

**SDS PODCAST
EPISODE 514
FIVE MINUTE
FRIDAY:
DOES CAFFEINE
HURT
PRODUCTIVITY?
(PART 2)**



(00:06):

This is Part 2 of a series of Five-Minute Fridays on Whether Caffeine Diminishes Productivity.

(00:19):

On last week's Five-Minute Friday, I described the reasons behind my suspicion that caffeine diminishes productivity. So here they are briefly. First, I'm consciously aware of my mind darting off on task-irrelevant thoughts the more recently I've had a cup of coffee. Second, the most productive person I know doesn't drink caffeine. And third, during the only period of my life where I can objective and quantitatively confirm that I was exceptionally productive, I didn't drink caffeine. Given those pieces of evidence, I decided I wanted to conduct an experiment. Conveniently going back to 2016, I have a daily record of a few dozen habits that I track. One of those habits is the number of cups of coffee that I drink. Another one is the number of Pomodoros, 25-minute sessions of uninterrupted, deeply focused work. Check out episode 456 for more detail on this Pomodoro productivity technique.

(01:22):

My hypothesis going into the experiment was that my productivity, as demonstrated by the number of Pomodoros I complete in a given Workday, would be higher on no coffee days than on coffee days. Well, wait no longer. The results of my experiment are in and I've analyzed them. So if you head to johncrone.com/coffee, You can check out my Jupiter notebook for yourself. It includes the experimental data, my code, and the results. Again, that's johncrone.com/coffee. And I'm going to walk you through the notebook right now. All right. So first let's talk about where the data come from. So I have two groups of data. One is from when I was still drinking coffee, so this is from the final two work weeks that I was drinking coffee,



Monday, August 2nd to Friday, August 20th. After that, I began weaning myself off of my two-cup per day coffee habit.

(02:23):

So then the second set of data, the no coffee group data are from the first two work weeks after I'd finished weaning myself off of coffee. So that was Monday, September 13th, through to Friday, September 24th. All right, so it would be boring for a podcast for me to read through the individual data points, but there are 10 of them in each of the two groups, corresponding to two work weeks. So the coffee group has 10 integers in it, and each of those integers represents the number of Pomodoros, again, 25-minute blocks of uninterrupted, deeply focused work that I completed on a given Workday. And the no coffee group, well, let's got another 10 integers. The coffee group has a mean of 10.9. So an average of 10.9 25-minute blocks of uninterrupted, deeply focused work in the baseline condition, in the coffee condition. And then after I'd weaned myself off of coffee, the mean jumped up to 15.4.

(03:32):

So, well, that's interesting. And if you go into the Jupiter notebook, you can see box plots that show the distributions of the coffee group relative to the no coffee groups, so you can compare those side by side. And so based on either the mean values, the averages, or these plots, it looks clear to my eye at least that the no coffee group has much higher values, that I tend to do significantly more Pomodoros than when I drink coffee. But we are armed with statistical tools. We don't just need to guess. We don't just need to say, "Okay, 15.4 sounds bigger than 10.9." And we don't need to just look at the distributions and say, "Well, it looks like the no coffee group is different from the coffee group. It looks like the no coffee group is higher." We can run a statistical test to compare the two.

(04:30):

So specifically, I used something called a Welch's t-test. And so this is a statistical test for comparing two distributions against each other to see if they have different means. And this particular test does not assume equal variance in the two groups. So that just means we're making less assumptions about what the real world data look like if we had an infinite number of samples, what the population data looks like. So after running that t-test, well, unsurprisingly, we see a very strong statistically meaningful relationship. So the no coffee group has a very statistically significant increase in the number of Pomodoros that I complete, specifically on those days where I had 15.4 Pomodoros in the no coffee group corresponding to 6.4 hours of deeply focused work. That was statistically significantly more than on the days that I did drink coffee, where I only had 10.9 Pomodoros corresponding to just 4.5 hours of deeply focused work.

(05:43):

So about two extra hours, or almost 50% more hours of deeply focused work on the days that I didn't drink coffee than on days that I did. And that is hugely valuable for me because that deeply focused work is where I do some of my most interesting things, like recording podcast episodes, videos for YouTube, writing books, and so on. So that time is hugely valuable to me, and I hadn't had any coffee since running this experiment. If you'd like to know the specific values that came out of that t-test, it had a t-value of 3.6, which corresponds to a p-value of 0.002, meaning that there is a less than a 1% chance that we would observe there's a difference by chance alone, making it seem crystal clear it's extremely likely that not drinking coffee is much better for my productivity.

(06:43):

Now, I'm just one person. So while these results are interesting, it's not a huge sample size. And the other issue is that I was also the experimenter. So I was the subject and the experimenter in this experiment, and that is



not a great design if you can avoid it. So it's better scientifically to have double-blind experiments, which means that not only is the subject unaware of whether they're in the experimental condition or the control condition, but the experimenter also isn't aware. So in this case we had the opposite, I was the subject and the experimenter, and I was definitely aware whether I was in one condition or the other. So that could definitely be a factor. Although I didn't think to do this experiment until after I'd collected the data. So I don't know. So I don't think I was deliberately fudging the results here. So anyway, because of these limitations, just having me running this experiment and not being double blind, next week in part three of this series on whether caffeine diminishes productivity, I'll come back with a quick review of what the experimental evidence from the scientific literature suggests on this topic.

(08:05):

All right, that's it for today's episode, keep on rocking it out there, folks, and catch you on another round of SuperDataScience very soon.