

Intervention Externalities due to Limited Attention *

Hannah Trachtman

Yale University

October 9, 2019

Abstract

The potential welfare benefits of motivating people to vaccinate their children, consume healthy foods, or use clean cookstoves are enormous. Recent research has uncovered many interventions that cost-effectively improve such behaviors, as well as many that do not. But most research evaluates one intervention in isolation on target outcomes. As such, we have little understanding of how interventions might interact with one another, or whether they generate spillovers to other behaviors. This paper explores the hypothesis that behavior change interventions might generate negative externalities due to limited attention. I propose a simple framework, focusing on three types of limited attention that have distinct policy implications. I test the predictions of the model using an online experiment in which individuals receive combinations of messages and incentives for two healthy behaviors, meditation and meal tracking, which are measured daily via phone applications. I find that messaging and incentive interventions generate negative spillovers of 2.8 and 2.4 percentage points on base rates of 9.4 and 11.8 for meditation and meal tracking, respectively. Estimating the parameters of the model reveals that effective interventions do not necessarily generate larger negative spillovers than ineffective interventions, all else being equal. Specifically, suppose a low-effectiveness intervention (0.2 SDs) is scaled so that, in the absence of spillovers, it is equally cost-effective to a high-effectiveness intervention (1 SD). In the presence of spillovers driven by limited attention of the observed type, the former intervention is predicted to cost 28% more than the latter. Thus, for policymakers who care about multiple outcomes, small-scale, highly-effective interventions may be preferable to large-scale, less effective ones, once spillovers are taken into account.

*I am so grateful to my dissertation committee for their help at every stage of this project: Dean Karlan, Chris Udry, Costas Meghir, and Jason Abaluck. For helpful discussions and comments, I thank Joseph Altonji, Gaurav Chiplunkar, Taha Choukhmane, Marvin Chun, Ori Heffetz, John Eric Humphries, Ro'ee Levy, Cormac McCarthy, Mushfiq Mobarak, Mark Rosenzweig, and Jaya Wen. I would like to thank the Economic Growth Center and Sylff Foundation, the Russell Sage Foundation, the Yale Economics Department, and the National Science Foundation for their generous support. This research was conducted under full IRB approval. All errors are my own.

1 Introduction

Seemingly small changes in behavior can lead to big economic benefits. A healthy diet reduces the risk of heart disease, diabetes, and cancer: in 2017, 11 million deaths globally were linked to dietary risk factors, amounting to 15% of the global disease burden ([Afshin et al., 2019](#)). Routine childhood immunization in the United States has been estimated to prevent 42,000 early deaths in every birth cohort ([Zhou et al., 2014](#)). The use of clean, non-traditional cookstoves drastically reduces rates of respiratory and cardiovascular disease: in 2017, household air pollution from solid fuels was estimated to be responsible for 2.4% of the global disease burden ([Institute for Health Metrics and Evaluation, 2018](#)).

Given the enormous returns to behaviors like these, low take-up is seen as a puzzle. Many traditional instruments like “sin taxes” ([O’Donoghue and Rabin, 2006](#)), as well as non-traditional instruments like “nudges” ([Thaler and Sunstein, 2009](#)) have been proposed. Many of these interventions have shown to be effective and cost-effective with respect to their target outcomes. Soda taxes have shown to reduce sugar consumption due to soda by 18% on average, and by 40% among young people age 13-21 ([Dubois et al., 2019](#)). Prompting people to write down a plan for getting the flu vaccine raised immunization rates by 13% ([Milkman et al., 2011](#)). Other interventions have proven less effective: encouraging the use of clean cookstoves has proved to be much more difficult than expected ([Hanna et al., 2016](#)). But the possibility of shifting behaviors with such high stakes has inspired an outpouring of interventions and evaluations. In particular, the use of nudges—which are often cheap, unimposing, and relatively easy to implement—has grown dramatically, in developed and developing countries, across many domains, and implemented by public and private actors alike.

At the same time, studies in neuroscience and psychology have produced an enormous amount of evidence that people have limited attention. We have evidence that people are only able to process finite amounts of information at once ([Duncan et al., 1997](#)). We have evidence that when something captures the brain’s attention, other things are set aside ([Yantis and Jonides, 1984](#)). And we have evidence that certain tasks are mentally or cognitively effortful or exhausting ([Shenhav et al., 2017](#)). We also have evidence that these manifestations of limited attention can have impor-

tant consequences for economic outcomes. For example, used car consumers do not fully attend to the right-most digits of mileage numbers ([Lacetera et al., 2012](#)), grocery shoppers do not fully account for sales taxes unless they are made salient ([Chetty et al., 2009](#)), and investors do not fully respond to earnings announcements that occur right before the weekend ([DellaVigna and Pollet, 2009](#)).

But we know less about how limited attention mediates interventions themselves, and in particular, about the ways in which limited attention might cause interventions to impose negative externalities on other interventions or behaviors. If people have only a limited capacity for information processing, interventions that rely on information provision or high-frequency stimuli might interfere with other similar interventions. If people have limited space at the “top of mind,” or if mental effort is costly, interventions that promote one behavior might generate negative spillovers on other behaviors. A paper by [Medina \(2017\)](#) provides suggestive evidence that interventions can indeed impose negative externalities on other behaviors via limits to attention. She finds that sending SMS reminders to pay credit card bills does effectively reduce late fees paid by 11%, but it also increases overdraft fees paid by 11%, resulting in a net loss for about 20% of her sample. The spillover is unsurprising given that both behaviors (paying bills and not going into overdraft) draw from the same budget, but the fact that some participants were made worse off by the intervention points to some sort of attentional failure.

This paper aims to answer three questions. First, can limited attention cause behavior change interventions to impose negative spillovers on one another? Second, to what extent are these spillovers caused by different types of limited attention? Third, what are the policy implications?

I begin by reviewing the psychology literature on limited attention, with the goal of identifying types of limited attention that will have important and distinct policy implications. I arrive at a taxonomy with three types of limited attention, which I call “overload,” “depletion,” and “diversion.” Overload captures limits to information processing, which might result in interventions interfering negatively with one another, by overwhelming the individual with stimuli. Depletion captures the idea of costly mental effort; the idea that engaging in certain behaviors depletes cognitive resources. Importantly, with depletion, an intervention’s target effects will determine its

spillovers, since it is the act of doing the target behavior that depletes resources. Finally, diversion captures limits to executive function, short-term, or working memory that prevent people from keeping more than one behavior at the top of mind. A key characteristic of diversion is that, unlike depletion, an intervention's target effects do not necessarily determine its spillovers. Interventions may divert attention away from other behaviors regardless of the extent to which they improve the target behavior. Thus, spillovers driven by diversion act as "fixed costs" of the intervention.

Potential negative externalities imposed by interventions will have implications that depend importantly on the type of limited attention that drives them. First, if externalities are driven by overload, we can reduce them by shifting toward interventions that place fewer demands on information processing. For example, an educational campaign about the negative effects of sugar consumption requires a great deal of information processing, while a soda tax does not. Second, if externalities are driven by diversion, we can potentially raise welfare by shifting toward high-cost, high-effectiveness interventions, and spreading the fixed spillover costs over a larger target effect.

I incorporate these ideas into a simple framework, which will motivate the experiment design. A decision-maker (DM) has two behaviors available to her, x and y . Doing each behavior generates a return, but requires attention, which is costly. A benevolent social planner can subsidize attention to a particular behavior with incentives or SMS messages. The model's comparative statics with respect to messages and incentives generate two non-parametric predictions that can be tested using the reduced form results from the experiment. With additional structure, the model can be estimated to identify parameters that link to all three types of limited attention.

The experiment design has five treatment groups: a control group, a group that gets messages about behavior x , a group that gets messages about behavior y , a group that gets both sets of messages, and a group that gets incentives for behavior y . I recruit 3,845 individuals via Facebook Ads that promote a study about daily meditation (behavior x) and nutritional monitoring (behavior y). Participants take a baseline survey and download two smartphone apps, one for tracking meditation, and the other for logging meals. Upon verifying that they downloaded both apps,

they are enrolled, randomized, and informed of their treatment assignment via email. Treatments begin immediately thereafter, lasting four weeks, and I continue to measure behavior for an additional four weeks after the end of treatment. At the end of the follow-up period, participants are informed of the end of the study and sent a final survey.

The reduced form results show large target effects of both message and incentive interventions. Meditation messages raised the rate of meditation by 8.8 percentage points (almost doubling the rate) and nutrition messages raised the rates of meal logging by 16.6 percentage points (more than doubling the rate). Incentives for meal logging had an even larger effect, raising rates of meal logging by 38.1 percentage points (more than tripling the rate). But all three treatments imposed substantial externalities on the opposite behavior, as measured by comparisons with the control group. Messages about meditation reduced meal logging by 2.4 percentage points (19%), messages about nutrition reduced meditation by 2.8 percentage points (29%), and incentives for meal logging reduced meditation by 2.5 percentage points (27%). The group with both sets of messages also did worse than the group with just meditation or just meal-logging messages, by 2.2 and 5.0 percentage points, respectively. There is no evidence of an interaction effect between the two sets of messages, suggesting that there is no evidence of overload in this context.

Two facts embedded in the reduced form results foreshadow the importance of diversion as a key mechanism. First, the fact that all three interventions generate similarly sized spillovers suggests that depletion alone cannot drive the results. If it did, the more effective interventions should have depleted more resources and generated larger spillovers. Second, we see that the covariance between meditation and meal logging does not vary across treatments. If depletion were driving the results, any intervention that successfully promoted the target behavior should also cause depletion and reduce the covariance between the actions. I confirm these conjectures by estimating the model, using simulated minimum distance to estimate 15 parameters with 22 moments. In my benchmark specification, both overload and depletion are not statistically different from zero. The estimates of both message and incentive diversion are negative, but only the former is statistically significant. These estimates suggest that diversion is a key driver of spillovers.

Other sources of data provide additional support for the key role of diversion. Evidence from

a raffle sent by SMS to all participants receiving messages suggests that participants assigned to both sets of messages were less likely to read them, relative to participants assigned to just one set. A simple depletion story cannot explain why spillovers should appear already at the reading stage; they should only appear at the action stage. A more complex depletion story, where participants anticipate depletion and decide in advance not to read meditation messages (or subsequently meditate), could explain this result. But data from a survey administered immediately after treatment assignment suggests that participants assigned to meal logging messages did not expect spillovers on meditation, nor were they more likely to opt-out of meditation messages. We thus have no evidence that participants made a decision, neither upon treatment assignment nor during treatment, to focus on one behavior and not on the other. These data are therefore consistent with the finding that depletion does not play a large role.

The fact that I found no evidence of overload suggests that we have no reason to worry about high stimulus interventions. The finding that diversion is a key driver of spillovers, however, has important implications for the cost-effectiveness of different types of interventions. Interventions can divert attention even when they do not positively affect the target behaviors, implying that the resulting spillovers act as “fixed costs.” A natural implication of this is that two equally cost-effective interventions will not be equivalent once spillovers are taken into account. Specifically, suppose we scale a cheap, low-effectiveness intervention (0.2 SDs) so that it is equally “locally” cost-effective to a high-effectiveness intervention (1 SD). In the presence of diversion, the former intervention is predicted to cost 28% more than the latter. Thus, policymakers who care about multiple outcomes should not be indifferent between equally “locally” cost-effective interventions: small-scale, highly-effective interventions may be preferable to large-scale, less effective ones, once spillovers are taken into account.

I begin by describing the theory and evidence for limits to attention in Section 2. In Section 3, I propose a simple framework that formalizes these notions, and generates predictions for behavior with respect to incentives and messages, which I will vary experimentally. In Section 4, I describe the experiment. In Section 5 I present orthogonality tests, descriptive statistics, and reduced form results. In Section 6 I estimate the structural model and test the model fit. In Section 7 I show

additional evidence for the proposed mechanisms, and examine alternative explanations for the findings. In Section 8 I explore the policy implications, and in Section 9 I conclude.

2 Limited Attention

Perhaps the most widely used taxonomy of limited attention goes back as far as 1890, when William James distinguished between “passive” and “active” attention. This basic division has persisted over the years, with several variants—bottom-up versus top-down attention, stimulus-driven versus goal-driven attention, exogenous versus endogenous attention, and most recently, external versus internal attention, the categorization used in a review paper by [Chun et al. \(2011\)](#). External attention refers to the selection and modulation of external information, and the storing of that information in the brain. It can be directed to the sensory modalities of sight, hearing, touch, smell, and taste, and it can be used to perceive the world across space (“spatial attention”) or time (“temporal attention”). Internal attention, on the other hand, refers to the selection and modulation of content that has already been stored in the brain. It includes the attention required to think about, plan, and make decisions about an action—including executive function, working memory, and long-term memory. It also includes the attention required to carry out a task—including things like cognitive effort and self-control.

I will build on this taxonomy to describe the possible ways in which limited attention could cause interventions to impose negative externalities, summarized in Figure 1. Suppose there is an intervention, say a text message, about some behavior x . We use our external attention to modulate that stimulus and store it in our brains. We then use internal attention to occupy ourselves with x , and ultimately do x . Now suppose there is also a text message about some different behavior y . Three things might happen. First, limits to external attention, or limited information processing, might cause the y stimulus to interfere with the x stimulus. Henceforth I will call this possibility “overload.” Second, limits to working or short-term memory might cause the y stimulus to divert attention toward y and away from x , reducing the likelihood of doing x (regardless of whether or not there is any x stimulus). I will call this possibility “diversion.” Finally, if the y stimulus works, causing us to do y and to exert costly cognitive effort, we might be subsequently

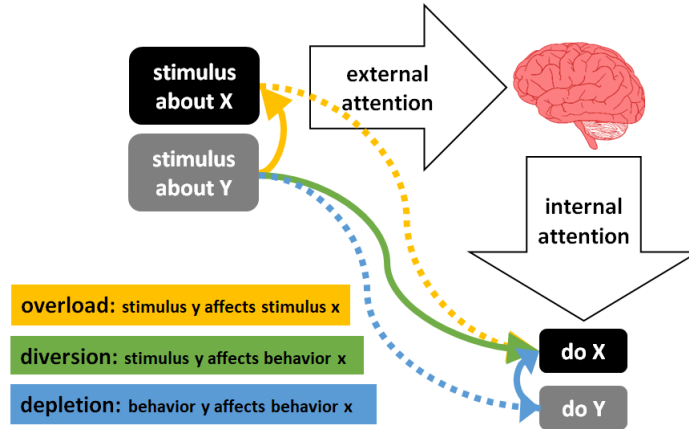


Figure 1: Taxonomy of Limited Attention

less likely to do x . I will call this possibility “depletion.” In short, overload can be summarized as “stimulus y affects stimulus x ,” diversion can be summarized as “stimulus y affects behavior x ,” and depletion can be summarized as “behavior y affects behavior x .”

