Thought: The Great Coffee Experiment

7/7/16
By Kevin DeLuca
ThoughtBurner

Opportunity Cost *of Reading this ThoughtBurner post:* ***$1.97*** *– about* ***8.95 minutes***

I drink a lot of coffee.

In fact, I'm drinking a cup as I write this article. I drink coffee because I think that it makes me more productive and gives me a slight boost in energy which allows me to work more on whatever it is I happen to be working on. As a ***Daily Optimizer,***[i] it makes perfect sense for me to drink more coffee if it truly does make me more productive. And while I suspect that any boost in work productivity from coffee-drinking is small, sub-optimal coffee drinking habits could lead to large productivity losses over a lifetime.

I'm not alone in my belief in the power of coffee. An old survey found that 65% of workers drank coffee on the job and 38% of people say they ***"couldn't live without it."***[ii] Another study showed that ***half of all people say that they are less productive without coffee.***[iii] Clearly, many people believe in the productivity-boosting power of coffee. What I worry about, however, is that people tend to drink coffee precisely when they *have* to be more productive. They sip that late night java as they cram for an exam, or chug a second cup on a particularly busy morning at work. Does the coffee actually help, or is it all in their head? Causation or correlation?

From what I could find online, there has been a lot of research done on the effects of coffee-drinking. But the evidence looks pretty mixed, and it is less about work productivity and more about health. It seems like every few weeks or so, an article goes around social media talking about a new study (often with a ***sensationalist headline***[iv]) that proves that coffee is bad for you. Then the next week, one will say that coffee is ***cancer-curing,***[v] or, at worst, harmless. Some research has shown that caffeine boosts performance on a number of simple cognitive tasks. Many of these studies, however, take place in highly controlled environments and test performance on ***"memory, reasoning and concentration tasks,"***[vi] and it's not clear whether these lab results generalize to the real world. These health effects are interesting for sure, but people don't worry about ***minimizing their risk of dying within 17 years***[vii] when they switch on their coffee pots in the morning. They drink it because they believe it's optimal to drink.

At least, that's certainly true for me. I drink coffee because I *think* it makes me more productive. But I don't really *know* that it does. This issue has plagued me ever since I became an avid coffee drinker. Rather than letting this problem keep me up at night (or is that all the caffeine?), I designed an experiment to get to the bottom of this coffee controversy.

The Great Coffee Experiment

My coffee-drinking self-experiment was a blind study meant to investigate the causal impacts of drinking coffee on my work output. I randomized the timing, amount, and type of coffee (regular or decaf) that I drank during workdays for a six-month period.

Experiment Design

The randomization of my coffee-drinking schedule was carried out in a number of steps. First, each day's coffee type was determined – either regular or decaf. These two conditions happened with equal probability. Next, the number of cups was randomly decided. The number of cups was restricted to being an integer between zero and five inclusive, with each value having the same chance of being chosen. Then, the day was randomly chosen to be either a "random" day or a "free" day. On free days, I was told the number of cups of coffee I was assigned for the day, and was allowed to choose when I drank each cup. On random days, each cup was assigned a time at which it was required to be drank. This random time was restricted to morning and work hours (for example, I did not allow for cups to be assigned a time in the middle of the night while I was most likely asleep). Times were chosen by selecting a random minute within the allowed time interval (6:30am-6:00pm).

For the experiment, I exclusively used Maxwell House instant coffee (either regular or decaf, depending on the assigned type). Coffee was given to me in a plastic bag at the beginning of each day by my research assistant (special shout out and big thank you to Ellen Kim, my research assistant and fiancée, who patiently put up with all of my weird experimental requests). The bags were pre-filled with enough instant coffee for five standard 8-oz cups. Notice that this means that I could not guess how much coffee I would have to drink on a given day based on the amount of instant coffee I observed in the bag each morning. In addition to being "blind" to the total amount, I was also "blind" to the coffee type assignment - the bags did not indicate whether the coffee was regular or decaf. After the timing for each cup was randomized, I used a web service called ***OhDontForget***[viii] to send myself a text message at each of the randomly assigned times.

