Framing Flexible Spending Accounts: A Large-Scale Field Experiment on Communicating the Return on Medical Savings Accounts

Jessica Leight and Nicholas Wilson*

Current draft: April 11, 2019

Abstract

Tax-preferred health savings devices such as Flexible Spending Accounts (FSAs) and Health Savings Accounts (HSAs) offer employees potentially valuable financial instruments for directing pre-tax earnings to eligible medical expenses. Despite their increasing popularity as an employee benefit, however, there is no causal evidence around individual demand for these accounts. This paper seeks to address this gap in the literature, reporting on a randomized controlled field experiment conducted with over 11,000 U.S federal employees in 2017 in order to evaluate the effectiveness of targeted messages designed to increase FSA contributions. Our results suggest that the provision of basic information about FSAs delivered via an emailed employee newsletter did not affect the likelihood of contribution or the contribution level. The addition of statements about the absolute returns or relative returns offered by the accounts similarly had no significant effects, and these null effects are observed despite relatively high email open rates. We discuss explanations for the null results and the policy implications of findings from what appears to be the first health economics experiment analyzing tax incentives around health care savings.

Keywords: health; savings; taxation
JEL codes: D14; D91; H24; I11; I18

---

* Leight: Department of Economics, American University, Washington, DC, USA and Fellow, Office of Evaluation Sciences, Washington, DC, USA. Wilson: Department of Economics, Reed College, Portland, Oregon, USA, and Fellow, Office of Evaluation Sciences, Washington, DC, USA, nwilson@reed.edu. We thank Kelly Bidwell, Jake Bowers, Lula Chen, Michael Hand, Mark Pauly, and seminar participants at the University of Pennsylvania and the Office of Evaluation Sciences for many helpful comments. We thank Allyson Barnett for assisting in the email design, Dan Corker for guidance on the GSA HR system, Ben Kistler for pre-study data work, and Jeff Woodworth for helping us place FSA communications in the GSA Today newsletter. This research would not be possible without the insights, support, and guidance of Danielle Fisher. The findings, interpretations, and conclusions expressed in this paper are those of the authors and do not necessarily represent the views of the aforementioned individuals or agencies. All errors are those of the authors.

1 Introduction

Tax-preferred health savings devices such as Flexible Spending Accounts (FSAs) and Health Savings Accounts (HSAs) offer employees financial instruments for spending pre-tax earnings on eligible medical expenses. These devices reduce the effective price of eligible medical expenses, and thus may increase the quantity of health services demanded for households that are price-sensitive or otherwise constrained in investing in such services; from the policymaking perspective, promoting the use of such accounts may be a viable strategy to increase household investment in health inputs. FSAs are a common employee benefit in the U.S., with three-quarters of large firms offering the accounts (Kaiser Family Foundation 2016). Yet fewer than half of eligible employees participate (Mercer 2008), and there is little large-scale evidence on methods designed to increase participation in these savings devices. Addressing this gap in the literature, this paper reports on a large-scale randomized controlled field experiment conducted with over 11,000 federal employees in 2017 in order to evaluate the effectiveness of three different targeted messages delivered via an employee email newsletter and designed to increase FSA contribution levels.

FSAs allow employees to contribute pre-tax earnings from their paycheck to an account from which they subsequently request reimbursement for eligible expenses. The United States created FSAs in the Revenue Act of 1978 with the goal of increasing household purchasing power for medical goods.[1] In 2010, the passage of the Patient Protection and Affordable Care Act (ACA) created a contribution limit, raised periodically and currently equal to $2,650 for an individual or married couples, and allowed for up to $500 to be rolled over into the subsequent year. The Kaiser Family Foundation Employer Health Benefits 2016 Annual Survey of U.S. firms indicates that 13% of all firms and 75% of firms with 200 employees or more offer FSAs (Kaiser Family Foundation 2016).[2] Participation rates among eligible employees appear to be below 50%, and are closer to 25% at large firms (Mercer 2008).

[1] The United States tax code also allows employers to offer Dependent Care FSAs and Limited Expense FSAs. Health Care FSAs appear to be more common than these other two types of FSAs. In our analysis, when we refer to FSA we are restricting our focus to Health Care FSAs.

[2] Take-up at large firms appears to be lower than take-up at small firms. In or around 2008, it was approximately 22 and 40 percent, respectively (Mercer 2008).

This evaluation sought to analyze the effectiveness of three different messages delivered to employees and designed to increase FSA contributions. The first message, denoted the basic treatment, stated that an employee could "save money" by using a FSA and provided information about the ease of enrollment in the benefit as well as its utilization. The second message, denoted the absolute return treatment, provided the same information about the benefit, but also highlighted that an average employee would "save $949", given a specified estimated level of out-of-pocket expenses and an estimated marginal tax rate. The third message, denoted the relative return treatment, again provided basic benefit information, but framed the savings as "savings of 36.5%" for the same hypothetical employee. These messages were disseminated through a tri-weekly electronic mail newsletter on the same day at the start of the annual open enrollment period and again four weeks later; each employee assigned to a treatment arm received the same information twice. The fourth study arm received none of these communications and constitutes our control arm.

