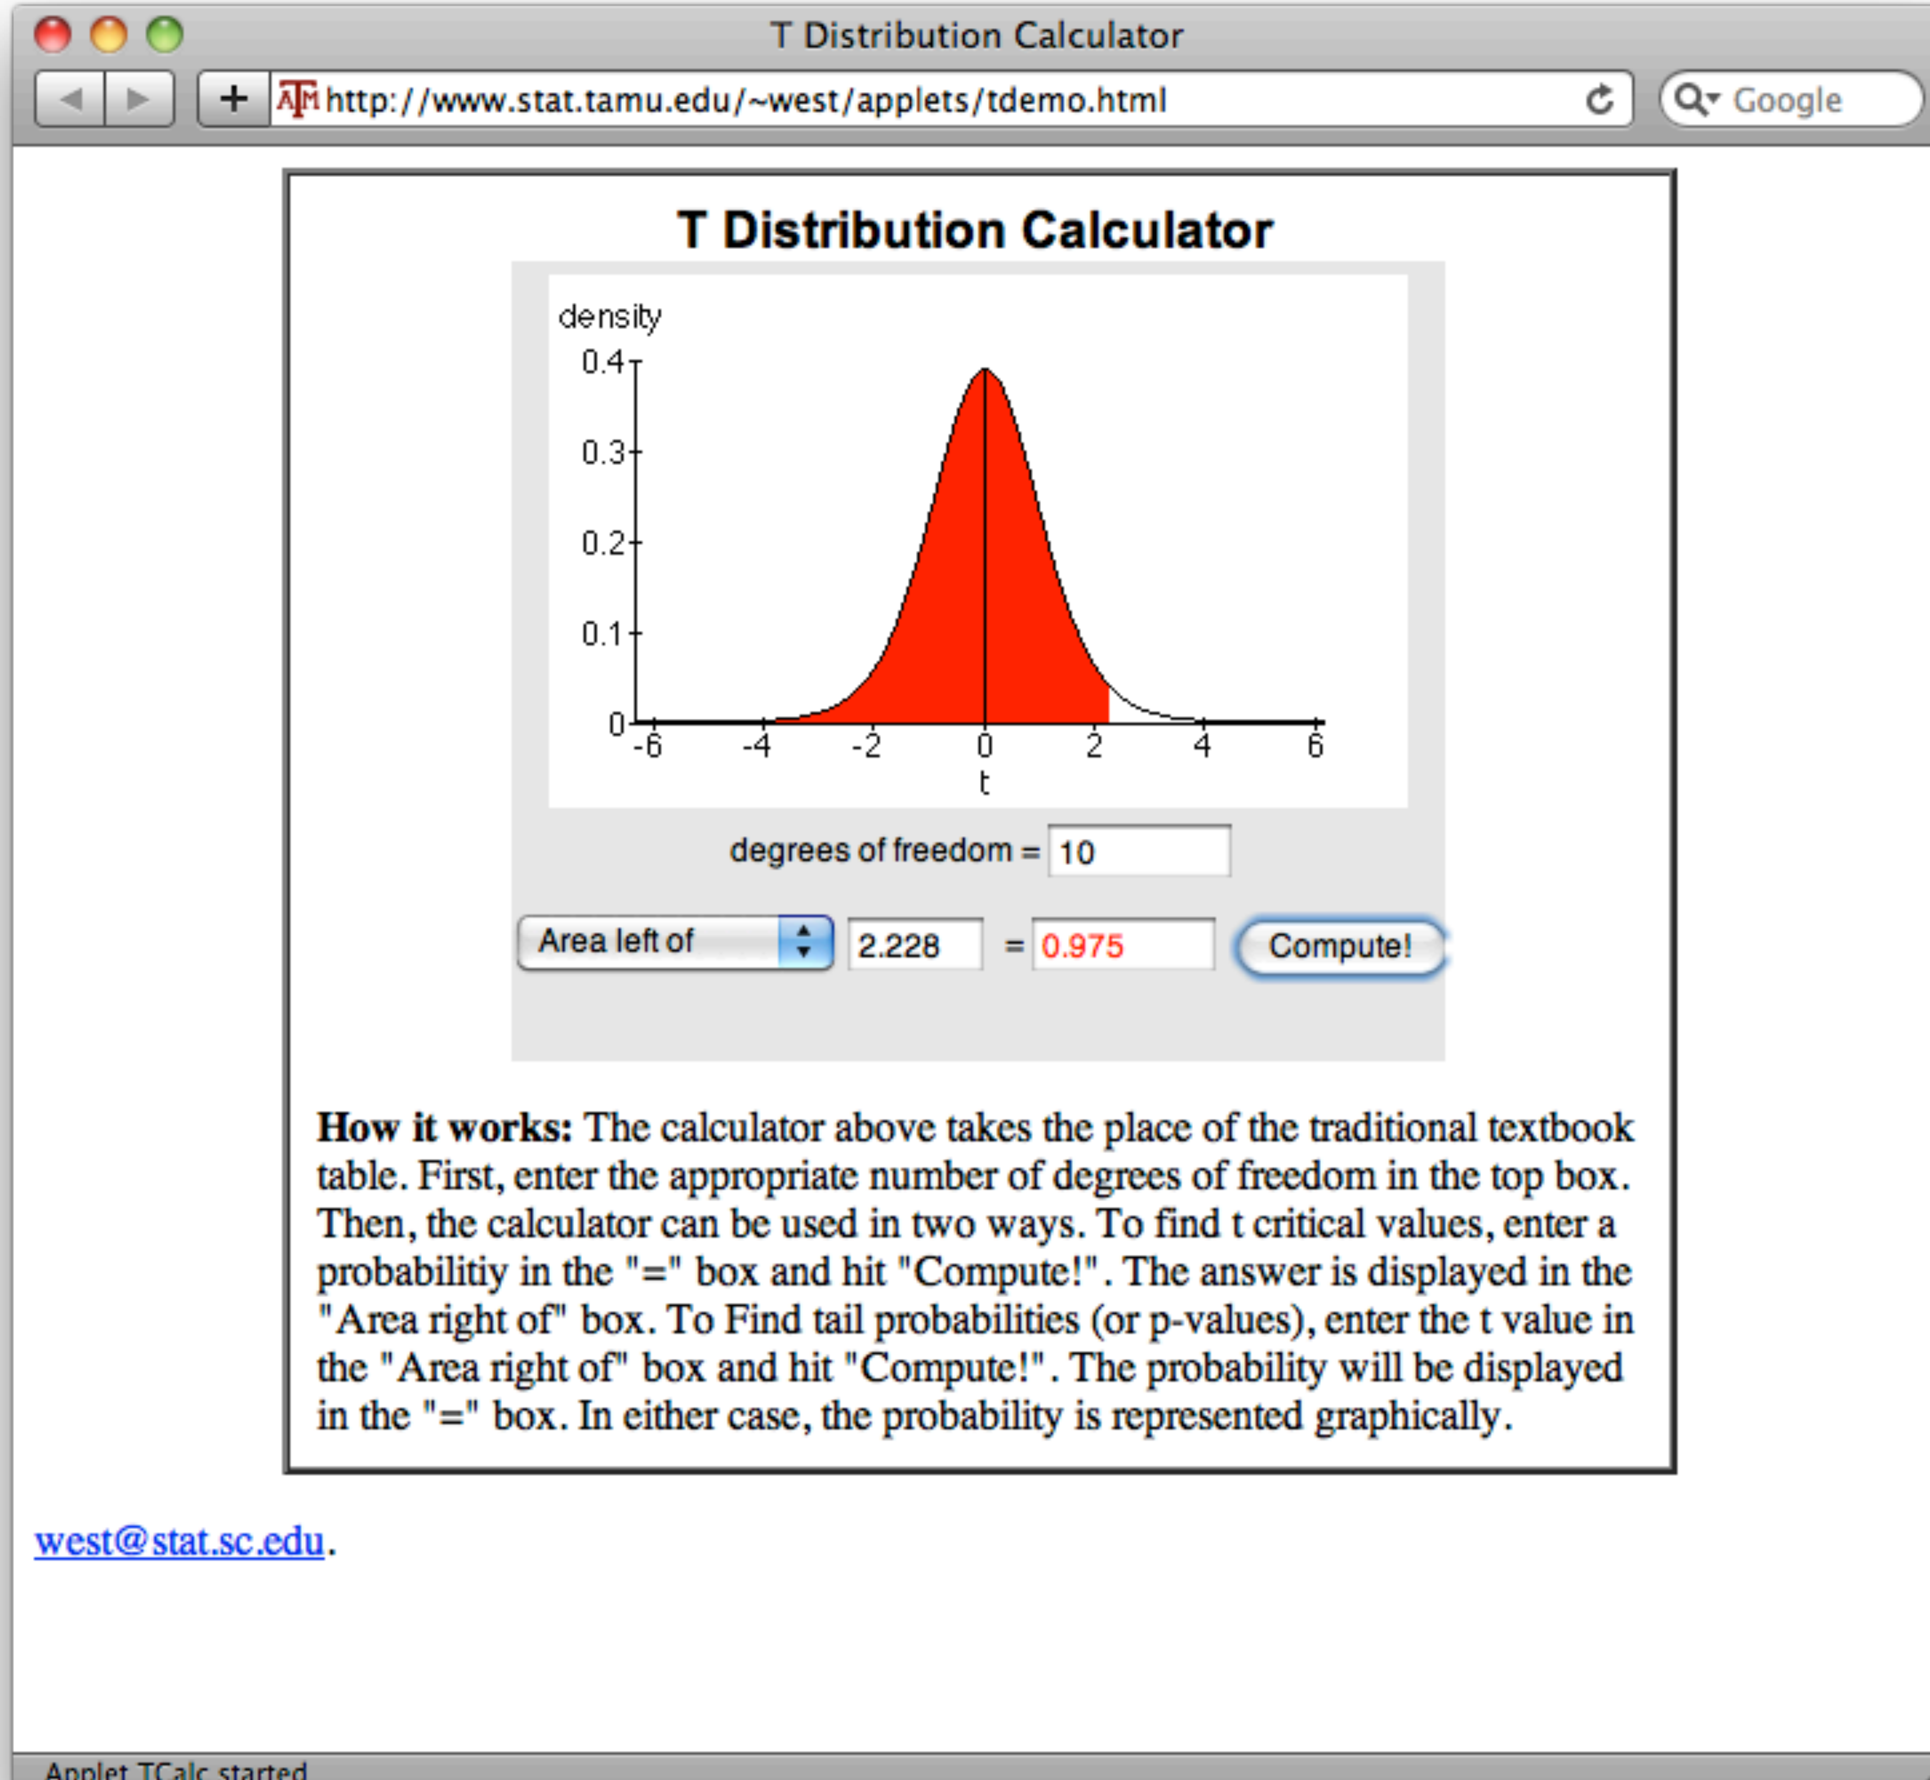

T Distribution Calculator

How it works: The calculator above takes the place of the traditional textbook table. First, enter the appropriate number of degrees of freedom in the top box. Then, the calculator can be used in two ways. To find t critical values, enter a probabilitiy in the "=" box and hit "Compute!". The answer is displayed in the "Area right of" box. To Find tail probabilities (or p-values), enter the t value in the "Area right of" box and hit "Compute!". The probability will be displayed in the "=" box. In either case, the probability is represented graphically.

west@stat.sc.edu.

Your turn

We perform an experiment to measure the speed of sound and repeat it 10 times: 340 333 334 332 333 336 350 348 331 344 (mean: 338, sd: 7.01)

Assuming $X_i \sim \text{Normal}(\mu, \sigma^2)$, what is an estimate of the speed of sound? What is the error (sd) of this estimate? Give an interval that we're 95% certain the true speed of sound lies in.

(Repeat the steps from the handout)

Example

340 333 334 332 333 336 350 348 331 344 (mean: 338, sd: 7.01)

If not known: (333, 342) (2.23)