On free days, I was sent a text message at 6am that told me the number of cups I was required to drink that day. On random days, I received a text message exactly at each of the randomly assigned times. My research assistant entered in all of the times so that I did not know beforehand when I would be required to drink each cup. On random days, whenever I received a text message I was *required* to make a cup of instant coffee from the bag I had been given that morning, and I had to finish the entire cup within 30 minutes. Compliance to the experimental conditions was good; for 92% of days, I stuck perfectly to the randomized schedule.

This design solves two major endogeneity problems present in observational studies of coffee-drinking. First, by randomizing the coffee types and having the study be a blind study, we can control for any placebo effects. In a non-blind study, I would know how many caffeinated cups of coffee I drink on a given day, and therefore would know my caffeine dosage. In this case, we might worry that I might change my behavior based on how many cups I get to drink (i.e. increase my work effort if I get to drink more coffee), creating what seems to be a causal relationship between caffeine and work output. In this

experiment, however, the placebo effect will influence both the regular and decaf work output equally due to the blind assignment. Therefore, we can control for (and measure) this placebo effect.

Second, by randomizing the timing and number of cups of coffee, we eliminate any endogeneity between choosing to drink coffee and the amount of work I have to complete. For example, it seems likely that, if people believe that drinking coffee increases their productivity or energy, that they would drink more coffee precisely when they have a lot of work to complete. This would bias any estimates of coffee drinking on productivity upward, since people would choose to drink more coffee at times during the day when they are required to work more. The randomized amount and timing of coffee-drinking eliminates that bias.

In addition to solving these two endogeneity issues, the design of the experiment also allows us to measure how good I am at naturally optimizing my own coffee-drinking in terms of the timing of my coffee consumption. If people are good at deciding the best times for them to drink coffee, we might imagine that when my coffee-drinking schedule is randomly determined the positive effects of drinking coffee may not be as strong. For example, coffee may boost productivity more if I drink it when I start to feel tired during the day (and choose to drink it) rather than when I am randomly assigned to drink it at a time when I may or may not be feeling tired. By comparing the estimates of the coffee effects on random days versus free days, we can get a sense of how good I am at being a ***Daily Optimizer***.[ix]

The outcome of interest was my daily labor supply – the number of minutes I worked in a day. I measured this by taking a stopwatch with me to my office. Whenever I would start working on a task, I would start the timer. Whenever I stopped working – whether it was to go have lunch, read a ***WaitButWhy***[x] article, or go to the coffee shop across the street – I would stop the timer and mark down how long I had worked. This outcome measure, along with a record of the randomized coffee drinking assignment, is what I use to estimate the causal effects of drinking coffee on my work output and work efficiency.

Results

Figure 1 shows the best linear fit line for average total daily work time conditional on coffee type and number of cups, both for free and random days. For free days (left side of Figure 1), I worked more on regular type days as compared to decaf days, as indicated by the caffeinated line being above the decaf line. For both coffee types, total work output increases as the number of cups increases, which is visually suggestive of a positive placebo effect. Additionally, the slopes of both the regular and decaf lines seem to be approximately the same, which suggests that there is no positive effect of drinking an additional cup of regular coffee (as opposed to drinking an additional cup of decaf). In other words, there is no visual evidence that caffeinated coffee increases my labor supply any more than decaf coffee on free days.

