Coffee Consumption and Labor Supply: A Self-Experiment

Kevin DeLuca

December 12, 2016

Abstract

Does drinking coffee actually make workers more productive? Many people believe that, through the ingestion of caffeine, drinking coffee helps them work harder/better while they are at work. This paper presents the results of a coffee-drinking self-experiment which was meant to test whether there is a causal relationship between drinking caffeinated coffee and work output. The experiment randomized the amount, timing, and type of coffee (decaf or regular) I drank during workdays for a six-month period. The results provide weak evidence of a positive effect of drinking caffeinated coffee as well as a positive placebo effect of drinking coffee. Drinking a cup of coffee - regular or decaf - was associated with an increase in my total daily work time of approximately 17 minutes per cup, and when the coffee was regular (as opposed to decaf) my total daily work time was an *additional* 17 minutes higher per cup. While both estimated effects are positive, neither of the effects are statistically significant at conventional levels.

1 Introduction

This project began due to my personal interest in what I call "Daily Optimization" problems.[1] In a nutshell, daily optimization problems are questions that could be solved with the analytical framework and toolkit used by economists, but are "not worth" solving. They are the daily questions regarding optimization that we often ask ourselves, but often never answer.

Imagine, for example, that you want to know whether drinking coffee makes you more productive at work and, if so, how many cups is optimal to drink. While you might assume that any daily increases in work productivity due to drinking coffee are likely small, over a lifetime sub-optimal coffee-drinking habits could lead to large productivity losses. If you had data on your total work output or productivity, coffee-drinking habits, and a lot of other controls, you could use an economic model to estimate the effect of drinking coffee on your productivity. This would allow you to describe and quantify the effect that drinking coffee has *specifically for you*, and modify your behavior accordingly. While an interesting idea, this sort of daily optimization project is typically never performed. This is due to a combination of the high costs of data collection in these types of situations, and the relatively low expected payoff of "solving" the daily optimization problem.

These problems, however, are exactly the sort of problems that I like to answer. So, I had the idea to conduct a coffee-drinking self-experiment to see whether drinking coffee increased my work productivity or labor supply. I designed an experiment that randomized when, how much, and what type of coffee (regular or decaf) I drank throughout the day in order to estimate the causal effect of drinking coffee on my personal labor supply.

[1] I coined this term over a year ago when I first created my blog, ThoughtBurner. See www.ThoughtBurner.org for the original article that describes what I call daily optimization problems, and for more projects and analyses with a similar flavor as this paper.

2 Coffee and Labor Supply

2.1 Previous Studies

There have been a number of studies and surveys that document people's attitudes towards coffee consumption. An old survey found that 65% of workers drank coffee on the job and 38% of people said they "couldn't live without it."[2] Another study showed that half of all people say that they are less productive without coffee.[3] Clearly, many people believe in the productivity-boosting power of coffee. The problem with these surveys, of course, is that people tend to drink coffee precisely when they have to be more productive, which makes it difficult to tease out the causal effect of coffee on work productivity (if there indeed is one).

From what I could find online, there has also been a lot of research done on the effects of coffee-drinking in general. The evidence, however, looks pretty mixed, and it is less about work productivity and more about health. For example, some studies warn that coffee may be "killing" people,[4] while others point to evidence that suggests coffee may be cancer-curing, or, at worst, harmless.[5] Some studies show that caffeine improves performance on a number of different cognitive tasks. Many of these studies, however, take place in highly controlled environments and test performance on "memory, reasoning and concentration tasks,"[6] and it's not clear whether these lab results generalize to the real world. These health effects are interesting, but people don't worry about minimizing their risk of dying within 17 years[7] when they switch on their coffee pots in the morning. They drink it because they believe it's optimal to drink. My experiment was designed, in part, to see whether drinking coffee truly does make people more productive.

[2] http://www.businessinsider.com/the-one-office-perk-you-must-splurge-on-2011-3
[3] http://fizzler.com/167/can-coffee-add-to-productivity/
[4] http://www.atlantamagazine.com/health/new-study-coffee-can-kill-you/
[5] http://well.blogs.nytimes.com/2015/01/22/coffee-may-cut-melanoma-risk/
[6] http://www.telegraph.co.uk/news/health/news/7710780/How-office-coffee-breaks-make-staff-work-harder.html
[7] http://www.npr.org/sections/thesalt/2013/08/17/212710767/how-many-cups-of-coffee-per-day-is-too-many

2.2 Experiment Design

I first randomized the amount, timing, and type of coffee I drank for a period of six months, from November 2015 through the end of April 2016. This randomization is necessary in order to estimate the causal effects of drinking coffee on my total daily work output. I recorded my work time for the entire six-month period - using a stopwatch, I timed the number of minutes I spent working on different tasks. Whenever I would start working on a task, I would start the timer. Whenever I stopped working – whether it was to go have lunch, talk with coworkers, or go to the coffee shop across the street – I would stop the timer and mark down how long I had worked. The cumulative daily total of this work time measurement is the main measure of labor supply that I use in this paper.