Our experiment reveals three key facts about FSAs and framing the return on savings via this tax device in this study population. First, participation rates are relatively low; approximately 25% of employees contributed to a FSA, and the mean contribution in the full sample was approximately $450. Second, communicating FSA eligibility and the return on savings in the basic treatment did not increase use or amount contributed. Third, highlighting the absolute or relative savings similarly had no additional effect on increasing contributions. The estimated treatment effects for all three intervention arms are precisely estimated zeros.

Beyond the analysis specified in our pre-analysis plan, we conducted an exploratory analysis of email open and click-through rates.[3] Although approximately 50% of recipients opened the electronic mail newsletter, only around 1% clicked through on the FSA links embedded in the emails. (There is no evidence of differential effects for the small number of employees who clicked to access further information, though these coefficients are noisily estimated.) Accordingly, the null result could reflect the fact that recipients

[3] At the time at which we had to file a pre-analysis plan we had not been able to confirm that we could access email open and click-through data. We discuss our pre-analysis plan in more detail in the Experimental Design section.

did not find the framing statements particularly salient, or were already fully informed. Alternatively, the experimental evidence is consistent with the hypothesis that the perceived utility of the FSA among those not already using a FSA is limited, especially for employees whose out-of-pocket medical expenses are low or variable. We explore these hypotheses further in the Discussion section.

Our study provides the first causal evidence on demand for medical savings accounts of any form. We build on a small set of studies of health savings accounts, including FSAs, HSAs, and other medical savings accounts. Existing economic analyses of FSAs study optimal participation levels (Bhattacharya et al. 2002, Cardon and Showalter 2007), actual participation and its correlates (Levy 1998, Hamilton and Marton 2008, Cardon et al. 2012), and the effects of FSAs on coinsurance (Jack et al. 2006). The HSA literature examines the effects of HSAs on savings and insurance (Steinorth 2011, Ye 2015, Peter et al. 2016). Literature on other medical savings accounts examines the welfare effects of medical savings accounts (Zabinski et al. 1999, Pauly and Herring 2000, Hurley et al. 2008) and whether medical savings accounts in Singapore are sufficient to meet retirement needs (Chia and Tsui 2005).

In addition, we expand the economic literature on framing the return on savings in health accounts. Schmitz and Ziebarth (2015) study the effect of a 2009 German regulatory reform in which insurers were required to change from expressing premium differences in terms of percentage point payroll tax differences to stating premium differences in absolute terms (i.e. euros). It found that the switch to absolute terms (i.e. euros) greatly increased willingness to switch plans by increasing the salience of the differences. We find that framing savings in absolute terms is no more and no less effective than framing savings in percentage terms and that neither communication was more effective than no communication. The magnitudes of the potential gains – on the order of 10-20 euros per month in Schmitz and Ziebarth (2015) and as much as $80 per month in our setting – do not appear to explain the differences in our findings. One possible explanation is that the decision to choose a health insurance plan is much more salient than the decision to choose a medical savings account such as a FSA. Although Schmitz and Ziebarth (2015) find that the absolute framing increased the salience of the premium differences, the baseline attention to the decision to participate in health insurance at all is likely quite high. In contrast, our findings are consistent

with high baseline inattention to FSAs and continued inattention despite direct communication to eligible employees, possibly because among the pool of individuals who do participate in a FSA the potential benefits of FSAs are not large.

Finally, we add to the broader economic literature on nudges designed to increase take-up of public benefits, including tax-based benefits (e.g., Bhargava and Manoli 2015). These studies largely test the effects of paper-based nudges. We demonstrate that email-based nudges in and around the brief window of open enrollment are not sufficient to increase take-up. Furthermore, we demonstrate that re-framing the return on savings was similarly insufficient to stimulate take-up.

Our study appears to be the first health economics experiment using the income tax system, complementing other experiments and nudges in tax systems (Finkelstein et al. 2009, Chetty and Saez 2013). It also appears to be one of the first large-scale health consumer nudge field experiment in economics, complementing existing large-scale experimental evidence on nudging health outcomes using implementation prompts (Milkman et al. 2011) and on nudging health providers using social comparisons (Sacarny et al. 2016). Similarly, our analysis of demand for a government health program complements large-scale experimental evidence on other government health programs such as Medicaid (Finkelstein et al. 2012).

The rest of the analysis is organized as follows. Section 2 describes the experimental design. Section 3 discusses the data and statistical methods. Section 4 presents the results. Section 5 discusses the main findings and evaluates several hypotheses about the mechanism underlying the null results. Section 5 concludes.

2 Experimental Design

The experiment analyzed in this section is a randomized controlled trial conducted by the researchers as fellows with the Office of Evaluation Sciences (OES), working in collaboration with a large federal agency. OES is housed at the U.S. General Services Administration (GSA), and works to translate and tests evidence-based insights into concrete recommendations for how to improve government. In this

project, OES partnered with the GSA Human Resources (HR) department in order to test strategies to increase utilization of and contribution to FSAs.

