

Inference for Quantitative Data: Means

AP Statistics

Should I Worry About Error?

Inference for a **mean** works like inference for a proportion, with one change: we rarely know the population standard deviation σ , so we estimate it with the sample s . That extra uncertainty means we use the ***t*-distribution** instead of the normal – a **distribution** 分布 that is bell-shaped but with heavier tails, and it depends on the **degrees of freedom** 自由度 $df = n - 1$; as n grows it approaches the normal.

Confidence Interval for a Mean


A one-sample t interval for μ :

$$\bar{x} \pm t^* \frac{s}{\sqrt{n}}.$$


t^* is the critical value with $df = n - 1$. Conditions: **random** sample, **Normal/Large Sample** (population normal, or $n \geq 30$ by the CLT, or a roughly symmetric sample with no outliers), and the **10%** condition. Interpret the interval and the confidence level in context.

Worked example. A random sample of $n = 25$ has $\bar{x} = 50$ and $s = 8$. For a 95% interval, $df = 24$ gives $t^* = 2.064$:

$$50 \pm 2.064 \cdot \frac{8}{\sqrt{25}} = 50 \pm 2.064(1.6) = 50 \pm 3.3 = (46.7, 53.3).$$



The t -distribution has a lower peak and heavier tails than the normal



"95% confident" describes the method, not one interval: over many samples about 95% of the intervals contain μ and about 5% miss it.

Justifying a Claim About a Mean

As with proportions: a claimed mean **inside** the interval is plausible; **outside** the interval, the data give evidence against it. Answer in context using the plausible range.

Setting Up a Test for a Mean

State hypotheses about μ : $H_0 : \mu = \mu_0$ versus $H_a : \mu \neq \mu_0$ (or $<$, $>$). Check the same conditions. The one-sample t **statistic**:

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}, \quad df = n - 1.$$

Worked example. Test $H_0 : \mu = 45$ against $H_a : \mu \neq 45$ for the sample above ($\bar{x} = 50$, $s = 8$, $n = 25$):

$$t = \frac{50 - 45}{8/\sqrt{25}} = \frac{5}{1.6} = 3.13, \quad df = 24.$$

This t is far out in the tail (two-tailed $p < 0.01$), so **reject** H_0 – strong evidence the mean is not 45. Notice 45 also falls outside the 95% interval (46.7, 53.3), the same conclusion by two routes.

Carrying Out a Test for a Mean

Find the p -value from the t -distribution with $df = n - 1$, compare to α , and conclude in context – reject or fail to reject H_0 , then state what that means for the claim. Show the test name, statistic, df , and p -value.

Confidence Interval for a Difference of Two Means

For **independent** samples, estimate $\mu_1 - \mu_2$:

$$(\bar{x}_1 - \bar{x}_2) \pm t^* \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}.$$

Conditions must hold in both samples. (Use technology for the df ; do not pool the variances on the AP exam.)



Randomisation underpins fair comparison of two groups in a mean difference test

Justifying a Claim About Two Means

If the interval for $\mu_1 - \mu_2$ contains 0, the data are consistent with equal means; if it excludes 0, there is evidence of a difference in that direction. Interpret in context.

Setting Up a Test for a Difference of Means

Hypotheses: $H_0 : \mu_1 = \mu_2$ versus $H_a : \mu_1 \neq \mu_2$ (or $<$, $>$). Distinguish **two independent samples** from **paired data** 配对数据 – for paired data (before/after, matched subjects), first take the **differences** and run a **one-sample t** procedure on them.

Carrying Out a Test for a Difference of Means

The two-sample t statistic:

$$t = \frac{(\bar{x}_1 - \bar{x}_2) - 0}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}.$$

Get the p -value (technology for df), compare to α , conclude in context.

Selecting and Communicating a Procedure

The hardest exam skill is **choosing** the right procedure: one or two samples? proportion or mean? paired or independent? confidence interval or test? Read the question for what is being estimated or claimed, then name the procedure, check its conditions, carry it out, and communicate the conclusion clearly with numbers and context.

Exam tips

- Use **t-procedures** for means (population σ unknown) — the t-distribution has heavier tails than normal.
- Check conditions: random, independent, and roughly normal (or large n).
- Interpret an interval and a test in context, always tied to the parameter (the true mean).
- Match the right procedure: one-sample, two-sample, or paired (look for a natural pairing).
- State the degrees of freedom; for a two-sample t -test use technology's value (or, by hand, the conservative smaller $n - 1$).