One-sample z-test for proportions

Use case

We have a single binary categorical variable. We want to perform inference on the true proportion of the population,
p, that fall into one value (e.g. "yes" on a survey). We cannot know this value with certainty, so we estimate by drawing a sample of size
n. We record
X successes from our sample. Let
$\hat{p} = \frac{X}{n}$, our sample proportion.
$\hat{p}$ is an *unbiased estimator* of
p.

Preconditions

Preconditions for hypothesis testing

This technique relies upon the normal approximation to the binomial distribution. For this to hold, the following conditions must hold:

1. $np_0 \geq 10$
2. $n(1 - p_0) \geq 10$

Where
p_0 is our supposed value of
p under the null hypothesis.

Preconditions for confidence interval

If we are constructing a confidence interval for
p, the preconditions are slightly different. Here, the normal approximation to the binomial distribution still needs to hold, but it applies to our sample proportion,
$\hat{p}$ rather than
p_0. The following conditions must hold:

1. $n\hat{p} \geq 10$
2. $n(1 - \hat{p}) \geq 10$

Note that this is equivalent to having at least 10 observations corresponding to either value (e.g., at least 10 "yes" and at least 10 "no").

Hypotheses

Null hypothesis,
H_0**:**
$p = p_0$

Alternative hypothesis,
H_a**:**
$p \neq / > / < p_0$

Test statistic

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$

p-value

Where Z is a standard normal random variable,

1. For
   $H_a : p < p_0$:
   $p = P(Z \leq z)$
2. For
   $H_a : p > p_0$:
   $p = P(Z \geq z)$
3. For
   $H_a : p \neq p_0$:
   $p = 2P(Z \leq -|z|)$

p-value (Python)

1. For
   $H_a : p < p_0$: `p = stats.norm.cdf(z)`
2. For
   $H_a : p > p_0$: `p = 1 - stats.norm.cdf(z)`
3. For
   $H_a : p \neq p_0$: `p = 2 * stats.norm.cdf(-abs(z))`

Associated confidence interval

$$\text{C\% confidence interval} = \hat{p} \pm z^{\star}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$
$$\text{choose } z^{\star} \text{ s.t. area on standard normal distribution from } (-z^{\star}, z^{\star}) = \text{C}$$

See the preconditions for the confidence interval.

Minimum sample size to achieve margin of error:

$$n = (\frac{z^{\star}}{m^{\star}})^2 p^{\star}(1-p^{\star})$$

$p^{\star}$ is an "educated guess" about the value of
p. If you have access to a sample proportion,
$\hat{p}$, set
$p^{\star} = \hat{p}$. Otherwise, a *conservative approach* is to set
$p^{\star} = 0.5$.

Related tests

If you want to compare proportions across two independent populations, a two-sample z-test for proportions is appropriate.

Considerations

A one-sample z-test for proportions conducted at confidence level
α will reject the null hypothesis if and only if the value corresponding to the null hypothesis,
p_0, is completely outside of the
$C = 1 - \alpha$ confidence interval for the true proportion.

Neither the z-test nor the confidence interval use the standard deviation of the sample proportion,
$\sigma_{\hat{p}}$.

1. The hypothesis test uses the standard deviation of the sample proportion under the null hypothesis,
   $\sigma_{\hat{p}}^{H_0}$.
2. The confidence interval uses the standard error of the sample proportion,

$SE_{\hat{p}}$.