The sample for the evaluation included the universe of GSA employees, and each employee was randomly assigned to one of four experimental arms. Randomization was conducted by the research team in Stata using the list of recipients of the agency's tri-weekly information newsletter, GSA Today. Individuals assigned to the three treatment arms received an additional message included in the GSA Today newsletter focusing on FSAs, while individuals assigned to the control arm received no additional information about FSAs. The messages were sent twice in the first and last week of open enrollment, approximately one month apart.

In the first, basic treatment arm – an arm in which employees received information about the FSA benefit – the primary statement included in the body of the newsletter included the following information.

> "Do you have medical, dental or vision expenses? Save money by enrolling in a Health Care FSA today! A Health Care FSA (HCFSA) is a pre-tax benefit account that's used to pay for eligible medical, dental, and vision care expenses that are not covered by your health care plan or elsewhere. Enroll by December 11, 2017 to claim your savings!"

Employees could then click through to a more detailed one-page summary that provided some examples of expenses eligible for a FSA, noted that contributions were pre-tax, and highlighted that up to $500 in the account could be rolled over into the next year.

In the second and third treatment arms, employees received the same basic information, but the messages also sought to explicitly highlight the potential savings from utilizing a FSA. In the second absolute returns arm, the text included the same basic information about the account's functioning, but also noted that "A GSA employee with a marginal state and federal income tax rate of 36.5% spending $2600 per year for out of pocket medical expenses would save $949 over the course of a year." In the third relative returns arm, the text noted that "A GSA employee with a marginal state and federal income tax rate of 36.5% would save that rate on out of pocket medical expenses." All three messages also highlighted the deadline for enrollment. Importantly, in each of the treatment arms, the email subject heading was aligned with the distinct content of the email. Thus, the subject heading in the basic arm was "Save money", while

the absolute returns arm used the heading "Save up to $949" and the relative returns arm used the heading "Save up to 36.5%".

These communication strategies were designed to address two key factors affecting employees' utilization of FSAs. The first is limited salience of the accounts or low levels of information; some employees may not have been aware of the benefits of enrollment. The second is procrastination, as employees may have been aware of FSAs and had the objective of enrolling, but failed to do so. In this agency, information about enrollment in FSAs is disseminated during on-boarding for new employees, and thereafter is shared periodically via employee message boards. However, it was hypothesized that these messages may not have been widely read. At the beginning of our study, around 29% of GSA employees were also enrolled using a specific benefits manager (FSAFEDs), and this vendor also directly contacted employees with information about open enrollment. While the presence of other communications is not a source of bias in our experimental design, it potentially narrows the scope for these targeted communications to have a significant effect.

3 Data and Statistical Methods

3.1 Data

Our study utilizes administrative data available to the human resources department of the participating agency, and anonymized for the purposes of the evaluation. While this data has the advantage of not relying on self-reports by employees, a relatively limited set of variables are available as compared to standard individual and household surveys.

More specifically, data available includes whether the employee contributed to a health spending FSA for the benefit year of 2018 – the cycle that is the target of this intervention – as well as the amount of the contribution. We also have an indicator variable for whether the employee contributed in the previous year, though the amount of the contribution is unreported. The only additional covariates available are basic characteristics of the employee's position at the agency (duration of employment, office, salary rank, etc.) Demographic characteristics such as income, familial status, etc. are not available. As is standard

OES practice, a time-stamped pre-analysis plan specifying the outcomes of interest for this evaluation was pre-filed on our website.[4]

Table 1 reports summary statistics for FY18 FSA contributions. Approximately 25% of employees contributed to a FSA in FY18 and the mean contribution was approximately $450. Figure 1 displays the distribution of contribution amounts conditional on contributing more than $0. Conditional on contributing a positive amount to the FSA, the mean contribution amount was approximately $1,775. A notable fraction of the sample contributed the maximum amount for one person for a Health Care FSA (i.e. $2,650). There is also a smaller mass at $500, the maximum amount that can be rolled over into the following year. Many employees contributed multiples of $500, consistent with employees using this as a heuristic for choosing the contribution amount.[5]

3.2 Statistical Methods

The primary regression specification can be written as follows.

$$contribution_i = \gamma + \beta\, treatment_i + \varepsilon_i \qquad (1)$$

In this specification, $contribution_i$ is the amount the individual contributed to a FSA (defined equal to zero if s/he did not contribute), $treatment_i$ is an indicator variable equal to one if individual i was assigned to the treatment group, and ε_i is an idiosyncratic error term. We use ordinary least squares (OLS) regression to estimate the parameters of Equation (1) and employ heteroskedasticity-robust standard errors. We also estimate the same equation using two additional dependent variables: an indicator variable equal to one if the employee reports any contributions to a FSA, and an indicator variable equal to one if the employee

[4] Our pre-analysis plan (Project 1733) may be found here: https://github.com/gsa-oes/office-of-evaluation-sciences/tree/master/assets/analysis

[5] Our data are administrative data on exact contribution amount, not self-reported contribution amount. Thus, "heaping" at multiples of $500 cannot reflect employee-driven measurement error in contribution amount, as is often the case in self-reported expenditures.

reports contributing more than $500. In addition, we estimate a parallel specification including indicator variables for all three treatment arms, and we estimate separate effects for employees who did and did not report contributions to a FSA in the previous year.

