

Is Attention Produced Optimally?

Theory and Evidence from Experiments with Bandwidth

Enhancements*

Erin T. Bronchetti, Judd B. Kessler, Ellen B. Magenheim

Dmitry Taubinsky, Eric Zwick

June 2022

Abstract

This paper investigates whether people optimally value tools that reduce attention costs. We call these tools bandwidth enhancements (BEs) and characterize how demand for BEs varies with the pecuniary incentives to be attentive, under the null hypothesis of correct perceptions and optimal choice. We examine if the optimality conditions are satisfied in three experiments. The first is a field experiment ($n = 1373$) with an online education platform, in which we randomize incentives to complete course modules and incentives to utilize a plan-making tool to complete the modules. In the second experiment ($n = 2306$), participants must complete a survey in the future. We randomize survey-completion incentives and how long participants must wait to complete the survey, and we elicit willingness to pay for reminders. The third experiment ($n = 1465$) involves a psychometric task in which participants must identify whether there are more correct or incorrect mathematical equations in an image. We vary incentives for accuracy, elicit willingness to pay to reduce task difficulty, and examine the impact of learning and feedback. In all experiments, demand for reducing attention costs increases as incentives for accurate task completion increase. However, in all experiments—and across all conditions—our tests imply that this increase in demand is too small relative to the null of correct perceptions. These results suggest that people may be uncertain or systematically biased about their attention cost functions, and that experience and feedback do not necessarily eliminate bias.

*Bronchetti and Magenheim: Swarthmore College. Kessler: University of Pennsylvania and NBER. Taubinsky: UC Berkeley and NBER. Zwick: University of Chicago Booth and NBER. We thank Andrew Caplin, Mark Dean, Xavier Gabaix, Stephen O’Connell, Devin Pope, three grant reviewers at the Russell Sage Foundation, and seminar and conference participants for helpful comments and advice. We thank Alexander Hirsch, Stephanie Nam, Laila Voss, and Caleb Wroblewski for excellent research assistance. We gratefully acknowledge Mike Walmsley and CodeAvengers.com for their support with the education experiment. We gratefully acknowledge research funding from the Russell Sage Foundation, Swarthmore College, the Boettner Center, the Wharton School, the Wharton Behavioral Lab, and the Alfred P. Sloan Foundation. The first experiment was approved by the Swarthmore (covering Haverford and Muhlenberg), Bryn Mawr, Lafayette, and Ursinus IRBs, numbers: 14-15-065, R17-042, AY1617-12, 01-18-17. The second and third experiments were approved by the University of Pennsylvania IRB, number 832335. The opinions expressed in this paper are solely the authors’, and do not necessarily reflect the views of any individual or institution listed above.

A large and rapidly growing body of work in economics and cognitive science shows that attention is costly and that people at least partly control it (for recent reviews in economics, see Caplin, 2016; Maćkowiak et al., forthcoming; Gabaix, 2019). There is a growing recognition that, like many other types of costly effort decisions that economists have analyzed for decades, mental effort is also costly and deliberately deployed (Shenhav et al., 2017). A powerful modeling approach taken by the *rational inattention* literature is to assume that individuals choose their attention strategies optimally.

In this paper, we move beyond the question of whether people optimize attention within an environment to investigate whether people optimally choose their attention environment. Specifically, we study whether individuals optimally invest in tools that reduce attention costs. This is a fundamental question for three reasons. First, decisions over attention environments are pervasive. For example, people can choose whether to set reminders for themselves to reduce the cognitive costs of keeping opportunities top-of-mind; how much to avoid distraction in their work environment to reduce the cost of focusing mental effort; whether to avoid sellers with opaque and hard-to-calculate fees; and whether to obtain outside advice for complex, cognitively demanding tasks such as managing an investment portfolio, selling a home, or filing taxes. Second, because any dynamic model of costly attention (e.g., Sims, 2003; Matejka et al., 2017; Mackowiak et al., 2018) must be closed with an assumption about people’s beliefs about their attention cost functions, empirical evidence on whether people are uncertain or systematically biased about their attention costs is a critical input into such modeling.¹ Third, if individuals fail to choose an attention environment optimally, this failure raises the possibility that they may fail to optimize their attention strategies even in static environments, due to uncertainty or systematic misperceptions about their attention cost functions.

The first contribution of this paper is a methodology for testing whether people properly value tools to reduce demands on their *mental bandwidth*, i.e., the cost of being attentive. To fix terminology, we call these tools *bandwidth enhancements* (BEs). We then deploy this methodology in three complementary experiments. The first two experiments focus on people’s ability to remain attentive to a future task. The third experiment explores an attention-demanding psychometric task, as in recent experiments on rational inattention models (Dean and Neligh, 2018; Caplin et al., 2020; Ambuehl et al., 2020). The BEs we study are a plan-making tool in the first experiment, a reminder tool in the second experiment, and making a

¹See also de Oliveira et al. (2017), who provide an axiomatic treatment of rational inattention models by studying choices over menus. Such a characterization mechanically requires the assumption that people optimally choose between attention environments.

psychometric task easier in the third experiment.