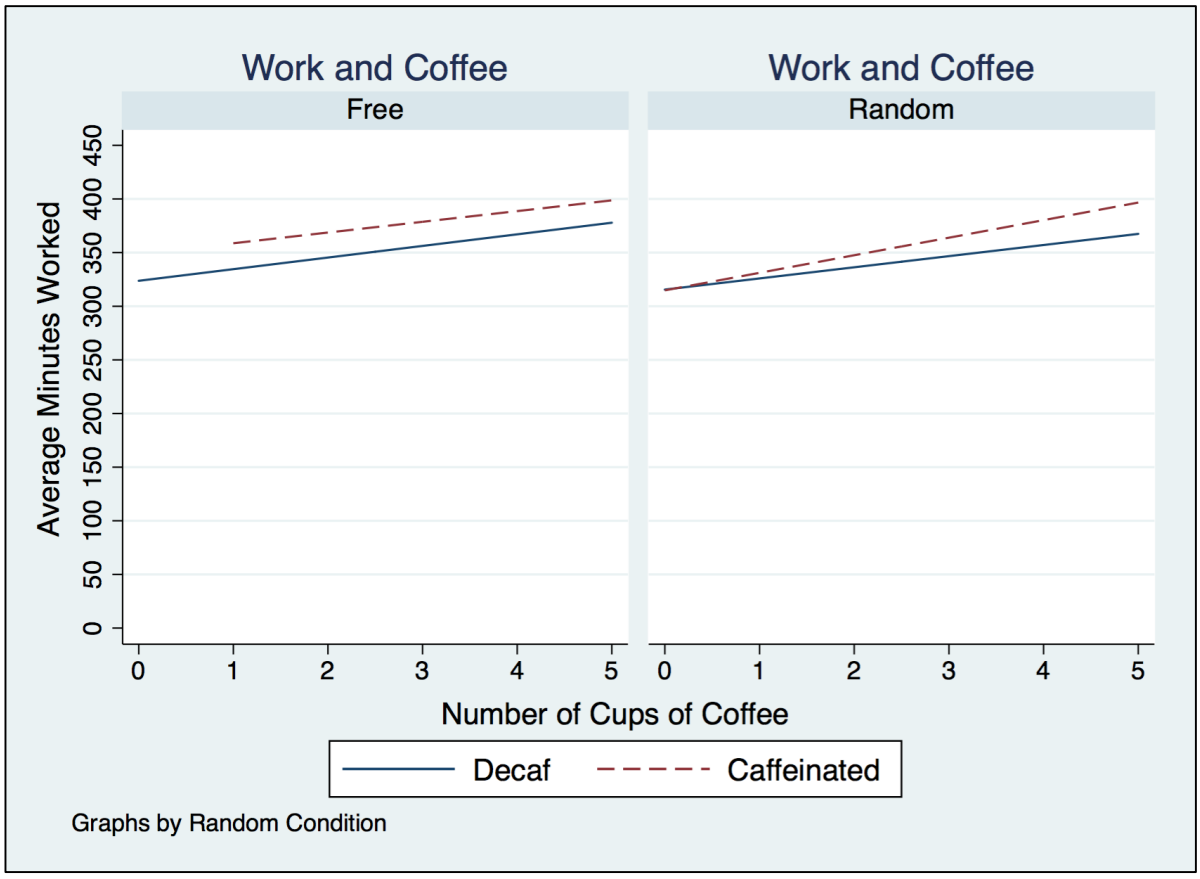

Figure 1: Time Worked and Coffee Drinking, Linear Fit

For random days (right side of Figure 1), the figure tells a slightly different story. Both the regular and decaf lines are still upward sloping, which again suggests a placebo effect. However, the slope for regular days looks steeper than the slope for decaf days. This difference in slope would suggest that the effect of drinking an additional cup of regular coffee increases my work output more than drinking an additional cup of decaf coffee. In other words, there is visually suggestive evidence that there is, indeed, a true effect of drinking caffeinated coffee on my total work output.

To more rigorously check these effects, I ran a regression of daily work time on the different coffee drinking variables. Specifically, I used the following model to test whether any of the effects were statistically significant:

$$Time\ Worked = \alpha + \beta_1 Regular + \beta_2 Cups + \beta_3 (Regular \times Cups)$$

where *Time Worked* is my daily time worked in minutes, *Regular* is an indicator variable equal to one when the coffee type was regular, *Cups* is the number of cups of coffee I drank that day (regular *or* decaf), and *(Regular x Cups)* is the number of regular cups of coffee I drank that day (decaf cups don't count).

The sign, magnitude, and statistical significance of the β coefficients are important. β_1 tells us whether I worked more on days where I drank regular coffee compared to days when I drank decaf (not

dependent on how many cups I had). β_2 measures whether an additional cup of coffee, either regular or decaf, increases the amount of time worked – in other words, β_2 measures any placebo effect. β_3 picks up any causal impacts of drinking an additional, regular cup of coffee on the total number of minutes I worked in a day. Most people believe that drinking more *caffeinated* coffee will help you get more work done. This is equivalent to believing that β_3 is positive and statistically significant. We will test whether this is true.