4 Regression Results

4.1 Main results

Figure 2 conveys the primary results in graphical format, depicting the mean level of contributions by experimental arm (in Figure 2a) and the probability of contributing more than zero or more than $500 by arm (in Figures 2b and 2c). It is evident that there are no significant differences in the level or the probability of contributions across arms.

Expanding on these results and presenting the results of formal hypothesis tests, Table 2 reports the results of estimating the simple specification using a single indicator variable for assignment to treatment. In Panel A, the specification of interest is estimated using the full sample; Panel B reports the results for the subsample of employees who had not contributed to a GSA FSA in the previous fiscal year, and Panel C reports results for the subsample of employees who had previously contributed.[6] It is evident that the coefficients of interest are consistently small in magnitude and statistically insignificant. The coefficient in Column (1) of Panel A suggests employees assigned to receive targeted information about FSAs contributed approximately $4.05 more on average relative to a mean contribution level of nearly $450; this is a proportional effect of less than 1%. Similarly, the coefficient estimated in Column (2) suggests a decline in the probability of any contribution of 0.6 percentage points; relative to a mean contribution probability of 25%, a proportional effect of around 2%. There is no discernible difference in this pattern evident in Column (3), capturing the probability of contributing at higher levels, or for the results estimated using subsamples.

[6] Approximately one-quarter of employees in our study contributed to a GSA FSA in FY17.

Importantly, the estimated coefficients are also relatively precise, allowing us to rule out even quite small positive effects of the messaging intervention on contributions. Examining the 95% confidence intervals, we can rule out an increase in contributions larger than around $42 for the pooled treatment, or an increase in the probability of making any contribution larger than around 1.4 percentage points.

Table 3 reports parallel results employing indicator variables for each treatment arm. The results reveal a consistent pattern: the point estimates are small in absolute terms, they are small relative to sample mean outcomes, and they are not statistically different from zero. In addition, the hypothesis that the treatment effects are identical comparing across different treatment arms cannot be rejected.

Finally, Table 4 examines heterogeneous effects with respect to employee characteristics: in particular, whether the employee is based in the capital (Washington, D.C. region); the employee's length of tenure; and an indicator variable equal to one for tenure of ten years or more. The objective of these tests is to explore whether employees with certain characteristics are more likely to be responsive to messaging around FSAs. For example, a location in Washington, D.C. may be associated with a more senior position and thus a higher marginal tax rate, rendering use of a FSA a more attractive benefit. Tenure is positively correlated with age, and this may be a (weak) proxy for a higher level of medical expenses. However, estimated coefficients on the treatment variable as well as the interaction term are consistently insignificant, suggesting that there is no evidence that particular subgroups of employees are differentially responsive to the messages provided around the benefits of FSAs.

4.2 Exploratory analysis

In addition to the data evaluated above following our pre-analysis plan, we have access to individual-level data capturing whether employees clicked on the embedded links in the informational treatments to access more information. Using the "click" variable, we estimate the following specification to evaluate whether there is any evidence of differential click-through rates by treatment arm.

$$click_i = \gamma + \beta_1 treatment_{1i} + \beta_2\, treatment_{2i} + \beta_3 treatment_{3i} + \varepsilon_i$$

Table 5 reports the results of these regressions. The first column employs a pooled variable for assignment to treatment, and indicates that the average click-through rate across all treatment arms is 0.7%.[7] The second column estimates the specification including three indicator variables for treatment assignment; the results indicate that the increase in click-through rates is observed only among individuals who received the basic message around the benefits of FSAs. Of the 59 unique individuals who clicked on a message, 58 of them received the basic message, and 1 received the message highlighting the relative savings from a FSA; there were no click-throughs among those who received the message highlighting absolute savings.

We then estimate the following equation in a two-stage least squares framework, where the independent variable $click_i$ is instrumented by a pooled treatment indicator.

$$contribution_i = \gamma + \beta\ click_i + \varepsilon_i \qquad (2)$$

The results of estimating this specification are reported in Table 6, and again the coefficients of interest are uniformly insignificant and noisily estimated. There is no evidence that individuals who accessed further information were more likely to contribute to a FSA.

5 Discussion

Our pre-specified analysis yielded three main facts. First, approximately 25% of employees contributed to a FSA and the mean contribution in the full sample was approximately $450. Second, communicating FSA eligibility and basic information about the accounts did not increase utilization of FSAs or the amount contributed. Third, highlighting the absolute savings or relative savings similarly had no additional effect on increasing contributions.

Available evidence suggests email open rates of approximately 50% for the first email campaign and approximately 50% for the second email campaign. Each recipient received our messaging twice, during the most read day of the tri-weekly newsletter, and our messaging was one of only four items highlighted in the main section of each newsletter. All of this suggests that recipients were aware that they

[7] In contrast, email open rates are quite high, around 50%. However, individual-level data on open rates are not available.

had received information about potential savings associated with using FSAs, and that the null effect does not simply reflect low levels of information penetration.

