A study of medication for knee osteoarthritis points the way to new methods for ranking drugs' effectiveness

February 24 2015, by Jacqueline Mitchell

“Using high-quality, large studies helped us to come up with very precise estimates, which allowed us to make conclusions with a lot of confidence,” says Raveendhara Bannuru. Credit: Depositphotos

Maybe you "trust Tylenol" or (like this writer) you're "all Advil." Research proves that both painkillers work, but many of us, including our doctors, can't help but have a preference shaped by experience and

perhaps even advertising. Which really does work better? That's what Raveendhara Bannuru, director of the Center for Treatment Comparison and Integrative Analysis at Tufts Medical Center and a research assistant professor of medicine, wanted to find out.

Bannuru and his colleagues compared the effectiveness of various treatments for pain caused by knee osteoarthritis, one of the most common complaints among older people. Using data from 137 studies, the researchers compared the relative efficacy of five oral pain pills, including acetaminophen and ibuprofen (the generic versions of Tylenol and Advil), and two injectable drugs, and oral and injectable placebos.

Their results, published in last month's *Annals of Internal Medicine*, were somewhat surprising. Every treatment worked better than acetaminophen (Tylenol) with one exception: celecoxib, an expensive, newer drug once hailed as a miracle treatment for joint pain. Overall, the injectable therapies outperformed the oral pain medications, a finding that runs contrary to the conventional wisdom. What's more, the placebo injectables—that is, a simple shot of saline solution—provided patients pain relief comparable to any oral pain medicine tested.

Tufts Now asked Bannuru, who just completed his Ph.D. in clinical and translational sciences at the Sackler School of Graduate Biomedical Sciences, to tell us more about his comparative effectiveness research.

Tufts Now: How is comparative effectiveness research different from other scientific research?

Ravi Bannuru: In clinical trials, you randomize people for one treatment—say Tylenol—versus no treatment or placebo. Then you go and find an effect. Those are called efficacy studies. From efficacy studies we know, for example, that Tylenol is better than taking nothing

or a placebo/dummy pill. And we also know that taking Aleve is better than taking nothing. But we don't know which of those two works better. With osteoarthritis alone, there are 20 to 30 treatments, and we don't always know which one to choose.

To answer that question, you need to do comparative effectiveness studies—that means you need to compare active treatments against each other.

How exactly do you compare treatments?

We look at existing research, including randomized clinical trials and observational studies, and we're able to estimate the differences between two drugs or therapies that were never compared before in a direct trial. We usually do meta-analyses and new kinds of analytical techniques like mixed treatment or network meta-analysis—when drug A is compared to drug B, and drug A is compared to drug C, we can estimate the difference between drugs B and C, even though they were never compared directly to one another. With this type of analysis, we can compare many treatments, and we can even come up with a ranking—for example, drug A is better than drug B, which is better than drug C.

Publication bias—the problem of some researchers submitting only positive results to publications—is an increasing concern among scientists. Since your work depends on others' research, does publication bias impact your results, too?

We try to address it, but we can't completely adjust for it. In this project, we went after and found many studies that were never published. I think that's the best way to address publication bias. We tried lots of tricks. We searched the FDA database for clinical trials that were submitted as part of the drug approval process, but were never published in scientific

journals. In recent years, the website clinicaltrials.gov is encouraging authors to publish their study results. Our recent study has a very high number of unpublished studies—at least 15 studies.

How many studies do you need to compare to get good results?

Some people will do a meta-analysis with two studies—which is ridiculous. I could just read both studies. Usually, people say you need at least five for a simple meta-analysis. For a good analysis, 10 would be ideal. But there is no strict rule. Also, fewer large, high-quality studies are better than many small, low-quality studies. The more, the merrier—we used 137. Using high-quality, large studies helped us to come up with very precise estimates, which allowed us to make conclusions with a lot of confidence.

Is it sometimes hard to convince practitioners about your results, say if they fly in the face of conventional wisdom?

I would say it depends. I led a project to develop clinical practice guidelines [PDF] for the Osteoarthritis Research Society International that were released last March. Those guidelines are well accepted. They conditionally recommended the injection therapies.

In the new study that just got published, we are saying that injection therapy is better than pain pills. That's really a surprising finding. Other studies compare oral pills with dummy pills and injection therapies with dummy injections. Comparing these two kinds of studies, people used to say oral treatments are better than injections.

Now we are saying that's not right. To assess two different treatments,

they need to be compared against a common treatment. For example, if A is better than B and C is better than D, we can in no way say that A is better than C. Instead, if A is better than B and B is better than C, then we can say that A is better than C by way of comparing both to B.

The study found that hyaluronic acid injections, used to replace natural fluids, ease pain more than many physicians thought. Were there other surprising results?

We found the placebo effects are not quite equal. We found injection placebos have more effect than oral placebos. We really want to explore that. As physicians, we can't separate the placebo effect from the treatment effect. As a patient, if I get that injection and my knee feels better, I don't really care whether it has active drug or placebo. So now we want to look into other placebos, like topical placebos and sham surgeries and sham acupuncture.

In sham surgeries and sham acupunctures, patients are put through experiences very much like surgery or acupuncture, but receive no treatment. In acupuncture, for example, the patient would have acupuncture needles inserted, but not in the specific ways prescribed by trained practitioners. Surprisingly, due to the placebo effect, these sham procedures sometimes produce pain relief. We want to measure that placebo response as well.

If placebo effects are such powerful painkillers, could they be used to treat knee pain?

That's a philosophical question. As a physician, to make it ethical, I'd have to disclose it. I'd have to tell my patient that it's a placebo, that it

doesn't contain any active drugs. But then if I disclose that, I don't know if it will still have the same painkilling effect. I'd like to study that. If it makes your knee feel OK for four weeks, why not just get a saline shot? That would be awesome. How we would pay—or get someone to pay—for the placebo treatment will also play a vital role in this decision-making process.

Provided by Tufts University