What evidence do we have of the possibility of overload? The fact that people are limited in their ability to process stimuli is so well-established that words like “selection” are commonplace in the literature; the relevant question is not whether we select which stimuli to process but how. The literature on this question is vast. One seminal paper showed that limits to external attention are modality-specific: people are unable to attend to two visual or two auditory streams, but better able to attend to one of each (Duncan et al., 1997). The “information overload” literature in management has documented a hump-shaped relationship between the quantity of information that consumers have about products, and the “decision quality” of their ultimate purchase (Hwang and Lin, 1999; Edmunds and Morris, 2000; Eppler and Mengis, 2004). A newer set of studies has examined habituation or desensitization to alerts over time. In medicine, as the use of electronic medical records and attendant automatic alerts to provide “decision support” have become widespread, there has been extensive discussion of “alert fatigue,” the idea that physicians become habituated to alerts over time. One SMS program designed to alert physicians to new clinical trials found that response rates declined 2.7% every two weeks (Embi and Leonard, 2012).

What evidence do we have of the possibility of diversion? Experiments on working memory

show that people are capable of holding only limited sets of digits or words in their heads (Miller, 1956; Luck and Vogel, 1997), implying that focusing on one thing often comes at the expense of something else. Studies on “attentional capture,” show that irrelevant stimuli can easily draw people’s attention away from a task at hand (Yantis and Jonides, 1984). The types of stimuli most likely to achieve attentional capture are novel stimuli (i.e. an unexpected SMS), emotionally salient stimuli (i.e. footage of a humanitarian crisis) and stimuli associated with rewards (i.e. a plate of cookies placed in front of you) (Fawcett et al., 2015; Chun et al., 2011). Our internal attention can be “captured” by external stimuli from the bottom-up, but it can also be re-directed from the top-down, deliberately or inadvertently, without any external stimuli. One important example is the phenomenon of “intention cost,” or reduced performance (and brain activity) in a current task as a result of thinking ahead to a future task (Burgess et al., 2003; Gonen-Yaacovi and Burgess, 2012).

Lastly, what evidence do we have of the possibility of depletion? Studies have shown that certain tasks are cognitively costly. For example, performance on difficult tasks tends to increase with incentives (Botvinick and Braver, 2015), and when given a choice between tasks that require high and low cognitive effort, participants tend to prefer the latter (Dunn et al., 2016). Neuroscientists are currently exploring the underpinnings of these costs, and the types of tasks that incur them. One insight they’ve made is that the more automatic the task—the closer it is to some “default” behavior—the less effort it requires (Shenhav et al., 2017). This is relevant because the policy goal of “behavior change” inherently asks people to move away from their defaults, and thus may also inherently require cognitive effort. Another reason tasks may be costly is because they require self-control. I will call limits to this type of internal attention “depletion.”

There are many economic implications of overload, which may cause agents to have incomplete information on products, prices, tax rates, news, and more. Sims’ (2003) model of “rational inattention” was one of the first to explore information processing constraints in economics, focusing on macroeconomics and monetary policy (Sims (2003)). Since then, theorists have proposed several models of limited information processing and its consequences (Falkinger, 2008; Eliaz and Spiegel, 2011; Masatlioglu et al., 2012; Manzini and Mariotti, 2012; Bordalo et al., 2012, 2013; Gabaix, 2014; Schwartzstein, 2014; De Clippel et al., 2014). Empirically, we have evidence

that people have difficulty processing all of the information about the products they buy (Lacetera et al., 2012), all of the dimensions of their production processes (Hanna et al., 2014), or all of the choices available to them (Chernev et al., 2015).

What are the economic implications of diversion? One important example concerns taxation. Evidence suggests that people are not fully responsive to sales taxes that are not salient (Chetty et al., 2009), and subsequent papers have drawn out the broader implications for optimal tax policy (Taubinsky and Rees-Jones, 2017; Farhi and Gabaix, 2018). Another example concerns behavior change: the fact that people respond to text messages reminders in a variety of contexts (Karlan et al., 2016; Taubinsky, 2013; Allcott and Rogers, 2014; Rogers and Milkman, 2016) suggests that such prompts indeed help people focus on or remember things they would have otherwise forgotten. These effects are consistent with the idea that well-designed prompts can divert attention from one place to another, mitigating some of the unfortunate consequences of limited attention or memory (Mullainathan, 2002; Holman and Zaidi, 2010).

Several theoretical papers in economics have explored the idea of depletion, typically in the form of costly self-control (Gul and Pesendorfer, 2001, 2004; Noor, 2007; Ozdenoren et al., 2012). A related idea is that of “moral licensing,” the idea that people aim to maintain their positive self image, and thus engaging in something “good” can license one to subsequently engage in something “bad,” or vice versa. (Dolan and Galizzi, 2015)¹ Moral licensing has been shown to have implications for consumption, environmental practices, and political behavior (Werthenbroch, 1998; Strahilevitz and Myers, 1998; Khan and Dhar, 2006). Finally, empirical work in finance (DellaVigna and Pollet, 2009) and development (Mani et al., 2013; Shah et al., 2012) has pointed toward the potential importance of cognitive exhaustion for economic outcomes.

3 Theoretical Framework

I use a simple framework to define each type of limited attention, and then to derive comparative statics of behavior with respect to messages and incentives, which I experimentally vary.

¹Although moral licensing can be modeled in the utility function, as Dolan and Galizzi (2015) do, for the purpose of the study I will include it in the broad category of limited internal attention, since the idea that self-image should be maintained but not maximized may arise from some underlying effort cost to doing so.

The purpose of the framework is twofold. First, it generates two non-parametric predictions that are immediately testable with the reduced form results. Second, with additional structure, it enables me to identify key parameters associated with each type of limited attention, allowing me to decompose spillovers by type of limited attention and to draw corresponding policy conclusions.

3.1 Set-Up

I consider an agent who chooses whether or not to take two actions, $x \in \{0, 1\}$ and $y \in \{0, 1\}$. Taking action $j \in \{x, y\}$ results in benefit u_j . At the start of the period, the agent chooses attention weights a_x and a_y , which determine how much she “thinks about” each action’s payoff. After this she is hit by a negative i.i.d. distraction shock $\xi_j \sim F$, so that net attention is $\tilde{a}_j = a_j + \xi_j$. The agent will perform action j as long as net attention \tilde{a}_j is non-negative; this means we can think of \tilde{a}_j as a latent variable. Finally, the agent faces some cost of attention, which I denote by the function C .

Agents maximize expected utility:

$$\max_{a_x, a_y} \left\{ \Pr(x = 1)u_x + \Pr(y = 1)u_y - C(a_x, a_y) \right\}$$

I will normalize a_x and a_y by assuming that both ξ_x and $\xi_y \sim_{iid} U[-1, 0]$. Then $\Pr(x = 1) = \Pr(a_x + \xi_x > 0) = a_x$. In other words, when the agent chooses her attention weights, she is also choosing the probability that she overcomes the distraction shock and does the action. I can thus rewrite the agent’s problem as:

$$\max_{a_x, a_y \in [0, 1]} \left\{ (a_x u_x + a_y u_y) - C(a_x, a_y) \right\}$$

Attention costs cause the agent to expend less attention than she otherwise would. However, an outside actor can introduce an intervention w_j for behavior j , where the intervention can be either messages m or incentives z ; i.e. $w \in \{m, z\}$. Recall that novel stimuli “capture” attention from the bottom-up, and rewards “capture” attention from the top-down (Chun et al., 2011). I thus model both messages and incentives as attention subsidies. Let the modified cost function be

$C(a_x, a_y, w_x, w_y)$. Let the marginal cost of a_x , C_1 , be denoted as $c^x(a_x, a_y, w_x, w_y)$ and the marginal cost of a_y , C_2 , be denoted as $c^y(a_y, a_x, w_y, w_x)$. For now I make the following assumptions about the cost function. First, I assume that $c_1^x > 0$ and $c_1^y > 0$ to ensure the existence of a local maximum. Second, I assume that both target messages and target incentives subsidize the target behavior, reducing the marginal cost of attention: $c_3^x < 0$, $c_3^y < 0$.²

I define *overload* to be the possibility that $c_{34}^x > 0$ or $c_{34}^y > 0$ in the case of a message intervention ($w = m$). This means that y messages interfere with the attention subsidy produced by x messages (and vice versa). I define *diversion* to be the possibility that $c_4^x > 0$ or $c_4^y > 0$. Diversion thus operates like a tax: messages or incentives about behavior y increase the marginal cost of attention to x . I allow for the possibility of both message and incentive diversion. Lastly, I define *depletion* to be the possibility that $c_2^x = c_2^y > 0$. This means that the marginal cost of attending to behavior x is increasing in the attention paid to behavior y , and vice versa. Since attention is also the probability of doing the action, this captures the idea that doing one behavior depletes internal resources, raising the cost of doing the other. (From now on I will write $c_2^x = c_2^y$ as simply c_2^x .)

3.2 Non-Parametric Predictions

I can then derive comparative statics of attention with respect to messages and incentives, the most important of which will be: $\frac{\partial a_x^*}{\partial m_y}$, $\frac{\partial a_x^*}{\partial z_y}$, and $\frac{\partial^2 a_x^*}{\partial m_x \partial m_y}$ (and their equivalents for a_y^*). I define a *spillover* to be the negative response of a behavior to a non-target intervention. If $\frac{\partial a_x^*}{\partial m_y} < 0$, it constitutes a *message spillover*, and if $\frac{\partial a_x^*}{\partial z_y} < 0$, it constitutes an *incentive spillover*. I define *interference* to be a negative interaction between two interventions. If $\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} < 0$, it constitutes *interference*, and if both $\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} < 0$ and $\frac{\partial^2 a_y^*}{\partial m_x \partial m_y} < 0$, it constitutes *interference in both directions*.

I focus on interior solutions, deriving comparative statics under the assumption that a_x^* and $a_y^* \in (0, 1)$. In the case of corner solutions, locally, comparative statics will be zero with probability 1. I obtain two key predictions.

Proposition 1. *Either message (incentive) diversion or depletion is a necessary condition for message*

²I assume that the outside actor will not implement both messages and incentives. The types of limited attention that I focus on do not have interesting implications for interactions between messages and incentives, so I do not implement this treatment in my experiment.

(incentive) spillovers, and the presence of both message (incentive) diversion and depletion is a sufficient condition for message (incentive) spillovers.

Proof. The proof is straightforward from the expression for $\frac{\partial a_x^*}{\partial w_y}$, Equation 1 in Appendix A.1. \square

Proposition 1 implies that if I find message or incentive spillovers, they must be due to some type of limited internal attention. If I find neither, then the conclusion is ambiguous.

The second prediction requires two additional assumptions. First, I assume that, with the exceptions of c_{34}^x and c_{34}^y , the second derivatives of c^x and c^y are zero.³ Second, I assume that $c_{34}^x = c_{34}^y$; namely, that the overload effect is symmetric across different behaviors. This implies that if two messages are sent, one about x and one about y , the first will interfere with subsidy generated by the second on y just as much as the second interferes with the subsidy generated by the first on x .⁴

Proposition 2. *Assume that $c_{34}^x = c_{34}^y$ and that all other second derivatives of c^x and c^y are zero. Then overload is a necessary condition for interference in both directions, and a sufficient condition for interference in one direction.*

Proof. See Equations 2 and 3 in Appendix A.1 as well as Appendix A.2. \square

The implication of Proposition 2 is as follows. If I find interference in both directions, it must be due to overload. If I find interference in neither direction, then there is no evidence of overload. And if I find interference in just one direction, the conclusion is ambiguous.

4 Experiment Design

The experiment design is displayed in Table 1. The control group received no intervention. Group 2 received only messages about behavior x , and Group 3 received only messages about behavior y . Group 4 received messages about behavior x as well as messages about behavior y . Group 5 received incentives for behavior y .

³I have no reason to believe that these derivatives are zero, but no reason to believe otherwise, as economic intuition tells us nothing about the signs of these derivatives.

⁴Ultimately I can check this assumption by testing whether $\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} = \frac{\partial^2 a_y^*}{\partial m_x \partial m_y}$, and indeed this test is not rejected at the 5% level ($p=0.09$).

Group	Description	Messages		Incentives	N
		$m_x = 1$	$m_y = 1$	$z_y = 1$	
1	ctrl				814
2	x messages	x			763
3	y messages		x		803
4	x & y messages	x	x		822
5	y incentives			x	643
					3845

Table 1: Experiment Design

Behaviors x and y were daily meditation and nutritional self-monitoring. These behaviors were chosen for four reasons. First, they are important health behaviors for the sample frame (young Americans). A recent meta-analysis in the Journal of the American Medical Association found that meditation programs improved anxiety by 0.38 SDs at 8 weeks (and 0.22 at 3-6 months), improved depression by 0.30 SDs at 8 weeks (and 0.23 at 3-6 months), and reduced pain by 0.33 SDs at 8 weeks (Goyal et al., 2014). The use of smartphone apps for nutritional self-monitoring and feedback have been linked to weight loss (Wharton et al., 2014), which is associated with many health benefits.

Second, both behaviors can be measured objectively at high frequency via pre-existing smartphone applications. The meditation application allows users to access a wide variety of guided meditations or meditate on their own, and records details about each meditation session. In the nutritional monitoring application, users to input information about the meals they eat and then track various measures of the nutritional quality of their diet.

Third, both behaviors require minimal amounts of time. Both behaviors can be measured via pre-existing smartphone applications. The average meditation session was 21 minutes, but meal logging only took 11 minutes per day on average. Thus any spillover effects are unlikely to be explained by the time constraint, but I will address this possibility in Section 7.2. Finally, these two behaviors are not obviously related to one another in any utility function, though this will not be an identifying assumption.