Contextual details and the existing literature suggest several possible hypotheses about why our messages did not change behavior. First, although recipients may have researched FSAs because of our emails, they may have concluded that low and/or uncertain medical expenses meant that the net benefit of participation was low, as suggested in Bhattacharya et al. (2002). Individuals who were not already contributing to a FSA may have found that it was not beneficial to participate. The fact that our messaging did not increase participation among previous non-contributors (e.g., see Table 2, Panel B) is consistent with this hypothesis.

Second, individuals may have received information about FSAs from other sources, either prior to or during our experiment. Employees already enrolled in a FSA may receive additional information from the FSA vendor (as distinct from the human resources department). There are also other sources of information about human resources benefits at the agency. Evidence from other settings suggests that more frequent reminders may not be more effective than less frequent reminders (e.g., Pop-Eleches et al. 2011, Damgaard and Gravert 2018); accordingly, in the context of widespread communication from GSA HR and the third party contractor, including in the brief window leading up to and during Federal open enrollment period, our additional communication may not have increased information or salience.

Alternatively, our results may reflect in part contamination bias, if employees in the treatment and control arms directly shared information. We considered a cluster randomized design, yet the only feasible level at which to cluster (i.e. employee federal region) included only 11 clusters and still risked contamination given that employees often collaborate across regions.[8]

[8] The emails may have stimulated attention to FSAs and recipients procrastinated on plans to "do it tomorrow". Evidence on procrastination in health (e.g., DellaVigna and Malmendier (2006) highlights this possibility, yet the low click through rate suggests any immediate attention to new FSA information was limited.

Power calculations conducted as part of our pre-analysis plan indicated minimum detectable effects (MDEs) for contribution likelihood of approximately 2.5 to 3 percentage points.[9] MDEs for contribution amount were approximately $30. Virtually all GSA employees were included in the study, with the goal of maximizing statistical power. Given that the nudge intervention was virtually free of cost, even very small effects could be of interest from a policy perspective, and our evaluation was not powered to detect these very small effects. However, the size of this evaluation was large enough to generate relatively precisely estimated zero coefficients for the coefficients of interest.

Our exploratory analysis (i.e. analysis not included in our pre-analysis plan) revealed one striking fact. Among the 59 individuals who clicked through on the embedded link, 58 were in the basic information study arm. Thus, framing savings as "save money" generated greater interest than framing savings in relative (i.e. percent terms) or in absolute (i.e. dollar terms). One explanation for this finding is that message recipients were already aware of an average level of savings, yet the "save money" framing suggested that it may be more than this amount because it did not state the maximum savings. An alternative explanation – one we find more persuasive yet cannot adjudicate empirically – is that the coarse framing of "save money" was more salient than the detailed framing statements.

6 Conclusion

Many policymakers view health savings devices as an important policy tool to be deployed as part of a broader agenda of health care reform (e.g., see Lieber 2017). Yet take-up among eligible employees in the United States appears to be low, with estimates of participation in Flexible Spending Accounts (FSAs), the most common health-savings device ranging from 10% to as high as 40% in some sub-groups (Mercer 2008). Our results indicate that email-based communication at the start of open enrollment period framing the return on savings in FSAs did not increase the likelihood of participation or contribution levels. We explore several explanations for this null result. Our inference is that individuals may already be informed

[9] Our pre-analysis power calculations assumed a baseline contribution likelihood of 23% (or contribution amount of $250 and standard deviation of $500), power of 80%, significance of 0.05, and two-sided hypothesis tests.

and do not appear to exhibit behavioral biases in health savings calculations. Future research examining the effects of changes in incentives to participate in medical savings accounts could reveal whether the magnitude of the financial benefits affects demand for these accounts. Framing the return on savings in medical saving accounts may not be sufficient to change health savings behavior.

References

Bhargava, Saurabh and Dayanand Manoli. 2015. “Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment." American Economic Review, 105(11): 3489-3529.

Bhattacharya, Jayanta, Michael Schoenbaum, and Neeraj Sood. 2002. “Optimal contributions to flexible spending accounts.” Economics Letters 76(1):129-135.

Cardon, James H. 2010. "Flexible spending accounts and adverse selection." Journal of Risk and Insurance, 77(1): 145-153.

Cardon, James H., Joel Moore, and Mark H. Showalter. 2012. "The distributional effects of health reform limits on flexible spending accounts." Forum for Health Economics & Policy, 15(2).

Cardon, James H., and Mark H. Showalter. 2007. "Insurance choice and tax-preferred health savings accounts." Journal of Health Economics, 26(2): 373-399.

Chetty, Raj, and Emmanuel Saez. 2013. "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients." American Economic Journal: Applied Economics, 5 (1): 1-31.

Chia, Ngee-Choon, and Albert KC Tsui. 2005. "Medical savings accounts in Singapore: how much is adequate?" Journal of Health Economics, 24(5): 855-875.

Damgaard, Mette Trier, and Christina Gravert. 2018. "The hidden costs of nudging: Experimental evidence from reminders in fundraising." Journal of Public Economics, 157: 15-26.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. “Paying not to go to the gym.” The American Economic Review, 694-719.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. 2012. "The Oregon health insurance experiment: evidence from the first year." The Quarterly Journal of Economics, 127(3): 1057-1106.

