STATISTICAL SUPPORT
SWISS CENTER FOR AFFECTIVE SCIENCES

Regarding *t*, *F*, *z* and χ^2

E-mail distributed on 29-09-2025

Dear all,

Today a brief refresher on the relationship between common test statistics. Basically one should keep in mind that **(a)** *t* is to *z* as *F* is to χ^2, and that **(b)** *t* is to *F* as *z* is to χ^2.

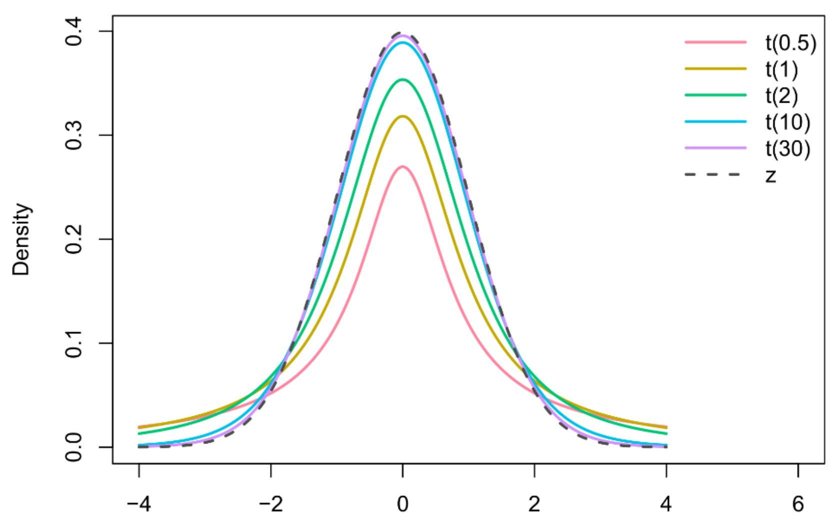

Figure 1. *Densities of the t-distribution compared to the standard normal distribution.*

1. From *t*/*F* to *z*/χ^2

A *t*-distribution is an approximation of the standard-normal distribution—or *z*-distribution—that incorporates the uncertainty added by the estimation of the unknown error variance. It is famously more "heavy-tailed" than the *z*-distribution. For higher degrees of freedom (DF), it will increasingly resemble the *z*-distribution, as shown on Figure 1. The approximation is remarkably fast, in that for DFs as low as 30, the two densities are virtually indistinguishable. Asymptotically, the two distributions become identical at infinite DF. Therefore, one could say that $t(\infty) = z$, just as $F(1,\infty) = \chi^2(1)$. In the graph, the densities were sampled with the `dnorm` and `dt` functions, respectively. It illustrates that these functions also accept fractional (or non-integer degrees of freedom). While values less than 1 would not commonly occur, fractional DFs are a familiar feature in multilevel models and heteroscedastic tests (e.g., Welch *t*-test).

In R, packages like `emmeans` sometimes print infinite DFs (see Figure 2). This occurs when the actual DF is deemed to be either too large or complicated to be calculated in a reasonable amount of time.[1] In this case, one essentially gets the corresponding *z*- or χ^2-test. When reporting these values, I would advise to report them as such, rather than write $t(\infty)$.

```
Note: D.f. calculations have been disabled because the number of observations exceeds 3000.
To enable adjustments, add the argument 'pbkrtest.limit = 50000' (or larger)
[or, globally, 'set emm_options(pbkrtest.limit = 50000)' or larger];
but be warned that this may result in large computation time and memory use.
Note: D.f. calculations have been disabled because the number of observations exceeds 3000.
To enable adjustments, add the argument 'lmerTest.limit = 50000' (or larger)
[or, globally, 'set emm_options(lmerTest.limit = 50000)' or larger];
but be warned that this may result in large computation time and memory use.

 X    emmean      SE  df asymp.LCL asymp.UCL z.ratio p.value
 0 -0.00667 0.00627 Inf   -0.0190   0.00562  -1.064  0.2875
 1 -0.00343 0.00630 Inf   -0.0158   0.00891  -0.545  0.5857

Degrees-of-freedom method: asymptotic
Confidence level used: 0.95
```

Figure 2. Infinite degrees of freedom printed by R package `emmeans`*.*

2. From t/z to F/χ^2

In linear regression, any *t*-test can be converted to its corresponding *F*-test simply by squaring the *t*-value, with numerator DF equal to 1, and denominator DF equal to the DF of the *t*-test. In R, one can verify this equivalence easily by comparing the `summary` output of a model to its ANOVA breakdown (e.g., `car::Anova`). While every *t*-test has an *F*-test equivalent, the reverse is not true, for example *F*-tests where the numerator DF is larger than 1.

The preceding equivalence remains true in generalized linear models (GLM) with non-Gaussian outcomes. Any *z*-test can be converted to a corresponding χ^2-test by squaring its value. However, when comparing the summary and ANOVA outputs in R, you will find that these values do differ. To obtain the correct equivalent, one should use:

```
Anova(model, test="Wald")
```

What is happening here? Many researchers are not aware that there are, in fact, three different ways to construct a test statistic from a model's likelihood function, that is, **(1)** the Wald test, **(2)** the likelihood ratio test (LRT), and **(3)** the score test. Of these, the Wald test is the most well-known, with the test statistic constructed by dividing the parameter of interest by its standard error (e.g., a regression slope divided by its SE). In models with Gaussian outcomes, the *t*- and *F*-tests are Wald tests, and usually the only ones implemented in software. For non-Gaussian outcomes, all three types have applications,

[1] Typically in multilevel models fitted to large amounts of data

although the score test is the rarest. For some procedures, the LRT is the default, such as ANOVA breakdowns for GLMs[2] (e.g., as implemented in `car::Anova`). However, regardless of which type of test is used, all three test statistics will follow either a *t*/*F* distribution or z/χ^2 distribution. On the surface, one would therefore not notice the difference, although most software will explicitly print the type somewhere in the output.

It is sometimes mistakenly believed that χ^2-tests are exclusive to categorical inference but the preceding should make clear that this is not true. In R, one can obtain *z*-tests in the summary output for Gaussian models by fitting a GLM and then applying the `summary` function with a specified dispersion parameter, e.g.:

```
model <- glm(Y~X)
summary(model, dispersion=1)
```

Which will return a summary of *z*-tests.[3] Here, the value of the dispersion parameter was taken from an earlier iteration of the `summary` function, which will estimate a dispersion value of none is given. By default `car::Anova` will print LRT χ^2-tests for `glm` objects, regardless of the outcome distribution, but it can be overridden by specifying `Anova(model,test="F")`. Note that the above is primarily intended for illustration. In practice, you should rarely need this output for a Gaussian model, since it is based on the assumption of known error variance.

--
Ben Meuleman, Ph.D.
Statistician
Swiss Center for Affective Sciences
University of Geneva | Campus Biotech
Chemin des Mines 9 | CH-1202 Genève
ben.meuleman@unige.ch | +41 (0)22 379 09 79

[2] Technically in GLMs such breakdowns should be referred to as an "analysis of deviance"
[3] I may misremember this but I believe at some point in R's history *z*-tests were the default output for all `glm` objects