Editorial

James V. Felicetta, MD

Editor-in-Chief

The Number Needed to Treat: You Can Do the Math!

This month, I want to talk a little bit about medical statistics. Wait! Before you groan and turn the page, give it a chance. I promise to make it as painless as humanly possible.

Although many of us feel fairly uncomfortable grappling with statistics, they are hugely important to the practice of medicine today. We're constantly bombarded by journal articles telling us that we should adopt such-and-such a medical practice and shun another. In view of the extensive, inherently biased role that the pharmaceutical industry often plays in both the design and reporting of clinical trials, it is critical that clinicians be able to separate the wheat from the chaff when deciding whether to adopt new practices. And statistics can be a big help in this task.

One statistic that is extremely useful in evaluating medical therapies is the number needed to treat (NNT). As the name suggests, this is the total number of patients who would need to be treated with a given therapy in order for a single patient to derive a defined benefit—typically, the avoidance of negative events, such as myocardial infarction or stroke.

We don't like to think of it this way, but the reality is we usually have to give a sizable number of patients a certain treatment for just one of them to have a positive benefit. How many is acceptable? A general rule of thumb is that a treatment associated with an NNT of 30 or fewer probably should be given very serious consideration. If the NNT is greater than 100, on the other hand, we probably want to forgo the therapy, unless it is both remarkably inexpensive (costing just pennies per patient) and essentially free of adverse effects. Between 30 and 100 is somewhat of a gray zone, in which the factors of cost, adverse effects, and general clinical experience hold sway.

Given the practicality of the NNT, you might think that every clinical trial report would include this statistic. Unfortunately, this is hardly the case. Could it be that some authors don't want to call attention to the fact that the therapeutic benefits they're reporting actually require a high NNT? Whatever the motive may be for leaving this statistic out, it turns out to be quite easy to calculate the NNT ourselves.

We start by looking at the absolute event rates for the two populations being compared. Let's say that traditional treatment X is associated with a 12% rate of a predefined adverse outcome, while promising new treatment Y has only a 10% rate over the same study period. The investigators may enthusiastically report that treatment Y reduced the relative risk of this adverse outcome by 16.7%, which they derived by dividing the difference between the two rates (2%) by the rate for the traditional treatment (12%).

But before we get too excited, let's calculate the NNT. To do this, we simply divide 100 by that 2% difference (dropping the percentage sign)—which gives us an NNT of 50 (100/2 = 50). An absolute benefit of 2% and an NNT of 50 put a serious damper on that 16.7% relative risk reduction, and they call into question the wisdom of abandoning treatment X in favor of treatment Y. The fact is that, when the absolute event rates are fairly low, it's easy to produce a relative risk reduction that looks good but has limited applicability once it's put into the context of the NNT.

Let's try one more example. Suppose that, over a defined study period, the proposed new treatment A results in a 20% absolute rate of adverse outcomes, while traditional treatment B results in a 25% rate. Based on these numbers, the authors report a relative risk reduction of 20% for the new treatment—not too far off from our first example.

But when we calculate the NNT, dividing 100 by the absolute difference of 5, we get an NNT of 20. This means that one need only treat 20 patients with the newer therapy as opposed to the older therapy to prevent an event. If the new treatment is not too expensive or too toxic in terms of adverse effects, an NNT of 20 could strongly incline us in its favor. Thus, once again, the NNT gives us a much more useful perspective on the potential worth of a new approach than does the relative risk reduction.

So, there you have it. You just look for the absolute event rates in the two groups, subtract one from the other, and then divide 100 by this difference. With this simple tool, we can all be a little more savvy when it comes to interpreting clinical trial results. ●

Author disclosures

The author reports no actual or potential conflicts of interest with regard to this editorial.