Hamilton, Barton H., and James Marton. 2008. “Employee choice of flexible spending account participation and health plan." Health Economics, 17(7): 793-813.

Hurley, Jeremiah, G. Emmanuel Guindon, Vicki Rynard, and Steve Morgan. 2008. "Publicly funded medical savings accounts: expenditure and distributional impacts in Ontario, Canada." Health Economics, 17(10): 1129-1151.

Jack, William, Arik Levinson, and Sjamsu Rahardja. 2006. "Employee cost-sharing and the welfare effects of flexible spending accounts." Journal of Public Economics, 90(12): 2285-2301.

Kaiser Family Foundation. 2016. Employer Health Benefits 2016 Annual Survey.

Lieber, Ron. 2017. “Republicans Want You in a Health Savings Account. So Now What?” The New York Times, February 24, 2017.

Mercer. 2008. “National Survey of Employer-Sponsored Health Plans 2008.” Mercer Incorporated, New York, New York. March 2009.

Milkman, Katherine L., John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian. 2011. "Using implementation intentions prompts to enhance influenza vaccination rates." Proceedings of the National Academy of Sciences, 108(26): 10415-10420.

Pauly, Mark V., and Bradley J. Herring. 2000. "An efficient employer strategy for dealing with adverse selection in multiple-plan offerings: an MSA example." Journal of Health Economics, 19(4): 513-528.

Peter, Richard, Sebastian Soika, and Petra Steinorth. 2016. "Health insurance, health savings accounts and healthcare utilization." Health Economics, 25(3): 357-371.

Pop-Eleches C, Thirumurthy H, Habyarmina J, Graff Zivin J, Goldstein M, DeWalque D, MacKeen L, Haberer J, Sidle J, Ngare D, Bangsberg D. 2011. "Mobile phone technologies improve adherence to antiretroviral treatment in resource-limited settings: a randomized controlled trial of text message reminders." AIDS, 25(6): 825-834.

Sacarny, Adam, David Yokum, Amy Finkelstein, and Shantanu Agrawal. 2016. "Medicare letters to curb overprescribing of controlled substances had no detectable effect on providers." Health Affairs, 35(3): 471-479.

Schmitz, Hendrik, and Nicolas R. Ziebarth. 2016. "Does price framing affect the consumer price sensitivity of health plan choice?" Journal of Human Resources, 52(1): 88-127.

Steinorth, Petra. 2011. "Impact of health savings accounts on precautionary savings, demand for health insurance and prevention effort." Journal of Health Economics, 30(2): 458-465.

Ye, Jinqi. 2015. "The effect of Health Savings Accounts on group health insurance coverage." Journal of Health Economics, 44: 238-254.

Zabinski, Daniel, Thomas M. Selden, John F. Moeller, and Jessica S. Banthin. 1999. "Medical savings accounts: microsimulation results from a model with adverse selection." Journal of Health Economics 18(2): 195-218.

Table 1: Summary Statistics

	Mean	Standard deviation
	(1)	(2)
Contribution amount ($)	445.87	876.64
Contribution > $0	0.251	0.443
Contribution > $500	0.224	0.417
Observations	11,192	

Notes: "Contribution > $0" and "Contribution > $500" are indicator variables equal to 1 if true and 0 if false.

Table 2: Pooled Effect of Messaging on FY18 FSA Participation

Dependent variable:	Contribution ($)	Contribution>$0	Contribution>$500
	(1)	(2)	(3)
Panel A: Full Sample			
Any messaging	4.046	-0.006	-0.008
	(18.952)	(0.010)	(0.009)
Observations	11,192	11,192	11,192
Panel B: FY17 Non-Contributors			
Any messaging	-6.172	-0.006	-0.007
	(10.183)	(0.006)	(0.005)
Observations	8,425	8,425	8,425
Panel C: FY17 Contributors			
Any messaging	36.572	-0.008	-0.01
	(43.484)	(0.016)	(0.018)
Observations	2,767	2,767	2,767

Notes: "Any messaging" is an indicator variable equal to 1 if the employee was assigned to receive any study messaging. "Contribution > $0" and "Contribution > $500" are indicator variables equal to 1 if true and 0 if false. FY17 Non-Contributors are individuals who did not contribute to a GSA FSA in the previous year.

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 3: Effects of Specific Messages on FY18 FSA Participation

Dependent variable:	Contribution ($)	Contribution>$0	Contribution>$500
	(1)	(2)	(3)
Panel A: Full Sample			
Basic message	-3.776	-0.012	-0.013
	(23.301)	(0.012)	(0.011)
Absolute savings	4.451	-0.005	-0.008
	(23.333)	(0.012)	(0.011)
Relative savings	11.46	-0.002	-0.002
	(23.347)	(0.012)	(0.011)
Observations	11,192	11,192	11,192
Panel B: FY17 Non-Contributors			
Basic message	-8.250	-0.007	-0.010
	(12.394)	(0.007)	(0.006)
Absolute savings	4.329	0.000	-0.003
	(12.886)	(0.007)	(0.007)
Relative savings	-14.587	-0.01	-0.009
	(12.035)	(0.007)	(0.006)
Observations	8,425	8,425	8,425
Panel C: FY17 Contributors			
Basic message	45.336	-0.010	-0.007
	(53.987)	(0.019)	(0.023)
Absolute savings	-10.019	-0.027	-0.034
	(53.878)	(0.020)	(0.023)
Relative savings	74.767	0.014	0.010
	(52.66)	(0.018)	(0.022)
Observations	2,767	2,767	2,767