The randomization of the coffee-drinking schedule for each day was carried out in a number of steps. First, the coffee type was randomized - either regular or decaf. These conditions occurred with equal probability. Next, the number of cups was randomly decided. The number of cups was restricted to integer values between zero and five inclusive, each value with 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 minute at random within the allowed time interval.[8]

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.[9] 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.

[8] In general, the time interval was from 6:30am to 6:00pm.

[9] Special shoutout and a big thank you to Ellen Kim, my research assistant and fiancée, who put up with all of my weird experimental demands for six months.

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 service called OhDontForget[10] to send myself text messages with coffee drinking instructions. 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 whether the day would be a free day or a random day and, if it was a random day, when I would be required to drink coffee. On random days, I had to finish the entire cup of coffee within 30 minutes of receiving a text message.

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 should 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 on random days, 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, 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 on days and at times 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

[10] http://ohdontforget.com/

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 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 my productivity more if I drink it when I start to feel tired during the day (and choose to drink it) compared to 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.[11] This part of the investigation is related to a deeper question about economics: how good are people at making optimal decisions, especially when facing these sort of everyday but low-stakes types of problems? One of the goals of this experiment and analysis is to provide a very small contribution to answering this question.

2.3 Method

With this experiment design in mind, we can run the following simple regression to investigate the causal effect of caffeinated coffee drinking on my individual labor supply:

$$Y = \alpha + \beta_1(regular) + \beta_2(cups) + \beta_3(regular * cups) + \epsilon \tag{1}$$

Where Y is my total daily labor supply in minutes, *regular* is an indicator equal to one when the coffee type is regular coffee, *cups* is the number of cups (regardless of coffee type), and $regular * cups$ is the number of regular cups of coffee.

The coefficient β_1 controls for any systematic differences between regular type and decaf type days. The coefficient β_2 controls for any effects of drinking more coffee, regardless of coffee type (i.e. measures the placebo effect). And the coefficient of interest, β_3, is a measure of the causal impact of drinking an additional cup of regular coffee on my labor supply. Most people believe that drinking more *caffeinated* coffee increases work productivity. If this is true, we would expect β_3 to be positive and significant.

[11] https://thoughtburner.org/2015/02/19/thought-daily-optimization/

By comparing the estimates of the coefficients of this regression run separately for random days and free days, we can also see if I am able to successfully optimize the timing of my coffee drinking. In other words, we can see if being able to choose when I drink coffee increases the effects of drinking coffee on measures of my individual labor supply. Specifically, if we expect that the timing matters and that people can strategically time their coffee-drinking to maximize the positive effects on work output, we would expect for β_3 to be *greater* (more positive) on free days.

2.4 Results

For the results that follow, I exclude weekends, winter break, spring break, and any other vacation days from the sample. These restrictions limit the analysis only to "typical" work days, which limits potential external confounders that might affect my labor supply. In total, there are 72 days that meet these criteria that are used in the results - 38 free days, and 34 random days. Compliance to the experimental conditions was good. On only two days were the experimental conditions broken (about 97% compliance).

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.

