

Inferential Statistics

Lecture 7: Confidence, prediction and tolerance intervals

Jérémy Hervelin

Pforzheim University
j.hervelin@lb.hs-pforzheim.de

9 November 2026 13:45–15:15
Winter semester 2026/27

Adapted from the teaching materials of Prof. Dr. Thilo Klein.

Our 90-minute route

- First 15 minutes: retrieve prerequisites and introduce the question.
- Next 35 minutes: worked examples, with short checks after each concept.
- Next 25 minutes: paired activity and interpretation.
- Final 15 minutes: synthesis, questions and an applied discussion.

Today: three different kinds of interval

- Estimate a mean when population SD is unknown.
- Estimate a proportion and plan approximate precision.
- Distinguish confidence, prediction and tolerance intervals.
- State the repeated-sampling interpretation for each target.

Replace a known SD by an estimated SD

For iid normal observations with unknown σ ,

$$T = \frac{\bar{X} - \mu}{S/\sqrt{n}} \sim t_{n-1}.$$

The extra uncertainty from estimating σ produces heavier tails than the standard normal distribution.

This result is exact for any $n > 1$ under the normal model. There is no universal sample-size threshold that makes non-normal data harmless.

Classroom check – Q1

Exam-style

1–2 minutes

A t interval based on 8 iid normal observations uses how many degrees of freedom?

Think individually, then compare your reasoning with a neighbour.

Lecture 7, classroom quiz Q1. Same question number as the printed quiz.

Confidence interval for an unknown mean

With observed mean $\bar{x}$, sample SD s and $n > 1$:

$$\bar{x} \pm t_{n-1, 1-\alpha/2} \frac{s}{\sqrt{n}}.$$

For 95% confidence use the 0.975 quantile, because the remaining 5% is split between two tails.

For strongly skewed samples, extreme outliers, dependent observations or heavy tails, investigate whether this method is appropriate.

Classroom check – Q5

Exam-style

1–2 minutes

Holding n and s fixed, increasing confidence from 95% to 99% makes a t interval...

- A. Narrower;
- B. Wider;
- C. Unchanged

Think individually, then compare your reasoning with a neighbour.

Lecture 7, classroom quiz Q5. Same question number as the printed quiz.

Original student-rent example

Eight rents: 270, 460, 512, 550, 360, 399, 419 and 390.

$$\bar{x} = 420, \quad s^2 = 7772.286, \quad n = 8.$$

$$420 \pm t_{7,0.975} \sqrt{7772.286/8} = [346.30, 493.70].$$

The interval targets the population mean rent under an iid normal sampling model. With only eight observations, that model needs substantive support.

One-sided confidence bounds

A 95% lower confidence bound for a normal-population mean is

$$\left[\bar{x} - t_{n-1,0.95}s/\sqrt{n}, \infty \right).$$

An upper bound is defined analogously.

Choose the direction from the question before looking at the results. Changing direction because the observed mean is high or low invalidates the nominal guarantee.

An approximate interval for a proportion

For independent Bernoulli observations, $\hat{p} = s/n$.

$$\hat{p} \pm z_{1-\alpha/2} \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}}.$$

This is the Wald interval: a normal approximation. It performs poorly near 0 or 1 and in small samples; it may even extend beyond $[0, 1]$.

For illustration, 60/100 gives $0.60 \pm 1.96 \sqrt{0.24/100}$, approximately $[0.504, 0.696]$.

Classroom check – Q4

Concept check

1–2 minutes

A Wald proportion interval is most problematic when...

- A. n is large and p is near 0.5;
- B. n is small or p is near 0 or 1;
- C. The units are percentage points

Think individually, then compare your reasoning with a neighbour.

Lecture 7, classroom quiz Q4. Same question number as the printed quiz.

A more stable alternative: the Wilson interval

Let $z = z_{1-\alpha/2}$. The Wilson interval has centre and half-width

$$c = \frac{\hat{p} + z^2/(2n)}{1 + z^2/n}, \quad h = \frac{z}{1 + z^2/n} \sqrt{\frac{\hat{p}(1 - \hat{p})}{n} + \frac{z^2}{4n^2}}.$$

Use $[c - h, c + h]$. It is obtained by inverting a score test.