Notes: "Basic message" is an indicator variable equal to 1 if the employee was assigned to receive the basic information about FSAs. "Absolute savings" is an indicator variable equal to 1 if the employee was assigned to receive messaging highlighting the absolute savings on FSA contributions. "Relative savings" is an indicator variable equal to 1 if the employee was assigned to receive messaging highlighting the relative return on FSA contributions. "Contribution > $0" and "Contribution > $500" are indicator variables equal to 1 if true and 0 if false. FY17 Non-Contributors are individuals who did not contribute to a GSA FSA in the previous year.

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 4: Other Heterogeneous Effects of Pooled Messaging on FY18 FSA Participation

Dependent variable:	Contribution ($)	Contribution>$0	Contribution>$500
	(1)	(2)	(3)
Panel A: Geographic heterogeneity			
Any messaging	16.64	0.002	0.001
	(23.250)	(0.012)	(0.011)
Capital region	14.755	0.009	0.011
	(34.667)	(0.017)	(0.017)
Any messaging * Capital region	-37.116	-0.025	-0.027
	(40.139)	(0.020)	(0.019)
Observations	11,192	11,192	11,192
Panel B: Experience heterogeneity			
Any messaging	-8.654	-0.014	-0.015
	(27.754)	(0.014)	(0.014)
Tenure	2.832*	0.000	0.001
	(1.569)	(0.001)	(0.001)
Any messaging * Tenure	0.93	0.001	0.001
	(1.811)	(0.001)	(0.001)
Observations	11,192	11,192	11,192
Panel C: Experience heterogeneity (alternative specification)			
Any messaging	-3.101	-0.011	-0.011
	(24.931)	(0.013)	(0.012)
Tenure>10 years	99.958***	0.026	0.038**
	(32.683)	(0.016)	(0.016)
Any messaging * (Tenure>10 years)	12.981	0.009	0.006
	(37.905)	(0.019)	(0.018)
Observations	11,192	11,192	11,192

Notes: "Any messaging" is an indicator variable equal to 1 if the employee was assigned to receive any study messaging. "Capital region" and "Tenure>10 years" are indicator variables equal to 1 if true and 0 if false. "Contribution > $0" and "Contribution > $500" are indicator variables equal to 1 if true and 0 if false. FY17 Non-Contributors are individuals who did not contribute to a GSA FSA in the previous year.

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 5: Effects of Messaging on Click-Throughs

Dependent variable:	Click-through	
	(1)	(2)
Any messaging	0.007***	
	(0.001)	
Basic message		0.021***
		(0.003)
Absolute savings		0.000
		(0.000)
Relative savings		0.000
		(0.000)
Observations	11,192	11,192

Notes: "Click-through" is an indicator variable equal to 1 if the individual clicked-through the link embedded in the email messaging and 0 if they did not. "Any messaging" is an indicator variable equal to 1 if the employee was assigned to receive any study messaging. "Basic message" is an indicator variable equal to 1 if the employee was assigned to receive the basic information about FSAs. "Absolute savings" is an indicator variable equal to 1 if the employee was assigned to receive messaging highlighting the absolute savings on FSA contributions. "Relative savings" is an indicator variable equal to 1 if the employee was assigned to receive messaging highlighting the relative return on FSA contributions.

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Table 6: Two-Stage Least Squares (2SLS) Effects of Exposure to Embedded Link

Dependent variable:	Contribution ($)	Contribution>$0	Contribution>$500
	(1)	(2)	(3)
Panel A: Full Sample			
Any messaging	575.458	-0.914	-1.157
	(2694.69)	(1.359)	(1.314)
Observations	11,192	11,192	11,192
Panel B: FY17 Non-Contributors			
Any messaging	-1146.7	-1.076	-1.343
	(1903.69)	(1.103)	(1.027)
Observations	8,425	8,425	8,425
Panel C: FY17 Contributors			
Any messaging	3034	-0.629	-0.844
	(3650.34)	(1.294)	(1.528)
Observations	2,767	2,767	2,767

Notes: Regression results come from a 2SLS procedure where assignment to receive any message is an instrument for click-throughs. "Any messaging" is an indicator variable equal to 1 if the employee was assigned to receive any study messaging. "Contribution > $0" and "Contribution > $500" are indicator variables equal to 1 if true and 0 if false. FY17 Non-Contributors are individuals who did not contribute to a GSA FSA in the previous year.

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level.

