

Simple power calculation for intervention studies

by Christian Gold, September 2012 (links updated in 2015)

Many beginning researchers (PhD students and others) have told me that my article "The use of effect sizes in music therapy research", published in *Music Therapy Perspectives* in 2004, provided an accessible overview of a difficult but important topic. It is indeed essential when planning an outcome study to understand the basics of power calculation. The article showed the relationships between effect size, sample size, and test power. What was missing in my article, however, was an accessible instruction on how to actually *do* a power calculation. I aim to provide that here and hope readers will find it useful.

This guide assumes the simplest of situations: a comparison between two groups using a t-test. It also assumes that readers have already thought about what level of effect size (i.e., large, $d = 0.80$, medium, $d = 0.50$, small, $d = 0.20$) is realistic to expect in their study. I will show how to find the power for a given sample size, and how to find the sample size for 80% power, with a two-sided 5% significance level.

You either need to have the statistics programme R installed on your computer (available freely from <http://www.r-project.org>), or alternatively, use an online version¹ of R if you do not want to install a programme. You do not need to know anything about the programme, but you need to be able to type in the commands below exactly as shown.

To find the required sample size for a study where you expect a large effect, type:

```
power.t.test(power=0.8, d=0.8)
```

You will see that you need to have 26 subjects (rounded upwards from 25.52). Note that this is the number needed in *each* group, so in total you need 52 participants. Note also that many studies have about 10% drop-outs. To take that into account, you should aim to include a few more, maybe 55 or 60.

To find the power of a given sample size, again for a large effect, type:

```
power.t.test(n=10, d=0.8)
```

Here, the output shows that test power is 39% (0.3949). You probably do not want to start a study where the chance of "succeeding" is that lower than the chance of failing. Again, the sample size is per group, and drop-outs are not considered.

Trying this out with different parameters will help you find the optimal sample size, or more often, the best compromise between what would be desirable and what is feasible. It is also a great way of learning how test power works. Remember that of the three parameters, effect size, sample size, and test power, you need to specify two and the third will be calculated.

You can also extend this to calculate power for one-sided tests,

```
power.t.test(power=0.8, d=0.8, alternative="one.sided")
```

significance levels other than 5%,

```
power.t.test(power=0.8, d=0.8, sig.level=0.01)
```

and paired t-tests,

```
power.t.test(power=0.8, d=0.8, type="paired")
```

but you should only do it if you understand what these options mean and where you can apply them. My experience is that as a practical tool for planning research, simple power calculations tend to work best. The remaining considerations are non-statistical - you may refer back to my article cited above to learn about what you can do if the power of your study is not what you hope it might be.

¹ These may change over time. Search the internet for "r online" to find one. In June 2015, www.r-fiddle.org was one that worked; you can simply type the commands into the lower box on that page.