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

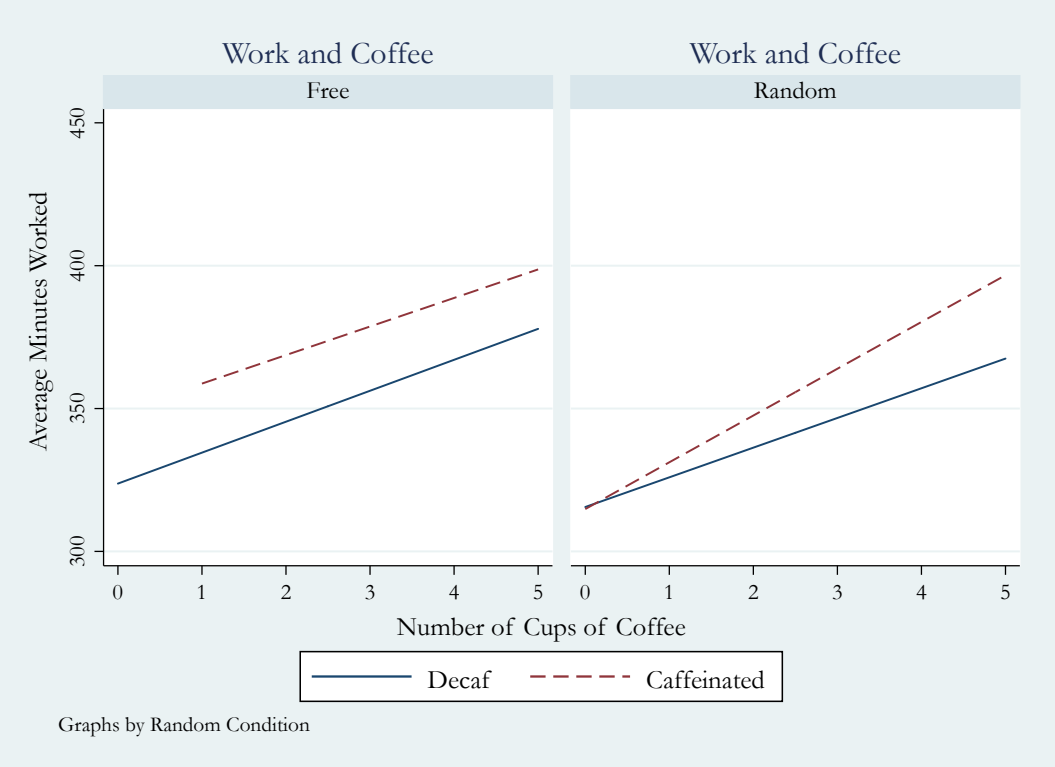

Figure 1: Average daily total work time conditional on number of cups of coffee and coffee type.

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. I use regression (1) to test whether these visual differences are statistically significant.

Table 1 shows the results of regression (1) on my daily total work supplied (in minutes), separately for both free and random days. 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

Table 1: Coffee Drinking and Labor Supply

	Total Work	
	Free	Random
Regular (β_1)	31.29	-16.27
	(37.73)	(48.74)
Cups (β_2)	5.47	16.87
	(7.50)	(11.58)
Regular Cups (β_3)	-5.08	17.07
	(12.23)	(17.08)
Observations	*38*	*34*

Notes: Total work is in minutes per day. Regular is an indicator equal to one when the coffee type is regular, Cups is the number of cups (regardless of coffee type), and Regular Cups is the number of regular cups of coffee. Significant effects are in bold; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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.

In retrospect, the variance in the number of minutes I work each day was higher than I expected, and given that I have a relatively small sample size my estimates are fairly imprecise.[12]

[12]In particular, no days were randomly assigned as free day by regular by zero cups, and this missing measurement may have significantly influenced the estimates, particularly if comparing the free vs random days results.

This would make it difficult to detect any labor supply effects due to coffee drinking, because these effects are likely to be small.

3 Discussion

Importantly, I did not want to influence my own behavior by previewing the results of my study or by observing my labor supply. As such, I did not begin analysis of the data until after the experiment concluded. Of course, there is always a concern that I may have acted differently if I had known that I *wasn't* being observed (i.e. Hawthorne Effect).

It could certainly be the case that these estimates are not generalizable to most workers. Caffeine and coffee can have very different effects on different people,[13] and I wouldn't be surprised if the effects were only applicable to me or a particular subset of people who are very similar to me. It should also be noted that I am a fairly avid coffee drinker, so if my body had habituated to drinking a lot of coffee before I started the experiment, the effects may be different compared to someone who rarely drinks coffee.

It could also be the case that the placebo effect would disappear if I *know* that the coffee is decaf. Maybe the placebo effect is only present if I believe there is a *possibility* that the coffee is caffeinated. Additionally, there is a chance that, as a result of learning the results of the experiment, the psychological factors that gave rise to the placebo effect (e.g. believing that caffeinated coffee makes me more productive) may have been changed in a way that could affect my future coffee drinking habits and/or the effect that coffee has on my labor supply.

In the future, it would be interesting to scale this experiment up to include many participants in many different workplaces. This would allow for more precision in the estimation, and would lead to more credibly generalizable results. It could also be interesting to see if coffee drinking affects certain jobs/tasks differently.

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

In terms of my personal results, I suspect that what is really going on here is that when the experiment forced me to drink less coffee than I normally would, I worked less. Naturally, I probably drink about 4 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.

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.