The 3845 participants were recruited using Facebook ads targeting adults age 18-35 living in the U.S. (see Appendix 4). Upon clicking the link, participants underwent a brief screening that en-

sured they (1) had an iPhone or android; (2) were over 18; (3) were interested in working on wellness habits like meditation and tracking nutrition; and (4) were willing to download the two free applications. They then provided informed consent and proceeded to Survey 1, which took about 15 minutes. The first part of Survey 1 provided instructions for downloading the two apps (which were also emailed upon survey completion). Participants were instructed that in order to enroll, they would need to download both apps within 24 hours. Participants were then asked questions on demographics, electronic notifications, and preferences/experiences surrounding meditation and nutritional monitoring.

Participants who were verified to have downloaded both apps were then randomized to treatments, re-randomizing on the following variables: gender, age, whether or not they had a college degree, daily notifications, whether or not they meditated in the last month, and whether or not they tracked their meals in the last month. These participants then received an enrollment confirmation email with their treatment assignment, a link to Survey 2, and other details about the study. Survey 2 required about five minutes and contained questions about participants' expectations of each behavior, conditional on their treatment assignment.

Importantly, when informed of their treatment assignment, participants were told that "this assignment was completely random, and has nothing to do with your survey responses or the relative importance of meditation, exercise, nutrition, and sleep." The purpose of this was to rule out an alternative potential source of spillovers or interference: the possibility that participants infer the relative benefits of the behaviors from their treatment assignment.⁵ We also tell participants that, "Depending on your above assignment, we may (or may not) be encouraging you to meditate and/or log your meals, but your ultimate use of the apps is entirely up to you." The purpose of this was to avoid experimenter demand effects, and prevent participants from feeling obligated to engage in behaviors assigned to be treated (perhaps with motives of reciprocity or adherence to some imagined authority).

Each message program included twice-daily text messages: one simple reminder to do the behavior, and one longer message with information about some proven benefits to the behavior,

⁵This mechanism is potentially important, but cannot be well studied in this kind of experimental context, since many participants already assume assignment is random.

as demonstrated in Table 4. Messages were sent at either 7am and 7pm or at 8am and 8pm, alternating on a daily basis, and scheduled so that the group that received both meditation and nutrition messages never received them at exactly the same time (it was always the case that one message was at 7 and the other at 8). The purpose of this was to avoid capturing mechanical interference due to the simultaneous arrival of messages.

Group	Time	Msg 1	Msg 2
med only	8AM	Remember to meditate today! Try the 3-minute breathing space by Mark Williams on Insight Timer!	
	3PM	A meta-analysis in a top medical journal reviewed 47 studies and found systematic evidence that meditation reduces depression and anxiety! (Goyal et al. 2014)	
med & nut	8AM	Remember to meditate today! Try the 3-minute breathing space by Mark Williams on Insight Timer!	Logging meals can help with weight loss (Burke et al. 2011)! And people are better at meal-logging when they use apps like FatSecret (Wharton et al. 2014).
	3PM	A meta-analysis in a top medical journal reviewed 47 studies and found systematic evidence that meditation reduces depression and anxiety! (Goyal et al. 2014)	Take one minute to log your meals using FatSecret today!

Table 2: Example Messages, Day 1

Incentives took the form of a raffle. Participants were informed in their enrollment email that they would earn one green lottery ticket for every day they successfully do the behavior, and one red lottery ticket every day that they don't. At the end of the four weeks, we would draw one of their tickets, and each winning ticket would be worth a \$10 Amazon gift certificate. Every Sunday during the program participants received an email updating them about the tickets earned the previous week. They also received an email informing them when the program ended.

Four weeks after the end of treatment, participants received Survey 3 via email. Survey 3 included questions about meditation and nutritional monitoring outside of the assigned apps, the timing of behaviors, some measures of mental health and diet, and quizzes about the information content of any message program they received. For further details about the experiment protocol see Appendix B. For further details about attrition, see Appendix C.

Table 3 shows means and standard deviations of key variables across treatments, as well as an F-test of the joint significance all treatment variables. The re-randomization procedure ensured

that the first variables were balanced across treatments, and the rest of the variables are highly balanced as well. Overall, the sample was overwhelmingly female (93%), mostly college educated (71%), with an average age of 27. Participants receive on average 51 notifications daily, 36 of which are messages, 10 of which are updates, 4 of which are reminders, and 1 of which was classified as “other.” (See Appendix D, Figure 6 for details.) On average, participants perceived meal logging to be slightly more important than meditation and slightly more difficult, but the main difference in the behaviors is that meditation is perceived to be much more “fun” than meal logging. Most participants had experience with both meditation and meal logging. With respect to meditation, 90% had meditated before, 57% had done so on a daily basis, and 46% had done so in the last month. With respect to meal logging, 90% had logged their meals before, 87% had done so on a daily basis, and 32% had done so in the last month. (See Appendix D, Figure 7 for details.) These are people who have strong prior interest and experience in both behaviors, but who, for whatever reason, have not been engaged in them recently.

5 Reduced Form Results

I estimate linear probability models of the outcome on treatments, where the outcome is 1 if the participant did the behavior on a given day. For meal logging, the behavior is having logged at least one meal and 0 otherwise. I define m_x (m_y) to be 1 if the individual received x (y) messages and 0 otherwise; $m_x * m_y$ is 1 if the individual received both sets of messages. I include a vector of controls that consists of the variables used for re-randomization: whether or not the participant is female, whether or not they completed college, daily notifications, whether or not they meditated in the month prior to the study, and whether or not they logged a meal in the month prior to the study. I also include day fixed effects, and cluster standard errors at the individual level. Coefficients on message treatments represent intent-to-treat effects, as some participants chose to stop receiving messages.

The results are shown in Table 4. (See Appendix E Table 12 for estimates reported as treatment effects, and Figures 8 through 12 for depictions of the raw data). During the treatment period, both sets of messages doubled the rates of their target behaviors: meditation messages raised the

	control	mx	my	mx & my	zy	F-test, joint sig
female	0.93	0.93	0.93	0.93	0.93	1.00
	0.25	0.25	0.25	0.25	0.26	
went to college	0.71	0.72	0.71	0.71	0.72	0.99
	0.45	0.45	0.45	0.45	0.45	
age	27.46	27.43	27.20	27.74	27.11	0.19
	5.72	6.05	5.22	5.48	4.90	
daily notifications	52.59	53.32	52.70	54.16	53.40	0.99
	70.13	83.67	78.54	74.26	70.23	
meditated daily, ever	0.56	0.59	0.56	0.56	0.57	0.79
	0.50	0.49	0.50	0.50	0.49	
meditated daily, last month	0.47	0.46	0.46	0.46	0.46	0.99
	0.50	0.50	0.50	0.50	0.50	
logged meals, ever	0.88	0.87	0.86	0.87	0.89	0.50
	0.33	0.34	0.35	0.34	0.32	
logged meals, last month	0.33	0.33	0.32	0.33	0.33	0.99
	0.47	0.47	0.47	0.47	0.47	
importance, $x - y$	-0.44	-0.49	-0.42	-0.45	-0.47	0.99
	3.36	3.30	3.39	3.28	3.26	
difficulty - fun, $x - y$	-2.48	-2.76	-2.64	-2.54	-2.54	0.83
	4.93	5.07	5.35	5.16	5.10	

Notes: Means and standard deviations of ten baseline variables. The first five variables were used in re-randomization procedure. F-test of joint significance reported in last column.

Table 3: Orthogonality Check

meditation rate from 9.4% to 18.2%, and nutrition messages raised the meal logging rate from 11.8% to 28.4%. Message treatments also had negative spillovers on non-target behaviors: participants getting only nutrition messages meditated 29% less than the control group (6.6% relative to 9.4%) and participants getting only meditation messages logged meals 19% less than the control group (9.4% relative to 11.8%). The group with both sets of messages did worse than the group with just meditation or just meal-logging messages, by 2.2 and 5.0 percentage points, respectively. There is no evidence of any interaction between the two sets of messages, however, neither for meditation nor for meal logging. Incentives for meal logging had large target effects, more than quadrupling the rate of meal logging (from 11.8% to 50%), as well as negative spillover effects on meditation of 2.5 percentage points. Given Proposition 1, the fact that m_x , m_y , and z_y all generated negative spillovers on the non-target behaviors implies that participants are subject to either

diversion, or depletion, or both. And given Proposition 2, the fact that we find no strong evidence of interference in either direction implies that there is no evidence of overload.

In the post-treatment period, effects of messages on target behaviors persisted, at 28% and 18% the size of their treatment-period magnitude for meditation and meal logging, respectively. Target effects of incentives also persisted, at 10% the size of their treatment period effect. Importantly, negative spillover effects of meal logging messages and incentives on meditation rates persisted at almost 100% of their treatment period effects. There is no evidence, however, that spillovers of meditation messages on meal logging persisted.

	<i>Treatment Period</i>		<i>Post-Treatment Period</i>	
	Meditated (x)	Logged Meal (y)	Meditated (x)	Logged Meal (y)
	(1)	(2)	(3)	(4)
mx	0.088*** (0.011)	-0.024** (0.010)	0.024*** (0.009)	-0.010 (0.006)
my	-0.028*** (0.008)	0.166*** (0.013)	-0.025*** (0.007)	0.029*** (0.008)
mx*my	0.006 (0.014)	-0.026 (0.018)	0.009 (0.011)	-0.002 (0.010)
zy	-0.025*** (0.009)	0.381*** (0.016)	-0.024*** (0.007)	0.037*** (0.008)
mx + mx*my	0.094 (0.009)	-0.050 (0.014)	0.034 (0.007)	-0.013 (0.008)
my + mx*my	-0.022 (0.011)	0.140 (0.012)	-0.016 (0.009)	0.027 (0.007)
Ctrl Mean	0.094	0.118	0.054	0.033
Ctrl SD	0.291	0.323	0.227	0.178
Obs	102905	102905	102499	102499

Notes: OLS regressions at the individual-day level of daily meditation (0/1) and daily logging of at least one meal (0/1) on treatments. mx (my) is 1 if the individual received x (y) messages and 0 otherwise; mx*my is 1 if the individual received both sets of messages. The specification includes controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month) as well as day fixed effects. Standard errors clustered at individual level. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 4: Reduced Form Results

Recall that I have defined depletion to be the possibility that cognitive resources (effort, self-control, executive function) required to take action are costly. If depletion is driving the results,

it should be reflected in two reduced form facts, which will be the basis for the identification of ρ in the structural estimation. First, since a spillover driven by depletion operates through the target action, we should expect interventions with large positive target effects to also have large negative spillovers, and interventions with small positive target effects to have small negative spillovers. I can check this in two ways by comparing two sets of spillover/target effect ratios. First, I can compare spillover/target ratios across messages and incentives, for behavior y , assuming that diversion effects are constant across interventions that promote y . This comparison gives spillover/target ratios of 0.07 and 0.17 for y incentives and messages, respectively, and I can reject that the ratios are equal ($p=0.03$). However, the assumption that messaging and incentive interventions generate equal diversion effects is not a plausible one, given the evidence that attentional capture varies widely by the type of stimulus. The second method is to compare spillover/target ratios of messages across behaviors, assuming that diversion effects are constant across behaviors. The second comparison gives spillover/target ratios of 0.27 and 0.17 for meditation and meal logging, respectively, which are different but not significantly so ($p=0.28$). These tests are thus not conclusive, but they suggest that depletion cannot fully explain the spillovers we see. Since the assumption required by the second comparison is more plausible than the first, I use the second comparison in the estimation, allowing diversion effects to vary by intervention but not by behavior (θ_m, θ_z).

Fortunately, the data provide an additional test (and source of identification) for depletion. Again, because I have defined depletion to be the possibility that cognitive resources required to take action are costly, it must affect the covariance between the two actions. Specifically, if there is no depletion, then the covariance between x and y should be the same across all treatments, since it only reflects any positive or negative relationships between the two behaviors in the utility function. However, in the presence of depletion, this covariance must vary across treatments, since treatments induce action, and action results in depletion.

Table 5 shows the effects of treatments on the covariance between x and y . The model is static and does not specify whether the predictions about the covariance refer to the covariance of x and y across people, or within a person over time. Both are plausible: depletion can conceivably

cause people who meditate to be unlikely to also log their meals; it can also conceivably cause people who meditate on one day not to log their meals on the same day. Therefore, I check both: Column (1) looks at the covariance over individuals (within days) and Column (2) looks at the covariance over days (within individuals). In the presence of depletion, we should expect both sets of messages, as well as incentives, to have negative effects. I see no evidence of this, supporting the conclusion that depletion cannot fully explain the spillovers we observe.

<i>Treatment Period</i>		
	Cov(x,y) over people (within day) (1)	Cov(x,y) over time (within people) (2)
mx	0.001 (0.004)	0.002 (0.001)
my	-0.002 (0.004)	0.001 (0.001)
mx X my	0.021*** (0.005)	0.007*** (0.002)
zy	-0.003 (0.004)	-0.001 (0.001)
Ctrl Mean	0.019	0.008
Ctrl SD	0.138	0.084
Obs	102905	102905

Notes: OLS regressions, including controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 5: Treatment Effects on Covariances

6 Structural Estimation

In this section, I estimate the parameters of the model in order to quantify the contribution of each mechanism to the observed spillovers and draw out the policy implications.

6.1 Structural Model

Let $C(a_x, a_y, w_x, w_y) = f(a_x, a_y) - s^x(w_x, w_y)a_x - s^y(w_x, w_y)a_y$, where $f(a_x, a_y) = \frac{1}{2}\alpha(a_x^2 + a_y^2) + \rho a_x a_y$. Let the attention subsidy in the case of x messages be $s^x(m_x, m_y) = \nu_x m_x + \gamma m_x m_y + \theta_m m_y$, and in the case of y messages, $s^y(m_y, m_x) = \nu_y m_y + \gamma m_x m_y + \theta_m m_x$. Let the attention subsidy in the case of incentives be $s^x(z_x, z_y) = \lambda z_x + \theta_z z_y$ for x and $s^y(z_y, z_x) = \lambda z_y + \theta_z z_x$ for y . I thus assume that attention subsidies targeting behavior j reduce the marginal cost of attention to j by fixed amount s^j . I allow target messages to have different attention subsidies depending on the behavior (allowing different ν_x and ν_y) but I assume that non-target messages impose the same attention tax regardless of the behavior (fixed θ_m, γ for x and y), as discussed above. I also impose that $c_1^x = c_1^y$: since I am allowing x and y to have different baseline returns and responses to target messages, all differences between x and y will be loaded onto the corresponding parameters, and differences between c_1^x and c_1^y will not be identified. In this parameterization, ρ captures depletion, θ_m and θ_z capture diversion for messages and incentives, respectively, and γ captures overload. In my experiment I will only have incentives for y , so z_x will always be zero and λ will only represent the target effect of incentives for y .