For 60/100 the 95% Wilson interval is approximately $[0.502, 0.691]$. This formula is enrichment; focus on why the Wald approximation can fail.

Planning precision before collecting data

For desired margin of error e in a proportion estimate:

$$n \approx \frac{z_{1-\alpha/2}^2 p_0(1 - p_0)}{e^2}, \quad \text{round up.}$$

Use a defensible planning value p_0 ; use 0.5 for the largest Bernoulli variance.

This is an approximate precision calculation for independent observations. Non-response, clustering and finite-population corrections require separate treatment. It is not a power calculation.

Classroom check – Q3

Exam-style

1–2 minutes

A target interval has total width 0.02. What is the margin of error?

Think individually, then compare your reasoning with a neighbour.

Lecture 7, classroom quiz Q3. Same question number as the printed quiz.

Original precision example: width versus half-width

A 99% interval with total width at most 0.02 has margin of error $e = 0.01$.

$$n = \left\lceil \frac{z_{0.995}^2 (0.25)}{0.01^2} \right\rceil = 16588.$$

Using the rounded critical value 2.576 gives 16590 instead.

Activity: what happens if 0.02 is mistakenly used as the half-width? The planned sample becomes about four times too small.

Predict a new observation

A mean interval estimates a population average. A prediction interval targets a single independent future observation from the same normal population:

$$\bar{x} \pm t_{n-1, 1-\alpha/2} s \sqrt{1 + 1/n}.$$

The extra 1 represents variation of the future individual around the population mean. Prediction therefore remains uncertain even with a very large estimation sample.

Classroom check – Q2

Exam-style

1–2 minutes

A 95% interval for one future commuting cost is a . . .

- A. Confidence interval for the mean;
- B. Prediction interval;
- C. Tolerance interval

Think individually, then compare your reasoning with a neighbour.

Lecture 7, classroom quiz Q2. Same question number as the printed quiz.

Original commuting-cost prediction

Using the original rounded summaries $n = 10$, $\bar{x} = 2403$, $s = 127.6$:

$$95\% \text{ prediction interval} = [2100.26, 2705.74].$$

A one-sided 95% upper prediction bound is 2648.32.

The coverage statement repeats both the original sample and a new independent cost. This is not a confidence interval for mean commuting cost.

Prediction without a normal model

For n iid observations from a continuous distribution, let $X_{(i)}$ denote the i th ordered observation. For an independent future Y :

$$P\{X_{(i)} < Y < X_{(j)}\} = \frac{j-i}{n+1}.$$

Also $P\{Y < X_{(j)}\} = j/(n+1)$ and $P\{Y > X_{(i)}\} = (n+1-i)/(n+1)$.

The new observation has any of the $n+1$ ranks with equal probability. Ties require separate treatment.

Tolerance intervals: population content

A tolerance interval aims to cover at least a chosen fraction γ of the population with confidence $1 - \alpha$.

For a normal population it often takes the form $\bar{x} \pm ks$. The factor k depends on n , γ , confidence and the chosen method.

Example question: with 95% confidence, contain at least 90% of individual commuting costs. This is a different guarantee from predicting one future cost.

Choose the target before choosing the formula

Question	Tool
What is mean earnings?	Confidence interval
What might one new person earn?	Prediction interval
Contain at least 90% of earnings?	Tolerance interval

Pair activity: formulate one question of each type for delivery times. State the assumptions and explain the uncertainty in words.

Independent practice

- Practise the five online questions for this lecture; 80% provides access to the next practice module.
- Course materials remain accessible regardless of your practice progress.
- Explain your reasoning before comparing answers.
- Review the method, assumptions and interpretation as well as the calculation.

Sources and teaching notes

Adapted from Prof. Dr. Thilo Klein's supplied Inferential Statistics materials.

Revised teaching edition: Jérémy Hervein, Pforzheim University, winter semester 2026/27.

New youth-employment, programme and interview-score examples are synthetic illustrations.