Figure 1: Distribution of FY18 Contributions ($) Conditional on Contribution>$0

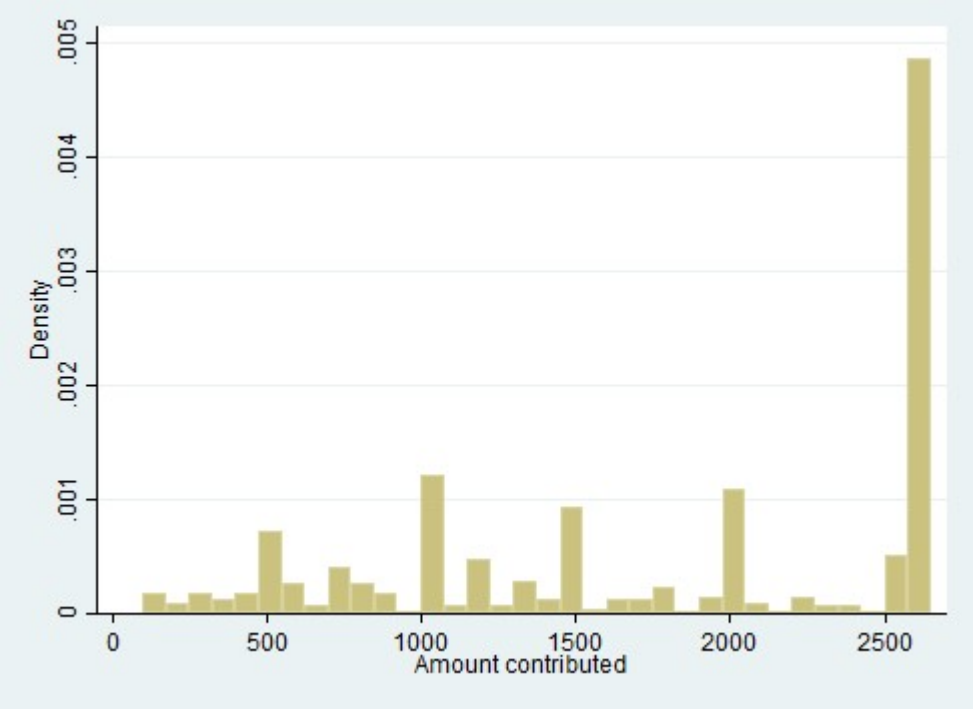

Figure 2a: Contribution Amount by Study Arm

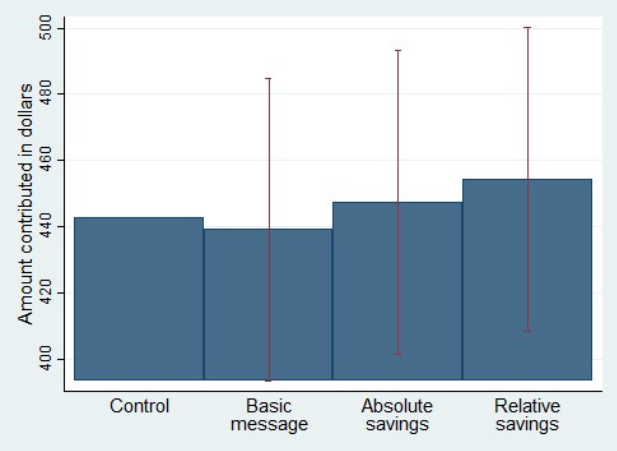

Figure 2b: Probability of Contribution>$0

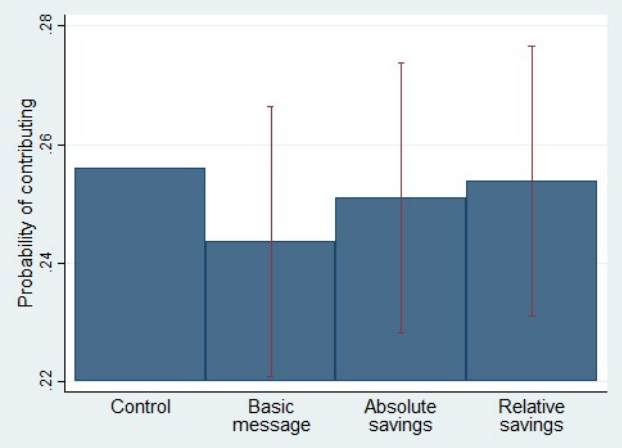

Figure 2c: Probability of Contribution>$500

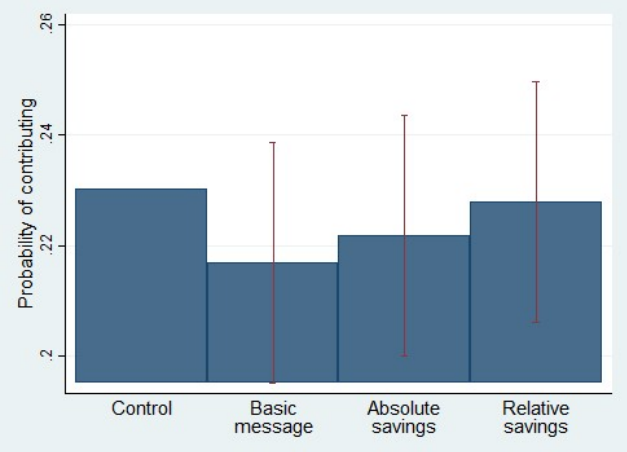

Notes: Whiskers are 95% confidence intervals.