In each period, I can thus write the agent's problem as:

$$\max_{a_x, a_y \in [0,1]} \left\{ a_x u_x + a_y u_y - \left(\frac{1}{2}\alpha(a_x^2 + a_y^2) + \rho a_x a_y \right. \right. \\ \left. \left. - a_x(\nu_x m_x + \gamma m_x m_y + \theta_m m_y + \theta_z z_y) - a_y(\nu_y m_y + \gamma m_y m_x + \lambda z_y + \theta_m m_x) \right) \right\}$$

Let a_{xit}^* represent the optimal attention paid to behavior x by individual i in period t . I allow for individual heterogeneity in u_x and u_y . I let $u_{xi} = \mu_x + \epsilon_{xi}$, and I normalize μ_y to 1, so that $u_{yi} = 1 + \epsilon_{yi}$. I assume that ϵ_{xi} and ϵ_{yi} are jointly normal, with mean zero, variances $\sigma_{\epsilon_x}^2$ and $\sigma_{\epsilon_y}^2$, and covariance $\sigma_{\epsilon_x \epsilon_y}$. I also allow for individual heterogeneity in the effects of target messages, so that $\nu_x = \phi_x + \delta_{xi}$ and $\nu_y = \phi_y + \delta_{yi}$. I assume that δ_{xi} and δ_{yi} are jointly normal, with mean zero, variances $\sigma_{\delta_x}^2$ and $\sigma_{\delta_y}^2$, and covariance $\sigma_{\delta_x \delta_y}$.

Recall that in the model, the only things that change over time are the distraction shocks ξ_{it}^x and ξ_{it}^y , which are re-drawn each period. I define $a_{xi}^* = E_{\xi_x}[a_{xit}^*]$, which I can estimate in the data

as: $\hat{a}_{xi}^* = \frac{1}{T} \sum_{t=1}^T x_{it}$.

6.2 Estimation

I estimate the model with a minimum-distance estimator. In the benchmark specification, I use the following 22 moments: the control means of a_{xi}^* and a_{yi}^* , the treatment effects on a_{xi}^* and a_{yi}^* , the control variances of a_{xi}^* and a_{yi}^* , the main and interaction effects of m_x and m_y on the variances of a_{xi}^* and a_{yi}^* , the control covariance of a_{xi}^* and a_{yi}^* , and the main and interaction effects of m_x and m_y on the covariance of a_{xi}^* and a_{yi}^* .⁶

The relationship between α and ρ is identified in two ways in the data. First, it can be identified using the four target and spillover effects $\frac{\partial a_x^*}{\partial m_x}$, $\frac{\partial a_x^*}{\partial m_y}$, $\frac{\partial a_y^*}{\partial m_y}$, and $\frac{\partial a_y^*}{\partial m_x}$. The intuition is the same as previously described: if ρ is zero, then we expect $\frac{\partial a_x^*}{\partial m_y} = \frac{\partial a_y^*}{\partial m_x}$. If ρ is positive, then the difference between $\frac{\partial a_x^*}{\partial m_y}$ and $\frac{\partial a_y^*}{\partial m_x}$ should reflect the difference between $\frac{\partial a_y^*}{\partial m_y}$ and $\frac{\partial a_x^*}{\partial m_x}$, since both are driven by differences between ϕ_x and ϕ_y . Second, it is identified by the differences in the covariance between a_{xi} and a_{yi} across treatments. Again, the intuition is the same as previously described: if ρ is zero, then there should be no difference in the covariance across treatments; if ρ is positive, then treatments that induce higher x or y should also induce a lower covariance.

Once I have pinned down the relationship between α and ρ , I can combine this with the control group means $E[a_x|ctrl]$ and $E[a_y|ctrl]$ to separately identify α and ρ . The remaining parameters are straightforward to identify once α and ρ are known. Specifically, ϕ_x , ϕ_y , λ , θ_m , and θ_z are identified by the three target effects and three spillover effects of messages and incentives. The identification of the diversion parameters depends critically on ρ having already been pinned down. γ is (over-)identified by the two interference effects; this arises mechanically from the way overload was defined in the model. σ_{ϵ_x} , σ_{ϵ_y} , and σ_{xy} are identified by the variances and covariances of a_{xi} and a_{yi} in the control group, and σ_{phi_x} , σ_{phi_y} , and $\sigma_{\phi_x\phi_y}$ are identified by the variances and covariances of a_{xi} and a_{yi} across treatments.

Let ζ represent the vector of q parameters and let $m(\zeta)$ represent the r moments as functions of the parameters. The minimum-distance estimator selects parameters $\hat{\zeta}$ that minimize

⁶I do not use the effects of incentives on variances and covariances because for the sake of simplicity, I have not allowed for heterogeneity in the incentive attention subsidy.

the expression $(m(\zeta) - \hat{m})'W(m(\zeta) - \hat{m})$. For the weighing matrix W I use the diagonal of the inverse of the variance-covariance matrix of the moments.⁷ I estimate the variance of $\hat{\zeta}$ as $(\hat{G}'W\hat{G})^{-1}(\hat{G}'W\hat{\Lambda}W\hat{G})(\hat{G}'W\hat{G})^{-1}$, where $\hat{G} = \Delta_{\zeta}m_n(\hat{\zeta})$ (the matrix of derivatives of the moments with respect to parameters, evaluated at the estimated parameters) and $\hat{\Lambda} = Var(\hat{m})$.

The maximized value of the objective function is asymptotically distributed as $\chi^2(r - q)$, so the critical value for an over-identification test of model fit is 2.17 (for the benchmark specification). The maximized value of the objective function is 12.3, so the test is rejected. Figure 6.2 compares the actual moments and predicted moments. The high test statistic is driven principally by moment 12, $\frac{\partial var(ax)}{\partial my}$. It turns out that in the data, this estimate is significantly negative, at -0.015 (0.005). In the model, the variance is indeed predicted to fall, but only by a very small amount. Since the economic magnitude of this deviation is small, and since this moment is not critical to the identification of the parameters of interest, I do not consider this to be strong evidence that the model is wrong.




Figure 2: Model Fit

I report parameter estimates in Table 6. There is no evidence of overload, as $\hat{\gamma}$ is not significantly different from zero. Interestingly, $\hat{\rho}$ is also very close to zero, though its standard error is

⁷With small samples, using the full variance-covariance (VC) matrix results in biased estimates (Altonji and Segal, 1996). Indeed, in this case, similar estimates are obtained using the diagonal of the VC matrix or the identity matrix produces, while different estimates are obtained using the full VC matrix.

very high. Message diversion is both large in magnitude and statistically significant; incentive diversion is large in magnitude but its standard error is high. These results confirm that message diversion is an important driver of spillovers. They do not rule out the possibility of depletion, but they do suggest that depletion alone cannot explain the results.

<i>Parameter</i>	<i>Description</i>	<i>Estimate</i>	<i>Standard Error</i>
β	slope of marginal cost of attention	6.145	0.619
ρ	depletion	-0.029	0.522
μ_x	return to x	0.799	0.081
ϕ_x	x message attn subsidy	0.594	0.126
γ	overload	-0.001	0.109
θ_m	message diversion	-0.209	0.085
λ	y incentive attn subsidy	2.059	0.277
ϕ_y	y message attn subsidy	0.963	0.173
σ_{ϵ_x}	S.D. of ϵ_x , heterogeneity in return to x	1.413	0.134
σ_{ϵ_y}	S.D. of ϵ_y , heterogeneity in return to y	1.865	0.093
θ_z	incentive diversion	-0.190	0.182
σ_{xy}	covariance of ϵ_x and ϵ_y	0.741	0.517
σ_{δ_x}	S.D. of δ_x , heterogeneity in message subsidy	1.000	0.240
σ_{δ_y}	S.D. of δ_y , heterogeneity in message subsidy	1.671	0.291
$\sigma_{\phi_x\phi_y}$	covariance of δ_x and δ_y	1.000	0.342

Table 6: Parameter Estimates

In Appendix F Table 13, I present estimates using the same moments and parameters, but using the identity matrix as the weighing matrix. The estimates differ in predictable ways, but all within 95% confidence intervals of the benchmark estimates, and the main conclusions are consistent. In fact, the estimates for θ_m and θ_z are even more negative, as is the estimate for ρ . In Appendix F Figure 13 I examine the sensitivity of the four key parameter estimates to moments as in Andrews et al. (2017).

7 Additional Evidence on Mechanisms

7.1 Expectations, Opt-Outs, and Reading of Messages

In this section I examine data collected at three different points throughout the experiment to check for consistency with the mechanisms identified above, and to potentially shed light on some miss-

ing details. First, I look at the two observed measures of the extent to which participants actually read and internalized the content of messages. The first measure is whether or not participants responded to a surprise raffle sent to participants with messages via SMS.⁸ The message said: "Hi from Remindful/eNOMerate! We are offering a surprise raffle for a USD 20 Amazon gift card. To enter, tap [link] and press send. Msg & dta rates may apply." Each message participant received a maximum of one raffle message over the course of the experiment. Roughly half of message participants received the raffle on day 10 or 11 (halfway through the second week) and the other half received the raffle on day 20 or 21 (at the end of the third week). Participants receiving both messages were randomly assigned to receive either the eNOMerate raffle or the Remindful raffle.

The second indicator of internalization of messages is knowledge about meditation and nutrition, as measured by the percentage of questions answered correctly on a quiz administered at the end of the study, one month after the end of treatment. The quiz consisted of true/false questions on information provided in the messages, with additional options to answer "I remember seeing this message but I do not remember the details" or "I do not remember seeing this message." Participants received 1 point for every correct answer, 0 point for every incorrect or "I do not remember seeing this message" answer, and 0.5 points for answering "I remember seeing this message but I do not remember the details." They were unaware of this scoring system, and had no explicit incentives to perform well. Both raffle response rates and quiz scores are restricted to participants with messages (I do not quiz groups on information they did not receive), so the omitted group will have one set of messages, and the treatment group will have both.

Table 7 displays the results. Raffle response rates were low even in the groups with one set of messages, at 31% and 26% for meditation and meal logging raffles, respectively (but higher conditional on not opting-out, at 36% and 31% respectively). The group with both sets of messages was about 30% less likely to respond to both raffles, suggesting that they were reading messages at a lower rate. This might explain why these groups also did slightly worse on the knowledge quizzes, as demonstrated in columns (2) and (4), though these effects are not significant. The coefficient of interest does not change substantially when we condition on having not opted out,

⁸Due to an implementation error, we are missing this data from 592 participants with messages. These participants accidentally received messages with a broken link, and so we do not know whether they responded or not.

suggesting that the bulk of the effect is driven by participants who continued to receive messages throughout the treatment period.

	Raffle Response Meditation (x) (1)	Knowledge Score Meditation (x) (2)	Raffle Response Nutrition (y) (3)	Knowledge Score Nutrition (y) (4)
<i>Panel A: Unconditional</i>				
mx & my	-0.089** (0.028)	-0.010 (0.006)	-0.081** (0.028)	-0.008 (0.007)
mx-only Mean	0.307	0.298		
mx-only SD	0.461	0.086		
my-only Mean			0.263	0.300
my-only SD			0.440	0.100
Obs	998	859	931	847
<i>Panel B: Conditional on Not Opting-Out</i>				
mx & my	-0.087** (0.033)	-0.005 (0.006)	-0.084** (0.034)	-0.005 (0.007)
mx-only Mean	0.356	0.307		
mx-only SD	0.479	0.079		
my-only Mean			0.309	0.308
my-only SD			0.463	0.098
Obs	825	762	771	748

Notes: Outcomes are whether individual responded to surprise raffle (Columns 1 and 3) and score on knowledge quiz (Columns 2 and 4). Omitted groups are mx-only (Column 1) and my-only (Column 2). Regressions include controls for five baseline variables on which re-randomization was based. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in table.

Table 7: Treatment Effects on Attention to Messages: Reading Rate and Memory of Content

A simple depletion story cannot explain the results from Table 7. Being cognitively depleted can explain why someone neglects to do a behavior, but it cannot explain why someone neglects to read a message. However, a more subtle depletion story, in which people anticipate depletion and therefore make an active decision to focus on one behavior or another, could explain why people with both sets of messages are less likely to read either set.

I can look for evidence of such a decision in two places. First, I look at data on expectations that we collected from participants in Survey 2, immediately after treatment assignment, but before treatment began. Table 8 shows regressions of individual-level expectations about rates of

meditation (Column (1)) and meal logging (Column (3)) on treatment assignment. Columns (2) and (4) show the differences between expected and actual rates of meditation and meal logging, respectively. Participants in the control group grossly over-predict their rates of both behaviors: they expect to meditate 38% of the time when actually they do so 9.4% of the time, and they expect to log their meals 52% of the time when they actually do so 11.8% of the time. Participants who receive only meditation messages predict their meditation rates to be even higher than the control group, but the level of over-prediction is similar, resulting in an “expected target effect” (11%) that is actually quite close to the true target effect (8.8%). Participants who receive only nutrition messages also predict their meal-logging rates to be higher than the control group, but here they over-predict less, resulting in an “expected target effect” (12%) that is significantly lower than the true target effect (16.6%).

“Expected spillover effects” of nutrition messages and incentives on meditation are small and positive (though neither of these effects are significant). Participants thus significantly over-estimate spillover effects (i.e. they do not expect negative spillovers) by 5 percentage points in both cases. This means we can be confident that these spillovers were not driven by a decision in advance to focus on meal logging, once treatment assignment is known. In the case of the spillover of meditation messages on meal logging, we cannot conclude that participants correctly predicted the negative spillover (the expected effect is about 70% of the true spillover, at -0.017, but the standard error is high), but there is no evidence that they under-estimated spillovers either.

To compare this data to Table 7, we need to look at the comparison between the groups that got one set of messages and the group that got both: the coefficient on $m_y + m_x m_y$ in Columns 1 and 2, and the coefficient on $m_x + m_x m_y$ in Columns 3 and 4. In the former case, participants with m_x over-estimate spillovers of m_y on x by 2.4 percentage points, but the effect is not significant. In the latter case, participants with m_y significantly over-estimate spillovers of m_x on y by 6.1 percentage points; this is mostly due to a positive expected interaction effect. Overall, the data on expectations does not provide any evidence that participants with both sets of messages chose to focus on one behavior or another upon treatment assignment.