Results of the regression are presented below in the table below.

Coffee Drinking and Total Work Time

	Total Work Time	
	Free Days	**Random Days**
Regular	31.29	-16.27
	(37.73)	(48.74)
Cups	5.47	16.87
	(7.50)	(11.58)
Regular Cups	-5.08	17.07
	(12.23)	(17.08)
Observations	*n=38*	*n=34*

Notes: Regular is an indicator for regular type days, Cups is the number of cups (regardless of coffee type), and Regular Cups is the number of regular (caffeinated) cups of coffee. All measurements are in minutes per day.

The estimates from the table generally confirm our visual intuitions, with the caveat that most of the estimates are not statistically significant. For both free days and random days, the estimate for β_2 is positive, indicating evidence of a placebo effect.

For random days, the estimates suggest that drinking an additional cup of coffee - regular or decaf - increases the total daily number of minutes worked by approximately 17 minutes. This is equivalent to a 4.8% increase from the daily mean total time worked. On random days, the estimate for β_3 is also positive, indicating that drinking an additional cup of regular cup of coffee has a stronger positive effect on total work supply than drinking an additional cup of decaf coffee. Specifically, the estimates suggest that drinking an additional cup of regular coffee has *double* the effect of drinking an additional cup of decaf coffee - when the coffee was regular, my total daily work output increased by an

additional 17 minutes per cup. However, neither the estimate for β_2 or for β_3 are statistically significant at conventional levels.

Interestingly, there is no evidence of an "amplifying" effect on free days. In fact, the results of running the regressions separately for free and random days suggest that, if anything, being able to choose the timing of my coffee drinking actually *decreased* the effect of drinking an additional cup of regular coffee, though I don't want to read too much into this result because the estimates are not statistically significant.

Discussion

In terms of my personal results, there seems to be slight evidence of both a placebo effect and "true" effect from drinking caffeinated coffee. I suspect that part of what is really happening, for me, is that when the experiment forced me to drink less coffee than I normally would, I worked less. Naturally, I probably drink about four regular cups of coffee per day. Whenever I had to drink less than this, I may have decreased my work output. It's not so much that coffee makes me work more, just that not being able to drink my "natural" amount of coffee makes me work less. This would explain why both regular and decaf coffee seemed to increase my total work time.

While these results are real for me, they may not be the same for you. Caffeine and coffee can have very different effects on different ***people,***[xi] and I wouldn't be surprised if the effects were only applicable for me or a particular subset of people who are very similar to me.

So what is the solution to the daily optimization problem of how much coffee to drink? Given that the results of the experiment look promising for coffee consumption, I have since lifted restrictions on my coffee-drinking in order to maximize my work output and productivity. I now drink a combination of caffeinated and decaf coffee in order to hedge against possible long-term negative health consequences of daily caffeine consumption, in part because it seems to be the case that I can still increase my total work output even while drinking decaf coffee (though, perhaps, slightly less than if I exclusively drank regular coffee). Further (self-)research is needed to get more precise estimates of these potential coffee-drinking effects, and to be certain whether this is truly an optimal course of action.

---

[i] https://thoughtburner.org/2015/02/19/thought-daily-optimization/
[ii] http://www.businessinsider.com/the-one-office-perk-you-must-splurge-on-2011-3
[iii] http://fizzler.com/167/can-coffee-add-to-productivity/
[iv] http://www.atlantamagazine.com/health/new-study-coffee-can-kill-you/
[v] http://well.blogs.nytimes.com/2015/01/22/coffee-may-cut-melanoma-risk/
[vi] http://www.telegraph.co.uk/news/health/news/7710780/How-office-coffee-breaks-make-staff-work-harder.html
[vii] http://www.npr.org/sections/thesalt/2013/08/17/212710767/how-many-cups-of-coffee-per-day-is-too-many

[viii] http://ohdontforget.com/
[ix] https://thoughtburner.org/2015/02/19/thought-daily-optimization/
[x] http://waitbutwhy.com/

[xi] http://earnworthy.com/caffeine-and-productivity/