It is possible, however, that participants only understood the consequences of depletion once

	Expected Meditation (x) (1)	Expected - Actual Meditation (x) (2)	Expected Meal Logging (y) (3)	Expected - Actual Meal Logging (y) (4)
mx	0.113*** (0.012)	0.008 (0.015)	-0.017 (0.017)	0.009 (0.018)
my	0.017 (0.012)	0.053*** (0.013)	0.122*** (0.014)	-0.042** (0.018)
mx X my	-0.027 (0.017)	-0.029 (0.021)	0.037 (0.021)	0.052* (0.025)
zy	0.021 (0.013)	0.053*** (0.014)	0.213*** (0.015)	-0.223*** (0.021)
mx + mxmy	0.085 (0.012)	-0.020 (0.014)	0.019 (0.013)	0.061 (0.018)
my + mxmy	-0.010 (0.012)	0.024 (0.016)	0.159 (0.015)	0.010 (0.018)
Ctrl Mean	0.384	0.278	0.519	0.382
Ctrl Mean S.E.	(0.009)	(0.010)	(0.012)	(0.012)
Obs	2891	2876	2891	2886

Notes: Expected rates of behavior at baseline over treatment period (Columns 1 and 3) and difference between individual's expected and actual rate (Columns 2 and 4). Regressions include controls for the five baseline variables on which re-randomization was based. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 8: Baseline Expectations of Behaviors by Treatment

the treatment period began, and only then made a decision to focus on one behavior or another. If this were the case, then they should have opted out of one or the other messaging treatment, assuming there is some cost to receiving extraneous messages. To opt-out, participants simply had to reply "STOP" to the same number from which they were receiving messages. They were informed of this in the consent form, by email upon treatment assignment, as well as in the first text message they received. Table 9 examines whether participants receiving meditation messages ever opted out (Column 1), and whether participants receiving nutrition messages ever opted out (Column 2), where the omitted group has one set of messages, and the treatment group has both. Note that here we are capturing the combination of spillover and interaction effects.

Baseline levels of opt-out were relatively high, at 14.5% for the meditation-only group and 16.8% for the nutrition-only group. But there is no evidence that participants receiving both sets of

messages opted out more frequently than those with just one, especially when it came to the meditation program. So there is no evidence that people receiving both sets of messages ever made a conscious choice to focus on one behavior and not the other—not at the beginning upon treatment assignment, nor during the treatment period. When taken together, then above results provide some support for the conclusion that depletion is unlikely to be a major driver of spillovers. A story of unanticipated diversion, on the other hand, is consistent with the above results. People receiving both sets of messages are less likely to read either set of messages, not as the result of an active choice in anticipation of depletion, but rather because one set of messages diverts their attention from the other, bumping it from the top of mind.

	Opted Out Ever, Meditation Msgs (x) (1)	Opted Out Ever, Nutrition Msgs (y) (2)
mx & my	-0.029 (0.017)	0.030 (0.019)
mx-only Mean	0.071	
mx-only SD	0.257	
my-only Mean		0.086
my-only SD		0.281
Obs	1585	1625

Notes: Outcome is whether individual ever opted out of messaging program. Omitted groups are mx-only (Column 1) and my-only (Column 2). Regressions include controls for five baseline variables on which re-randomization was based. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in table.

Table 9: Treatment Effects on Opting Out of Message Programs

7.2 Can the Time Constraint Account for Spillovers?

In theory, the observed spillovers could also be driven by time constraints. In my model, time constraints would be inseparable from depletion, since they operate through the actions themselves. The fact that I found little evidence for depletion, therefore, suggests that time constraints are not driving the results either. As an additional check of this possibility, I use my estimates to compute the implied elasticity of substitution between meditation and meal logging. Suppose agents have

CES utility, $u(h_x, h_y) = (\alpha h_x^\rho + (1 - \alpha) h_y^\rho)^{1/\rho}$, where h_x and h_y are daily hours spent on meditation and meal logging, respectively. It can be easily shown that the cross-price elasticity of meditation with respect to meal logging ϵ_{xy} is equal to $(\sigma - 1) * s_y$, where $s_y = \frac{h_y p_y}{Y}$, the share of income spent on y , and $\sigma = \frac{1}{1-\rho}$, the elasticity of substitution. In our case, $s_y = \frac{h_i}{24}$, if I let p_i represent the value of 1 hour of time.

I can use my estimates to approximate a lower bound for ϵ_{xy} . Table 15 shows that daily average minutes meditated fell from 1.8 to 0.98, a 45% decrease, in the presence of incentives for meal logging. The incentives had an expected value of \$0.37 per successful day (with at least one meal logged), so I assume the price of meal logging fell by this amount. I approximate the price of meal logging to be $20 * (3.78/60) = 1.26$, assuming that the average hourly wage for a college graduate is \$20.00⁹, and that the minimum necessary time spent on meal logging was 3.78 minutes (the time required to log one meal as reported in the final survey, which is a lower bound for the time spent on daily meal logging). Together, this implies that incentives for meal logging reduced the price of meal logging by 29.4%, and that the cross price elasticity, ϵ_{xy} , is equal to 1.53. To calculate h_y , I use the time required to log all of one's meals as reported in the final survey, 11.82 minutes (an upper bound for the time spent on daily meal logging) and multiply it by the average daily rate of logging at least one meal across all groups, which was 0.14. Combining everything, I find that in order to explain the spillovers we see, the elasticity of substitution would need to be at least 1331. In summary, these behaviors take so little time that the elasticity of substitution would need to be unrealistically high in order for the time constraint to explain the observed spillovers.

8 Policy Implications

I use the parameter estimates from Table 6 in two ways. First, I decompose the observed spillovers from Table 4 by type of limited attention. Second, I quantify the costs of different types of interventions once spillovers are taken into account.

Given that my benchmark estimates of ρ and γ are negative and indistinguishable from zero, respectively, the implication is that 100% of spillovers are driven by diversion. Standard errors,

⁹citation

computed using 500 bootstrapped samples of 600 participants, are 0.37 and 0.42 for the shares of spillovers explained by message and incentive diversion, respectively. Thus, the lower bound for the share of message spillovers explained by diversion is 26%, and the lower bound for the share of incentive spillovers explained by diversion is 14%.¹⁰

What does this decomposition of spillovers tell us about how to implement behavior change interventions? Since I found no evidence of overload (γ was significantly different from zero), and no significant differences between message and incentive diversion, there is no evidence to suggest that policymakers should be concerned about interventions with large amounts of stimuli or information. How externally valid is this conclusion? The messaging interventions I ran in the experiment were more intensive than average, with two daily messages per program. However, there are many dimensions of messaging programs that should be researched more thoroughly prior to implementation: the difficulty of content, the length, and the timing of the message, just to name a few. Though I found no evidence of heterogeneity by the baseline amount of notifications, the experiment was not powered to detect these effects. Finally, more research should be done before concluding that overload is not a concern for other, very different types of informational interventions, like mandatory information sessions or workshops.

The strong evidence for the existence of message diversion, on the other hand, has an important policy implication. Spillovers driven by diversion, unlike those driven by depletion, can exist even when the intervention is not effective. In other words, an intervention with a large target effect will not necessarily create larger spillovers than one with a small target effect. Thus, two equally cost-effective interventions that differ in effectiveness will not necessarily generate the same spillovers. I quantify this by asking the question: if it costs \$1 to increase x and y by one standard deviation in the absence spillovers, how much does it cost in the presence of spillovers? In other words, assuming that we care about increasing both x and y (and abstracting from the complicated question of the welfare benefits of different amounts of different behaviors), how much more costly are interventions once spillovers are taken into account? I first re-estimate the model, fixing θ_m at different values between its benchmark estimate and zero. I then plug in the

¹⁰This computation is ongoing, so final numbers may be different.

corresponding estimates for ρ and γ when I compute the cost of the intervention in the presence of spillovers. I only conduct this exercise for messaging interventions, as the uncertainty around incentive diversion is too high to allow for a reasonably confident policy conclusion.

An important caveat when interpreting this calculation is that although I know that the extent of diversion does not *necessarily* depend on the effectiveness of the intervention, I cannot say that the two are unrelated. In fact, it is plausible that on average, more effective interventions do generate more diversion. In the meantime, I assume that effectiveness and diversion are unrelated, so the results should be interpreted accordingly.

Figure 3 shows the results. On the x-axis I vary the effectiveness of the intervention (the target standard deviations achieved, which is captured by ϕ in the model). On the y-axis I vary the diversion parameter, θ_m . As to be expected, when θ_m is zero, there is very little variation along the y axis, since spillovers are driven mostly by depletion, and therefore intervention effectiveness does not matter.¹¹ When spillovers are driven primarily by depletion, interventions cost 8-16% more in the presence of spillovers. As diversion becomes stronger, intervention effectiveness matters more and more. When they are driven by diversion, even interventions with very high effectiveness (1 S.D.) still cost 32% more in the presence of spillovers, with a standard error of 11%. And interventions with low effectiveness (0.2 SD) can cost as much as 69% more in the presence of spillovers, with a standard error of 27%. This means that in the presence of spillovers driven by diversion, an intervention with a 0.2 SD effect is predicted to cost 28% more than an equally locally cost-effective intervention with a 1 SD effect.

9 Conclusion

A growing literature has documented a wide variety of interventions that shift behavior and generate meaningful economic impacts. This paper has demonstrated the importance of studying these interventions in a broader context, taking into account how they affect other interventions and other behaviors. In particular, limited attention can cause interventions to have negative

¹¹The variation that does exist is because when θ_m is forced to be zero, the estimate for γ becomes more negative, and in the presence of overload intervention effectiveness also matters




Figure 3: Quantifying Costs of Spillovers by Intervention Effectiveness

spillovers on seemingly unrelated behaviors. These spillovers do not necessarily grow with the effectiveness of the intervention, suggesting that small-scale, highly-effective interventions may be generally preferable to large-scale, less effective ones, cost-effectiveness held constant. This paper raises many additional questions about the details of these effects; their generalizability to other behaviors, contexts, and interventions; and their policy implications.

References

- Afshin, A., Sur, P. J., Fay, K. A., Cornaby, L., Ferrara, G., Salama, J. S., Mullany, E. C., Abate, K. H., Abbafati, C., Abebe, Z., et al. (2019). Health effects of dietary risks in 195 countries, 1990–2017: a systematic analysis for the global burden of disease study 2017. *The Lancet*.
- Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.
- Altonji, J. G. and Segal, L. M. (1996). Small-sample bias in gmm estimation of covariance structures. *Journal of Business & Economic Statistics*, 14(3):353–366.
- Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly journal of economics*, 127(3):1243–1285.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.
- Botvinick, M. and Braver, T. (2015). Motivation and cognitive control: from behavior to neural mechanism. *Annual review of psychology*, 66:83–113.
- Burgess, P. W., Scott, S. K., and Frith, C. D. (2003). The role of the rostral frontal cortex (area 10) in prospective memory: a lateral versus medial dissociation. *Neuropsychologia*, 41(8):906–918.
- Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.
- Chernev, A., Böckenholt, U., and Goodman, J. (2015). Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, 99(4):1145–77.

- Chun, M. M., Golomb, J. D., and Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual review of psychology*, 62:73–101.
- De Clippel, G., Eliaz, K., and Rozen, K. (2014). Competing for consumer inattention. *Journal of Political Economy*, 122(6):1203–1234.
- DellaVigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64(2):709–749.
- Dolan, P. and Galizzi, M. M. (2015). Like ripples on a pond: behavioral spillovers and their implications for research and policy. *Journal of Economic Psychology*, 47:1–16.
- Dubois, P., Griffith, R., and O’Connell, M. (2019). How well targeted are soda taxes?
- Duncan, J., Martens, S., and Ward, R. (1997). Restricted attentional capacity within but not between sensory modalities. *Nature*, 387(6635):808.
- Dunn, T. L., Lutes, D. J., and Risko, E. F. (2016). Metacognitive evaluation in the avoidance of demand. *Journal of Experimental Psychology: Human Perception and Performance*, 42(9):1372.
- Edmunds, A. and Morris, A. (2000). The problem of information overload in business organisations: a review of the literature. *International journal of information management*, 20(1):17–28.
- Eliaz, K. and Spiegler, R. (2011). On the strategic use of attention grabbers. *Theoretical Economics*, 6(1):127–155.
- Embi, P. J. and Leonard, A. C. (2012). Evaluating alert fatigue over time to ehr-based clinical trial alerts: findings from a randomized controlled study. *Journal of the American Medical Informatics Association*, 19(e1):e145–e148.
- Eppler, M. J. and Mengis, J. (2004). The concept of information overload: A review of literature from organization science, accounting, marketing, mis, and related disciplines. *The information society*, 20(5):325–344.
- Falkinger, J. (2008). Limited attention as a scarce resource in information-rich economies. *The Economic Journal*, 118(532):1596–1620.

- Farhi, E. and Gabaix, X. (2018). Optimal taxation with behavioral agents.(no. w21524).
- Fawcett, J., Kingstone, A., and Risko, E. (2015). *The handbook of attention*. MIT Press.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.
- Gonen-Yaacovi, G. and Burgess, P. (2012). Prospective memory: The future for future intentions. *Psychologica Belgica*, 52(2-3).
- Goyal, M., Singh, S., Sibinga, E. M., Gould, N. F., Rowland-Seymour, A., Sharma, R., Berger, Z., Sleicher, D., Maron, D. D., Shihab, H. M., et al. (2014). Meditation programs for psychological stress and well-being: a systematic review and meta-analysis. *JAMA internal medicine*, 174(3):357–368.
- Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69(6):1403–1435.
- Gul, F. and Pesendorfer, W. (2004). Self-control and the theory of consumption. *Econometrica*, 72(1):119–158.
- Hanna, R., Duflo, E., and Greenstone, M. (2016). Up in smoke: the influence of household behavior on the long-run impact of improved cooking stoves. *American Economic Journal: Economic Policy*, 8(1):80–114.
- Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353.
- Holman, J. and Zaidi, F. (2010). The economics of prospective memory. *Available at SSRN 1662183*.
- Hwang, M. I. and Lin, J. W. (1999). Information dimension, information overload and decision quality. *Journal of information science*, 25(3):213–218.
- Institute for Health Metrics and Evaluation (2018). Findings from the global burden of disease study 2017. Technical report, Seattle, Washington.
- Karlan, D., McConnell, M., Mullainathan, S., and Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, 62(12):3393–3411.

- Khan, U. and Dhar, R. (2006). Licensing effect in consumer choice. *Journal of marketing research*, 43(2):259–266.
- Krpan, D., Galizzi, M. M., and Dolan, P. H. (2019). Looking at spillovers in the mirror: Making a case for "behavioural spillunders". *Frontiers in psychology*, 10:1142.
- Lacetera, N., Pope, D. G., and Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5):2206–36.
- Luck, S. J. and Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657):279.
- Mani, A., Mullainathan, S., Shafir, E., and Zhao, J. (2013). Poverty impedes cognitive function. *science*, 341(6149):976–980.
- Manzini, P. and Mariotti, M. (2012). Categorize then choose: Boundedly rational choice and welfare. *Journal of the European Economic Association*, 10(5):1141–1165.
- Masatlioglu, Y., Nakajima, D., and Ozbay, E. Y. (2012). Revealed attention. *American Economic Review*, 102(5):2183–2205.
- Medina, P. C. (2017). Selective attention in consumer finance: Evidence from a randomized intervention in the credit card market. Technical report, Working Paper, Northwestern University.
- Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences*, 108(26):10415–10420.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.
- Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.
- Noor, J. (2007). Commitment and self-control. *Journal of Economic theory*, 135(1):1–34.

- O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, 90(10-11):1825–1849.
- Ozdenoren, E., Salant, S. W., and Silverman, D. (2012). Willpower and the optimal control of visceral urges. *Journal of the European Economic Association*, 10(2):342–368.
- Rogers, T. and Milkman, K. L. (2016). Reminders through association. *Psychological Science*, 27(7):973–986.
- Schwartzstein, J. (2014). Selective attention and learning. *Journal of the European Economic Association*, 12(6):1423–1452.
- Shah, A. K., Mullainathan, S., and Shafir, E. (2012). Some consequences of having too little. *Science*, 338(6107):682–685.
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., and Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual review of neuroscience*, 40:99–124.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.
- Strahilevitz, M. and Myers, J. G. (1998). Donations to charity as purchase incentives: How well they work may depend on what you are trying to sell. *Journal of consumer research*, 24(4):434–446.
- Taubinsky, D. (2013). From intentions to actions: A model and experimental evidence of inattentive choice. *Work Pap.*
- Taubinsky, D. and Rees-Jones, A. (2017). Attention variation and welfare: theory and evidence from a tax salience experiment. *The Review of Economic Studies*, 85(4):2462–2496.
- Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Werthenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing science*, 17(4):317–337.

- Wharton, C. M., Johnston, C. S., Cunningham, B. K., and Sterner, D. (2014). Dietary self-monitoring, but not dietary quality, improves with use of smartphone app technology in an 8-week weight loss trial. *Journal of nutrition education and behavior*, 46(5):440–444.
- Yantis, S. and Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. *Journal of Experimental Psychology: Human perception and performance*, 10(5):601.
- Zhou, F., Shefer, A., Wenger, J., Messonnier, M., Wang, L. Y., Lopez, A., Moore, M., Murphy, T. V., Cortese, M., and Rodewald, L. (2014). Economic evaluation of the routine childhood immunization program in the united states, 2009. *Pediatrics*, 133(4):577–585.

Appendices

A Mathematical Proofs

A.1 Comparative Statics

$$\frac{\partial a_x^*}{\partial w_y} = \frac{-c_1^y c_4^x + c_2^x c_3^y}{c_1^x c_1^y - c_2^x^2} \quad (1)$$

$$\begin{aligned} \frac{\partial^2 a_x^*}{\partial m_x \partial m_y} &= \frac{c_1^y c_{34}^x - c_2^x c_{34}^y - c_1^y \left((c_{11}^x \frac{\partial a_x}{\partial m_y} + c_{12}^x \frac{\partial a_y}{\partial m_x}) \frac{\partial a_x}{\partial m_x} + (c_{21}^x \frac{\partial a_x}{\partial m_y} + c_{12}^x \frac{\partial a_y}{\partial m_x}) \frac{\partial a_y}{\partial m_x} \right)}{c_1^x c_1^y - c_2^x^2} \\ &\approx \frac{c_1^y c_{34}^x - c_2^x c_{34}^y}{c_1^x c_1^y - c_2^x^2} \quad (2) \end{aligned}$$

$$\begin{aligned} \frac{\partial^2 a_y^*}{\partial m_x \partial m_y} &= \frac{c_1^x c_{34}^y - c_2^x c_{34}^x - c_1^x \left((c_{11}^y \frac{\partial a_y}{\partial m_x} + c_{12}^y \frac{\partial a_x}{\partial m_x}) \frac{\partial a_y}{\partial m_y} + (c_{21}^y \frac{\partial a_y}{\partial m_x} + c_{12}^y \frac{\partial a_x}{\partial m_x}) \frac{\partial a_x}{\partial m_y} \right)}{c_1^x c_1^y - c_2^x^2} \\ &\approx \frac{c_1^x c_{34}^y - c_2^x c_{34}^x}{c_1^x c_1^y - c_2^x^2} \quad (3) \end{aligned}$$

A.2 Overload Proof

Re-write Equations 2 and 3 as the following, incorporating the assumption that $c_{34}^x = c_{34}^y$.

$$\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} \approx \frac{-c_{34}^x (c_1^y - c_2^x)}{c_1^x c_1^y - c_2^x^2} \quad (4)$$

$$\frac{\partial^2 a_y^*}{\partial m_x \partial m_y} \approx \frac{-c_{34}^x (c_1^x - c_2^x)}{c_1^x c_1^y - c_2^x^2} \quad (5)$$

Recall that we have assumed $c_1^x > 0$ and $c_1^y > 0$. For the case where $c_2^x \leq 0$, we can quickly see from Equations 4 and 5 that $c_{34}^x = c_{34}^y > 0 \implies$ interference in both directions, and interference in either direction implies $c_{34}^x = c_{34}^y > 0$.

For the case where $c_2^x > 0$, recall that for the existence of a local maximum we have also

assumed that $c_1^x c_1^y - c_2^{x^2} > 0$ which implies that either $c_1^x > |c_2^x|$ or $c_1^y > |c_2^x|$. Then we can see that if there is interference in both directions ($\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} < 0$ and $\frac{\partial^2 a_y^*}{\partial m_x \partial m_y} < 0$), then it must be true that $c_{34}^x = c_{34}^y > 0$. And if $c_{34}^x = c_{34}^y > 0$, there must be interference in at least one direction ($\frac{\partial^2 a_x^*}{\partial m_x \partial m_y} < 0$ or $\frac{\partial^2 a_y^*}{\partial m_x \partial m_y} < 0$).

B Experiment Protocol

I recruited participants using the below Facebook ad, targeting people in the U.S. age 18-65 and allowing the algorithm to train to maximize “conversions,” or successful completions of Survey 1.




Figure 4: Facebook Ad Used to Recruit Participants

The first part of Survey 1 was an eligibility test in which participants had to verify six things: (1) ownership of iPhone or Android phone; (2) age 18 or over; (3) interested in working on wellness habits like daily meditation and tracking your nutrition; (4) willing to download two (free) wellness-related smartphone apps for the study; (5) comfortable potentially using a nutrition tracking app;¹² (6) have not already participated in the study. Participants then provided elec-

¹²We received feedback that some participants who had struggled with eating disorders or body image issues in the past were ultimately uncomfortable using the meal tracking app.

tronic consent, which included consenting to receive SMS messages associated with the study. The consent form also described the rewards for participation in the study: entrance into a raffle for a \$20 Amazon gift card for participation in Survey 1 and app download; and entrance into a raffle for a \$50 Amazon gift card for participation in Survey 2 or Survey 3. (If they completed both they were entered twice.)

The next part of Survey 1 was the app download. Participants were given two options: to either download the two apps now, or to download them after finishing the survey. Either way, they were told that they had to download both apps within 24 hours of completing the survey in order to be enrolled. They were told that if they downloaded the apps after 24 hours, they could still enroll, but they should email us to let us know. They were then shown a screen with instructions about how to download each app, which were also emailed to them upon survey completion. These instructions included a temporary assigned password, which enabled us to access their data for the duration of the study.

As described in the paper, the rest of the survey included basic demographic questions; questions on past meditation, exercise, meal tracking, and sleep; questions about the full set of notifications, across devices and apps, received by the participant; and questions about the participant's perceived importance, difficulty, and "fun" of meditation, exercise, meal tracking, and sleep. Upon completion of Survey 1, participants were sent an email repeating the instructions for how to download the apps, including their assigned passwords.

Every day, we verified whether new Survey 1 participants for whom 24 hours had elapsed since survey completion (plus old Survey 1 participants who emailed us) downloaded both apps. Those who did were randomized to one of the five treatments, using a script that re-randomized to ensure balance across the full sample. They were then sent an enrollment confirmation email, displayed in Figure 5.

For example, participants received to receive messages about meditation only were told, "you have been randomly assigned to receive messages about meditation with [app], as part of our Remindful program." The goal is to make clear that the behavior we have in mind is not meditation generally, but meditation specifically with the assigned app. Below the treatment assignment,

Welcome to the Yale Wellness and Technology Study! You successfully downloaded the apps and are officially enrolled. This email contains lots of information about the study. You can refer back to it throughout the study if you have questions.

Here is a brief summary of what it contains:

1. Your (random) assignment to messages or incentives for one or more wellness behaviors. **You were assigned:** . See below for details!
2. The [link to Survey 2](#), which will expire in 24 hours. This survey is not mandatory, but takes only a few minutes, and if you participate in time, we'll enter your name into our second raffle (for a \$50 Amazon gift card).
3. Survey 1 raffle results
4. A reminder of your password for [REDACTED]
5. Study duration and how to withdraw

Be well and let us know if you have questions.

Best,
Hannah

1. Your (Random) Assignment to Messages or Incentives

You have been randomly assigned to receive . [INSERT TREATMENT]

Remember, because this is an experiment, this assignment was completely random. It has nothing to do with your survey responses, or with how important we think meditation, exercise, nutrition, and sleep are. (They're all important!)

Regardless of any programs you were or were not assigned above, your ultimate use of [REDACTED] is entirely up to you. You are welcome but not obligated to use these apps for the study, so please use them as much or as little as you'd like. The accounts just have to stay active (with the correct email and password) for the duration of the study.

2. Link to Survey 2

[Here](#) is the link to Survey 2, which will expire in 24 hours. This survey is not mandatory, but it takes just 3-5 minutes, and if you fill it out, we'll enter your name into our second raffle for a

Figure 5: Enrollment Confirmation Email

participants were given a paragraph describing the benefits of each behavior they were assigned treatment for.

The link to Survey 2 was included in the enrollment email. It first reminded participants of their treatment assignment, and then asked them how many days per week they “hoped” and “expected” to meditate and log their meals using the study apps. Finally, the enrollment email included Survey 1 raffle results and information about the study duration and how to withdraw. It reminded participants that at any time, they can opt out any SMS message program by replying STOP, without withdrawing from the study.

Table 10 shows the full set of possible messages. The first column contains all of the messages received by any participant assigned to m_x , and the second column contains all of the messages received by any participant assigned to m_y . Each message was sent twice throughout the program (except for messages 14 and 28, which were sent just once). The first 14 rows contain the informational messages, and the second 14 rows contain the reminder/encouragement messages. (A participant assigned to, say, m_x received 2 messages per day—one informational, one reminder—

over 27 days, so 54 total messages.) As mentioned in the paper, the two daily messages were sent in the morning (either 7am or 8am) and in the evening (either 7pm or 8pm). The timing of meditation vs. nutrition messages and information vs. reminder messages alternated in a balanced fashion as shown in Table 4. Messages were sent using the platform Slicktext.

The incentive treatment was described initially in the enrollment email as the following. “You will earn a green raffle ticket from eNOMerate for every day that you log at least one meal with FatSecret, and a red raffle ticket for every day that you don’t. To receive a ticket, you must log a meal on the day that you ate it. Every Sunday, for the duration of the program, we will let you know via email how many tickets you’ve accumulated. At the end, we will pull one of your tickets, and if it’s green, you will win a \$10 Amazon gift certificate. So if you log your meals every day, you will definitely get the gift certificate. If you log your meals half of the time, you will get it with 50% odds. And if you never log your meals, you definitely won’t get it. (This is separate from the raffles for survey completion.) The program will begin tomorrow and will last exactly 27 days.”

Each Sunday, participants in the incentive treatment received an email informing them of the total green and red tickets they had accumulated. At the end of the treatment period, they were sent a final email informing them of their total tickets, and then later sent the results of the raffle. Ultimately 52% of participants won the raffle.

At the end of the treatment period, all participants received an email informing them that any treatment programs they were in would now end, but that they should keep their app accounts intact with their assigned passwords for another four weeks, when they would receive a wrap-up email from us with a link to Survey 3.

After four weeks a final email was sent, concluding the study and providing a link to Survey 3. In Survey 3, we first ask how much they meditated without the assigned apps, about the timing of their meditation, and whether they felt like meditation came at the expense of any other activity. We then do the same for meal logging, with the additional question of how long it took them to log their meals each day. We then ask whether they set up any additional notifications for either behavior. Next, we ask questions about their mental health and diet. Finally, we administer an

	Meditation	Nutritional Monitoring
1	Evidence from 47 studies suggests that meditation reduces depression and anxiety! (Goyal et al. 2014)	Fact: more than 102 million American adults have high cholesterol, and 35 million are at risk for heart disease as a result (CDC 2013).
2	Did you know that meditation actually changes the physical structures of the brain (Fox et al. 2014)?	Did you know that potassium helps keep your blood pressure low and your heart healthy? The CDC recommends 4700mg of potassium daily for adults age 19-50.
3	Fun fact: for people with insomnia, meditation improves nightly sleep time, and helps people fall asleep faster! (Gross et al 2011)	37.7% of Americans reported that they consume fruits less than once per day! 22.6% report the same for vegetables (CDC 2013). Make sure it's not you!
4	Aetna, a Fortune 500 company, claims that its meditation program made employees more productive, saving \$3,000 per employee per year!	90% of Americans consume too much sodium (NHANES 2009-2012), which is a risk factor for heart disease! Many more foods have salt than you might expect!
5	Did you know that meditation programs combat depression almost as effectively as antidepressants? (Kuyken et al. 2008)	Over 15 years, people who consumed >25% of calories as added sugar were twice as likely to die from heart disease as those who consumed <10% (Yang et al. 2014)
6	Did you know that people can use meditation to reduce their physical pain? (Zeidan et al. 2011)	38% of U.S. adults are obese today, relative to 15% in 1980 (NHANES 2013-2014). Log your meals to keep track of your diet!
7	Fun fact: evidence suggests that meditation improves relationship satisfaction! (Sedlmeier et al. 2012)	Logging meals can help with weight loss (Burke et al. 2011)! And people are better at meal-logging when they use apps like [meal tracking app] (Wharton et al. 2014).
8	Meditation programs have been shown to reduce stress levels for people with high blood pressure! (Rainforth et al. 2008)	Less than 3% of Americans meet the daily recommended fiber intake (NHANES 2003-2006). Fiber can lower cholesterol and reduce the risk of heart disease
9	Fun fact: the part of the brain responsible for memory actually looks different in people who meditate! (Fox et al. 2014)	The American Heart Association says daily consumption of added sugar should be <25g for women and <38g for men. Yet the average American consumes 82g daily.
10	Did you know that General Mills runs 7-week meditation programs for its executives? Participants say they work more productively and make better decisions.	A host of studies suggest that nutrition is the most important factor in weight management – much more important than exercise (e.g. Johns et al. 2014).
11	Meditation has so many health benefits that today, 79% of medical schools offer some element of mindfulness training (Buchholz 2015)	Are you eating enough whole grains? Find out! Whole grains reduce the risk of diabetes; refined carbohydrates actually increase the risk! (AlEsa et al. 2015)
12	Did you know that 18.1% of adults in the U.S. experience some type of anxiety disorder? Meditation has proven to help! (Goyal et al. 2014)	Moderately active women between 21-40 should be consuming 2200-2000 calories per day (and men 2600-2800). Do you? Find out by tracking meals with [meal tracking app]!
13	Did you know that 35% of firms had mindfulness classes in 2017, and another 26% are considering them for the future (National Business Group on Health)?	>100 million Americans have diabetes or prediabetes (Nat'l Diabetes Stats Report 2017). Eating whole grains, and reducing sugar & trans fats, reduces the risk
14	Fun fact: meditation increases the thickness of your prefrontal cortex, the area of your brain associated with attention and self-awareness (Fox et al. 2014)	Fact: many companies are having their employees track their nutrition via smartphone apps as part of wellness programs. Jump on the bandwagon!
15	Greetings from Remindful! Try Tara Brach's Vipassana (Basic) meditation on [meditation app]!	Greetings from eNOMerate! Remember to log your meals today with [meal tracking app], if you haven't already!
16	Hello from Remindful! We hope you had a great day. Try Manoj Dias' Basic Breath Meditation on [meditation app]!	Hello from eNOMerate! We hope you had a great day. Take 5 minutes to log your meals with [meal tracking app]!
17	Hope you had a healthy, happy day from Remindful. You'll feel great if you end the day with some meditation! [meditation app] makes it easy.	Hope you had a healthy, happy day from eNOMerate. You'll feel great if you end the day by logging your meals! [meal tracking app] makes it easy.
18	Remindful wishes you a great evening! Remember to take care of yourself, and find a few minutes to meditate with [meditation app].	eNOMerate wishes you a great evening! Remember to take care of yourself, and find a few minutes to log your meals with [meal tracking app]!
19	Good evening from Remindful! You told us you were interested in meditation! So let's get on it. Try something new on [meditation app]!	Good evening from eNOMerate! You told us you were interested in monitoring your nutrition! So let's get on it. [meal tracking app] makes it simple!
20	Hi from Remindful! Are you meditating daily with [meditation app]? Keep the habit up!	Hi from eNOMerate! Are you logging your meals daily with [meal tracking app]? Keep the habit up!
21	Just another friendly hello, and reminder to meditate with [meditation app], from Remindful. Try the 3-minute breathing space by Mark Williams on [meditation app]!	Just another friendly hello, and reminder to log your meals with [meal tracking app], from eNOMerate! ;)
22	Greetings from Remindful! Remember to meditate today with [meditation app], if you haven't already!	Greetings from eNOMerate! Remember to log your meals today with [meal tracking app], if you haven't already!
23	Hello from Remindful! We hope you had a great day. Take 5 minutes to meditate with [meditation app]!	Hello from eNOMerate! We hope you had a great day. Take 5 minutes to log your meals with [meal tracking app]!
24	Hope you had a healthy, happy day from Remindful. You'll feel great if you end the day with some meditation! [meditation app] makes it easy.	Hope you had a healthy, happy day from eNOMerate. You'll feel great if you end the day by logging your meals! [meal tracking app] makes it easy.
25	Remindful wishes you a great evening! Remember to take care of yourself, and find a few minutes to meditate with [meditation app].	eNOMerate wishes you a great evening! Remember to take care of yourself, and find a few minutes to log your meals with [meal tracking app]!
26	Good evening from Remindful! You told us you were interested in meditation! So let's get on it. Try something new on [meditation app]!	Good evening from eNOMerate! You told us you were interested in monitoring your nutrition! So let's get on it. [meal tracking app] makes it simple!
27	Hi from Remindful! Are you meditating daily with [meditation app]? Keep the habit up!	Hi from eNOMerate! Are you logging your meals daily with [meal tracking app]? Keep the habit up!
28	Just another friendly hello, and reminder to meditate with [meditation app], from Remindful! ;)	Just another friendly hello, and reminder to log your meals with [meal tracking app], from eNOMerate! ;)

Table 10: Full Table of Messages

informational quiz, asking a true/false question about each informational message the participant received. At the end of Survey 3, participants were told to change their passwords for the two apps.

C Attrition and Survey Participation

In total 5,845 people filled out Survey 1, meaning that 66% of Survey 1 participants ultimately downloaded both apps and enrolled in the study. Of the 3,885 participants who enrolled, 40 ultimately dropped out, amounting to 1%, and resulting in a final sample of 3,845. Table 11 shows that there was no evidence of differential attrition by treatment.

	control	mx	my	mx & my	zy	F-test, joint sig
attrited	0.009	0.008	0.011	0.012	0.012	0.870
	0.092	0.088	0.105	0.109	0.110	

Notes: Means and standard deviations. F-test of joint significance reported in last column.

Table 11: Attrition Rates by Treatment

In terms of survey participation, of our 3,845 study participants, 2,891 completed Survey 2 (75.2%), and 2,145 completed Survey 3 (55.8%).

D Additional Baseline Data

Figure 6 shows the distribution of daily notifications, as self-reported in the baseline survey. Participants were asked to list all apps that send notifications across all devices, and then to estimate daily notifications for each app. The top-left plot shows total notifications, and the subsequent plots break notifications down by type. In Figure 7 I show the distribution of baseline responses to questions about self-reported importance, fun, and difficulty of each behavior, on a scale from 1 to 10. The most notable difference between the two behaviors is that participants believe that meditation will be more “fun” than meal logging. In the bottom-right plot I depict the self-reported experience with each behavior. The first comparison shows the fraction of participants who ever

did the behavior before, the second shows the fraction of participants who ever did the behavior *daily* before, and the third shows the fraction of participants who did the behavior in the last month.




Figure 6: Daily Notifications (after winsorizing at 99%)

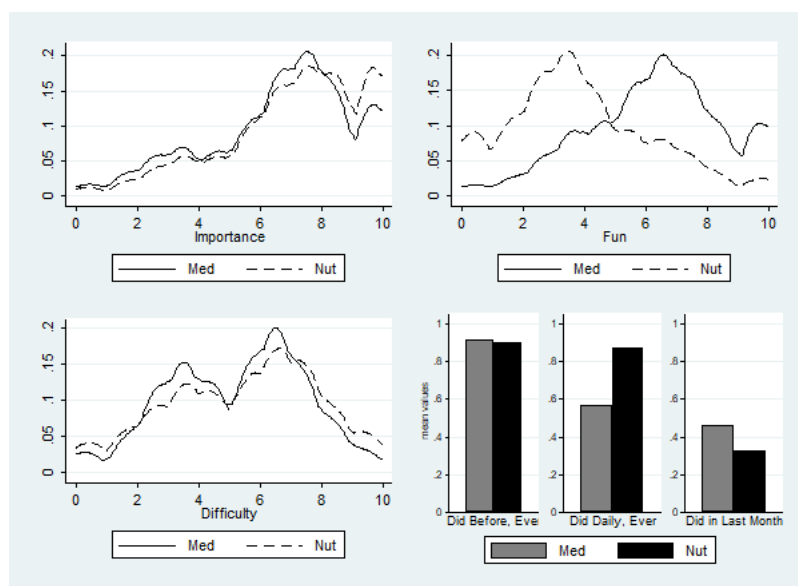


Figure 7: Preferences and Experience, Meditation & Meal Logging

E Reduced Form Results Reported as Treatment Effects

Table 12 contains the same results as Table 4, but depicts estimates for the four treatment groups as defined in the experiment design, rather than treatments and interactions.

	<i>Treatment Period</i>		<i>Post-Treatment Period</i>	
	Meditated (X)	Logged Meal (Y)	Meditated (X)	Logged Meal (Y)
	(1)	(2)	(3)	(4)
mx only	0.088*** (0.011)	-0.024** (0.010)	0.024*** (0.009)	-0.010* (0.006)
my only	-0.028*** (0.008)	0.166*** (0.013)	-0.025*** (0.007)	0.029*** (0.008)
mx & my	0.066*** (0.010)	0.116*** (0.012)	0.009 (0.008)	0.017** (0.007)
zy	-0.025*** (0.009)	0.381*** (0.016)	-0.024*** (0.007)	0.037*** (0.008)
mx only - mx & my	0.022 (0.011)	-0.140 (0.012)	0.016 (0.009)	-0.027 (0.007)
my only - mx & my	-0.094 (0.009)	0.050 (0.014)	-0.034 (0.007)	0.013 (0.008)
Ctrl Mean	0.094	0.118	0.054	0.033
Ctrl SD	0.291	0.323	0.227	0.178
Obs	102905	102905	102499	102499

Notes: OLS regressions with estimates reported as treatment effects, including controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month) as well as day fixed effects. Standard errors clustered at individual level. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 12: Reduced Form Results Reported as Treatment Effects

Figures 8 through 12 portray the raw data over the course of the treatment and post-treatment period.

F Estimation Robustness Checks and Sensitivity

In Table 13 I run the estimation using the identity matrix as the weighing matrix, instead of the diagonal of the inverse of the variance-covariance matrix as in the benchmark specification. The




Figure 8: Meditation Messages vs. Control

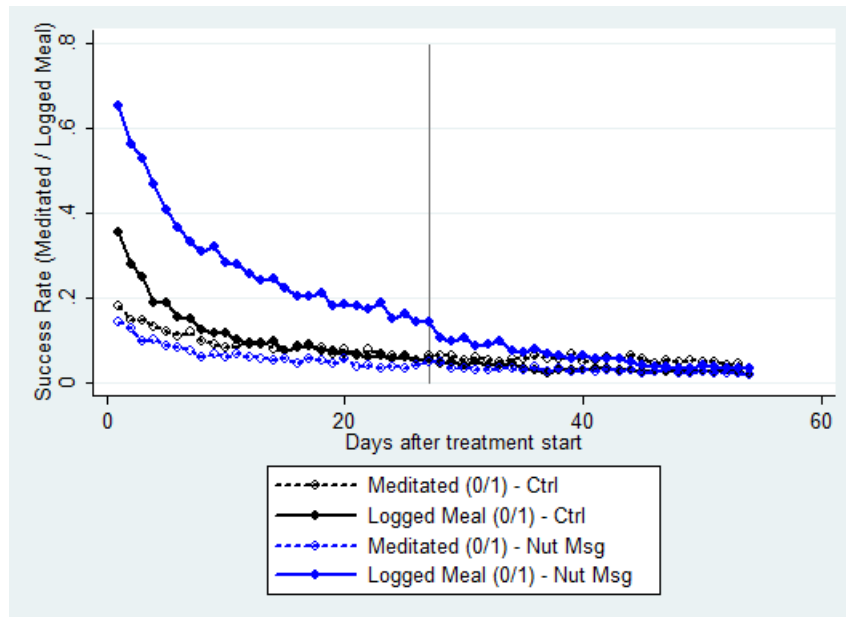


Figure 9: Nutrition Messages vs. Control




Figure 10: Meditation + Nutrition Messages vs. Meditation Messages Only




Figure 11: Meditation + Nutrition Messages vs. Nutrition Messages Only




Figure 12: Nutrition Incentives vs. Control

estimates differ in predictable ways, but all within 95% confidence intervals of the benchmark estimates, and the main conclusions are consistent. In fact, the estimates for θ_m and θ_z are even more negative, as is the estimate for ρ .

F.1 Sensitivity

In Figure 13 I plot the sensitivity matrix, as defined by [Andrews et al. \(2017\)](#), showing only the four key parameters. Each element of the matrix represents how a one percentage point increase in each moment affects each estimated parameter. We can see that ρ is most sensitive to the treatment effects of messages on the covariance, as well as to the main target effects, spillover effects, and interaction effects. As we would expect, θ_m and θ_z are sensitive to the same moments as ρ , but mostly in the opposite direction, since spillovers must be due to either diversion or depletion. (The bars go in the same direction because depletion is captured by $\rho > 0$ but diversion is captured by $\theta < 0$.) The only major exception is the effect of m_y on x , which pushes both θ_m and ρ toward zero. This is intuitive: the closer this spillover gets to zero, the smaller are both estimates of diversion and depletion. The effect of m_x on y pushes them in opposite directions (θ_m toward zero; ρ away

<i>Parameter</i>	<i>Description</i>	<i>Estimate</i>	<i>Standard Error</i>
α	slope of marginal cost of attention	9.569	0.802
ρ	depletion	-0.426	0.415
μ_x	return to x	1.149	0.130
ϕ_x	x message attn subsidy	0.925	0.162
γ	overload	0.097	0.196
θ_m	message diversion	-0.669	0.088
λ	y incentive attn subsidy	3.530	0.383
ϕ_y	y message attn subsidy	1.017	0.251
σ_{ϵ_x}	S.D. of ϵ_x , heterogeneity in return to x	3.824	0.343
σ_{ϵ_y}	S.D. of ϵ_y , heterogeneity in return to y	5.555	0.404
θ_z	incentive diversion	-0.552	0.167
σ_{xy}	covariance of ϵ_x and ϵ_y	-0.363	2.493
σ_{δ_x}	S.D. of δ_x , heterogeneity in message subsidy	4.242	0.684
σ_{δ_y}	S.D. of δ_y , heterogeneity in message subsidy	6.493	0.866
$\sigma_{\phi_x\phi_y}$	covariance of δ_x and δ_y	0.021	2.843

Table 13: Estimates using Identity Matrix

from zero) because this moment is being used to separate ρ from θ_m : the nearer to zero is this spillover, the smaller is the spillover-target ratio for x , bringing it closer to the spillover-target ratio for y and thus making ρ larger. As would be expected. γ is most sensitive to the message interaction moments, though it is sensitive to all of the other moments through ρ and α .

G Other Outcomes

G.1 Use of Other Apps

Table 14 shows the main results with an individual-level outcome variable that takes into account the use of other meditation and meal logging apps. At the final survey, we ask participants how many days they did the behaviors using other apps during the treatment and post-treatment period. Specifically, I inflate mean meditation and meal logging rates for the duration of the period according to the number of days in which other apps were reported to be used.

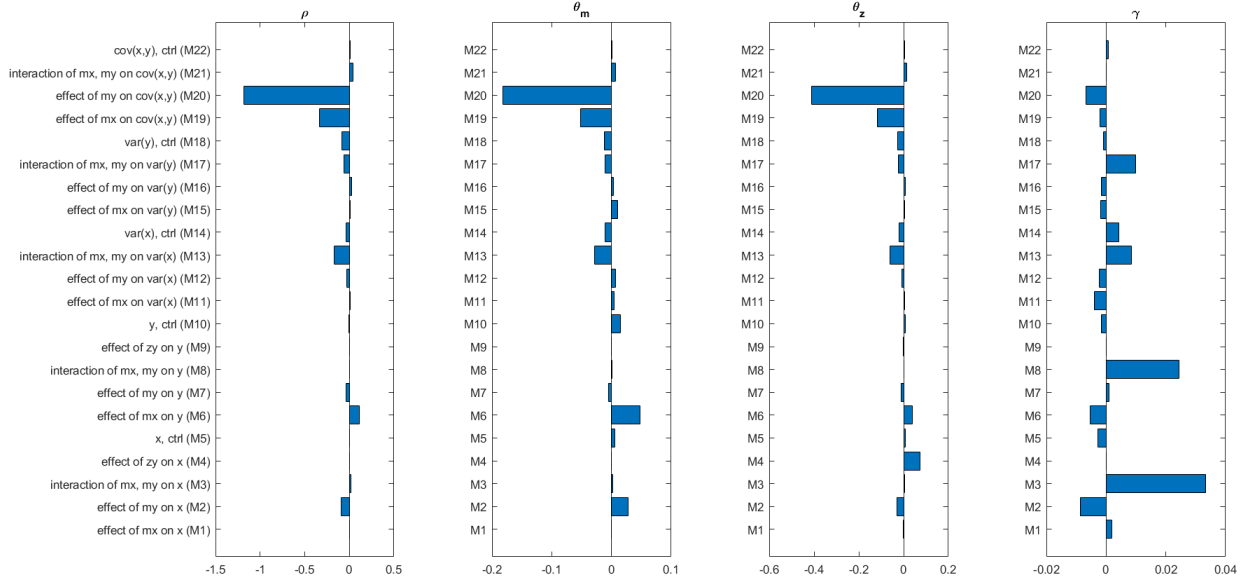


Figure 13: Sensitivity of Estimates to Moments

G.2 Health

Table 15 shows treatment effects on health. For meditation, I use total minutes meditated and a mental health index from the final survey. For meal logging, I use the fraction of the participant's weight goal that was achieved (self-reported) as well as a physical health index from the final survey.

H Heterogeneous Spillover Effects

H.1 Heterogeneity by Baseline Notifications

I look at two potential sources of heterogeneity in spillover effects. The first is by baseline notifications. Since the vast majority (92%) of notifications received by our sample participants are messages or updates, and not associated with a particular behavior, I simply modify the model to include a new type of message, m_w , but do not allow for any action associated with w (i.e. a_w). I make the same assumptions as before, again focus on interior solutions. The components of the attention cost function become: $f(a_x, a_y) = \frac{1}{2}\alpha(a_x^2 + a_y^2) + \rho a_x a_y$ and $s^x(m_x, z_x, m_y, z_y, m_w, z_w) =$

	<i>Treatment Period</i>		<i>Post-Treatment Period</i>	
	Meditated (x)	Logged Meal (y)	Meditated (x)	Logged Meal (y)
	(1)	(2)	(3)	(4)
mx	0.242*** (0.032)	-0.032 (0.037)	0.028* (0.014)	-0.005 (0.010)
my	-0.027 (0.029)	0.397*** (0.036)	-0.014 (0.011)	0.059*** (0.013)
mx X my	0.023 (0.046)	-0.084 (0.051)	0.015 (0.019)	-0.013 (0.018)
zy	-0.011 (0.030)	0.450*** (0.036)	-0.029** (0.011)	0.172*** (0.018)
mx + mxmy	0.265 (0.034)	-0.116 (0.036)	0.043 (0.013)	-0.018 (0.015)
my + mxmy	-0.003 (0.036)	0.313 (0.037)	0.001 (0.015)	0.046 (0.012)
Ctrl Mean	0.346	0.557	0.064	0.047
Ctrl SD	0.460	0.574	0.245	0.211
Obs	2131	2119	3805	3805

Notes: OLS regressions, including controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month). Outcome is mean individual rate of behavior over period, including self-reported use of other apps from final survey. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 14: Reduced Form Results, Including Use of Other Apps

$\nu_x m_x + \gamma m_x m_y + \gamma m_x m_w + \lambda z_x + \theta_m m_y + \theta_m m_w + \theta_z z_y$. I find the effect of m_w on spillovers to be:

$$\frac{\partial a_x}{\partial m_y} \Big|_{m_w=1} - \frac{\partial a_x}{\partial m_y} \Big|_{m_w=0} = \frac{-\rho\gamma}{\alpha^2 - \rho^2} \quad (6)$$

In the presence of both depletion ($\rho > 0$) and overload ($\gamma < 0$), notifications are predicted to push spillovers of m_y on x toward zero. The intuition is that distracting messages m_w interfere with m_y due to overload, causing them to generate less depletion, and thus less of a spillover, than they otherwise would. Since I found no strong evidence of either depletion or overload, I should not expect to see strong heterogeneous effects by baseline notifications. I do not attempt to predict the continuous effect of additional notifications, which would require many more assumptions.

	Avg Daily Min Meditated (1)	Standardized PHQ4 Score (2)	Standardized M. Health Score (3)	Fraction Weight Goal Achieved (4)	Standardized Diet Score (5)
mx	0.778** (0.288)	-0.010 (0.065)	0.079 (0.070)	-0.060 (0.057)	0.103 (0.065)
my	-0.861*** (0.225)	-0.007 (0.066)	0.100 (0.067)	0.012 (0.065)	0.224*** (0.070)
mx X my	0.528 (0.364)	-0.069 (0.092)	0.006 (0.101)	0.003 (0.078)	-0.120 (0.097)
zy	-0.818*** (0.234)	0.006 (0.068)	-0.058 (0.066)	-0.019 (0.055)	0.302*** (0.072)
mx + mxmy	1.306 (0.222)	-0.079 (0.066)	0.085 (0.072)	-0.057 (0.054)	-0.017 (0.072)
my + mxmy	-0.333 (0.286)	-0.076 (0.065)	0.106 (0.075)	0.015 (0.037)	0.104 (0.068)
Ctrl Mean	1.800	0.000	0.000	0.132	0.000
Ctrl Mean S.D.	5.633	1.000	1.000	0.964	1.000
Obs	3826	2145	2141	1651	2142

Notes: Health outcomes, including (1) average daily minutes meditated; (2) standardized score from the PHQ4, a four-item anxiety and depression questionnaire; (3) standardized response to "How would you describe your mental health now, relative to before you started the study?"; (4) fraction of weight goal achieved, and (5) standardized response to "How would you describe your diet now, relative to before you started the study?". Regressions include controls for the five baseline variables on which re-randomization was based. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 15: Treatment Effects on Health Outcomes

The most accurate test of the above prediction would be to look at people who have at least one other daily notification versus people who do not. Unfortunately only 12 participants have fewer than one other daily notification, and even cutting the data in half does not provide satisfactory power to detect small effects. Instead I opt for the test with the highest power, interacting each treatment with whether participants have notifications that are above or below the median. Table 16 shows the results. As expected, there is limited evidence of heterogeneity, though for all of the above reasons this is by no means a conclusive test.

H.2 Heterogeneity by Baseline Experience

I do not attempt to use the model to predict heterogeneity by baseline experience, as it is not clear how experience should enter the model. It likely reflects both preferences, which would enter

	Meditated (x) (1)	Logged Meal (y) (2)
mx	0.099*** (0.016)	-0.032* (0.015)
my	-0.043*** (0.012)	0.167*** (0.019)
mx X my	0.012 (0.021)	-0.005 (0.026)
zy	-0.028* (0.013)	0.397*** (0.024)
highnotif	-0.013 (0.014)	-0.012 (0.015)
mx X highnotif	-0.024 (0.021)	0.017 (0.021)
my X highnotif	0.031 (0.017)	-0.002 (0.026)
mx X my X highnotif	-0.010 (0.028)	-0.041 (0.035)
zy X highnotif	0.005 (0.018)	-0.031 (0.033)
mx + mxmy	0.111 (0.013)	-0.038 (0.021)
my + mxmy	-0.032 (0.017)	0.162 (0.018)
(mx + mxmy) X highnotif	-0.030 (0.020)	-0.020 (0.030)
(my + mxmy) X highnotif	0.020 (0.020)	-0.040 (0.020)
Ctrl Mean	0.094	0.118
Ctrl SD	0.291	0.323
Obs	102905	102905

Notes: OLS regressions of treatment-period behaviors on treatments and interactions with a binary measure of whether daily notifications are above or below the median. Includes controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month) as well as day fixed effects. Standard errors clustered at individual level. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 16: Heterogeneous Treatment Effects by Baseline Notifications

through u , as well as accumulated habits, which would likely affect the cost of attention. Instead, I try to run the simplest possible test of heterogeneity by baseline experience, in the hope that it might provide insight for future models.

I construct an experience score in which the participant gets 1 point if he/she has ever done the behavior, another point if he/she has attempted to do it daily before, and another point if he/she has done it in the last month, for a minimum score of zero and a maximum of three. Table 17 shows the results by whether participations are above or below the median in their experience with the outcome behavior in question. (In Column 1, experience represents meditation experience; in Column 2, experience represents meal logging experience.) I find no evidence of heterogeneity by experience, but again, the data is under-powered to detect small effects.

I Deviations from Pre-Analysis Plan

This study was registered at the AEA RCT Registry under the title "Nudges in Equilibrium" with RCT ID AEARCTR-0002435. In this section I describe any differences between the final paper and the pre-analysis plan.

Importantly, the main experiment design and sample size did not change substantially between the pre-analysis plan and the experiment. Slight differences in sample size across treatment groups are due to the fact that we randomized within cohorts using fixed proportions, and did not have full control over the total numbers. The slight rise in the total sample is also due to being unable to exactly control the size of the final cohort.

In terms of the analysis, the reduced form specifications are the same. One important difference is that ultimately we used meditation and meal logging *with* the assigned apps as the outcome in our main specification, rather than incorporating self-reports of meditation and meal-logging with other apps, as planned. The reason for this is twofold. First, it was actually a mistake to plan to use a measure that would include the action with and without the apps. The behavior we promoted in both message and incentive treatment was the behavior using the specified app, not the behavior generally, so it makes much more sense to use this as the outcome. Second, ultimately only 56% of participants completed Survey 3—much less than hoped—so attempting to incorpo-

	Meditated (x) (1)	Logged Meal (y) (2)
mx	0.093*** (0.021)	-0.021 (0.026)
my	-0.002 (0.014)	0.133*** (0.030)
mx X my	-0.045 (0.027)	0.015 (0.044)
zy	-0.009 (0.015)	0.358*** (0.045)
experience	-0.008 (0.011)	0.004 (0.028)
mx X experience	-0.003 (0.011)	-0.001 (0.012)
my X experience	-0.013 (0.008)	0.016 (0.014)
mx X my X experience	0.026 (0.015)	-0.020 (0.020)
zy X experience	-0.008 (0.009)	0.011 (0.020)
mx + mxmy	0.048 (0.017)	-0.006 (0.036)
my + mxmy	-0.047 (0.023)	0.148 (0.033)
(mx + mxmy) X experience	0.020 (0.010)	-0.020 (0.020)
(my + mxmy) X experience	0.010 (0.010)	0.000 (0.010)
Ctrl Mean	0.094	0.118
Ctrl SD	0.291	0.323
Obs	102905	102905

Notes: OLS regressions of treatment-period behaviors on treatments and interactions with a binary measure of whether baseline experience in the outcome behavior was above or below the median. The experience measure went from 0 to 3, where participants earned 1 point for having ever done it before, 1 point for having done it daily before, and 1 point for having done it in the last month. Includes controls for the five baseline variables on which re-randomization was based (female, college, daily notifications, whether individual meditated in last month, whether individual logged meal in last month) as well as day fixed effects. Standard errors clustered at individual level. One, two, and three stars indicate q-values of 1%, 5%, and 10% respectively; q-values calculated according to the Benjamini Hochberg step-down procedure, considering all tests in the table (but excluding linear combinations of coefficients).

Table 17: Heterogeneous Effects by Baseline Experience in Outcome Behavior

rate self-reports from this survey would have reduced our power significantly. Table 14 shows the results of the specification stated in the pre-analysis plan. The key coefficients of interest are not substantially different, but there is insufficient power to draw the same conclusions.

I do not include in the paper all of the sub-group analyses as described in the pre-analysis plan, since they are generally insufficiently powered. I also changed the measurement tool for mental health, substituting the PHQ4 for the General Well-Being Schedule, since feedback from the first participants suggested that the 18-item questionnaire was too annoying, and was reducing the likelihood of completion.

Finally, the model changed in several ways since the pre-analysis plan. The most important change was to ultimately include three types of limited attention (overload, diversion, and depletion), instead of just two (overload and depletion, originally called limited external and internal attention). I made this change because after getting the data, I realized that the original model simply did not fit because it did not include diversion, which is what I ultimately find to be the main driver of spillovers. As a result of this change, many smaller, subsequent changes had to be made to the model as well, which explains the rest of the differences.