Our theoretical approach builds on the Caplin et al. (2020) characterization of costly attention models with a competitive supply framework and the insights in Gabaix (2014) that link bounded rationality and Slutsky symmetry. The main idea of our approach is to examine how pecuniary incentives to complete a task affect demand for BEs that aid task completion. Our approach clarifies the difficulty with assessing whether people optimally value attention-cost reductions by just examining willingness to pay for BEs and how BEs affect behavior, as individuals may particularly like or dislike a given BE for reasons unrelated to its impact on behavior, such as the nuisance of additional reminders. The first prediction from our model is a precise condition on how willingness to pay for BEs changes with the pecuniary rewards for the task, under the null of correct perceptions. The second prediction is a form of a Slutsky symmetry condition, which states that the impact of task-completion incentives on take-up of the BE is equal to the impact of the price of the BE on the propensity to complete the task.² Our model is sufficiently general to apply to both dynamic and static settings, as well as under various assumptions about the set of feasible attention strategies, ranging from the flexible strategy spaces in rational inattention models to much more restrictive assumptions about how people choose attention. We obtain sharp results with minimal assumptions by using the generalized envelope theorems of Milgrom and Segal (2002).

Guided by this framework, we carry out three experiments. The first experiment was run in the field with 1373 students and alumni from six Philadelphia-area colleges who enrolled in an 8-week online coding course. The experiment randomized incentives to complete three 15-minute coding lessons each week and randomized incentives to make a plan to complete three 15-minute course modules each week. Making a plan, the BE in this experiment, involved clicking a link that automatically created three 15-minute events in the participant’s online calendar of choice for the following week and allowed the participant to rearrange the planned events to suit their schedule.

We document three key facts in our first experiment. First, use of our plan-making tool increased the likelihood of completing coding lessons, especially in the initial weeks. Second, take-up of our plan-making tool was elastic to the direct incentives for plan-making but remained below 100 percent, even with the incentives. The combination of incomplete take-up and the positive elasticity suggests that the use of our tool imposes internal or “nuisance” costs on at least some individuals. Third, we find that take-up of our plan-making

²See Gabaix (2019) for a discussion about exploiting Slutsky symmetry as an empirical strategy for testing limited attention. See Gabaix (2014) and Abaluck and Adams-Prassl (2021) for an implementation of such a test in the context of misperceived product attributes.

tool increased with incentives for *completing coding lessons*.

This third result is consistent with the *qualitative* prediction that optimizing individuals should value BEs more as the rewards for completing a task increase. At the same time, we estimate that the impact of completion incentives on plan-making is quantitatively too small relative to the Slutsky symmetry condition, suggesting that participants undervalue the plan-making tool. However, our confidence intervals are wide and do not permit us to reject the null of full optimality under correct perceptions.

Our second experiment is an online survey-completion experiment that elicits richer data that allows us to fully quantify the demand for a reminder technology and to test the first prediction from our model. The study was conducted on Amazon Mechanical Turk (MTurk) with 2306 participants. Study participants were offered a bonus (either \$3, \$4, \$11, or \$12) for completing a survey that would only be accessible for a week-long period after a delay (either 2 days, 1 week, 3 weeks, or 6 weeks). Prior to randomizing participants into one of the four possible delays and one of the four possible survey-completion rewards, we elicited participants' willingness to pay (WTP) for a set of three reminder emails, the BE in this experiment, for each possible delay-reward pair. Our procedure also generated exogenous variation in whether participants actually received the reminder emails, allowing us to estimate the effect of reminder emails on survey completion.

We find that survey completion increased with incentives and decreased with delay, while the impact of reminders decreased with incentives and increased with delay. The average impact of reminders on completing the survey was 29 and 16 percentage points for *low* (\$3 or \$4) and *high* (\$11 or \$12) task-completion incentives, respectively. Across the eight different possible delay-reward pairs, the impact of reminders ranged from -7 percentage points ($se = 6.00$) to 40 percentage points ($se = 5.88$).

This set of findings informs several hypotheses about attention in this type of setting. First, the negative effect of delay on task completion and the positive effect of delay on the impact of reminders is consistent with the attention/memory decay curves proposed by cognitive psychologists (see, e.g., Mullainathan, 2002; Ericson, 2017 for reviews). Second, at high incentives, task completion was higher but the impact of reminders was lower, which is consistent with individuals exerting more costly attention to keep the task top-of-mind when stakes are higher.

Our second set of findings from our second experiment is that, while WTP for reminders increased with the size of the bonus for survey completion, the increase was too small relative to the null of correct perceptions of attention costs. The theory implies that under this null, an extra \$1 of task-completion incentives should increase WTP by approximately \$0.29 and

\$0.16 in the *low* and *high* incentive groups, respectively.³ Instead, WTP increased by \$0.07 ($se = 0.017$) and \$0.02 ($se = 0.047$) in those two groups. Using the variation in WTP and the effects of reminders across all eight conditions, we estimate a model of how people’s perceived effects of reminders vary with the actual effects. We find that people uniformly underestimate the effects of reminders by 84 percent, rather than underestimating the effects when they are small and overestimating them when they are large.

Our third experiment, conducted on Prolific Academic with 1465 participants, expands the scope of our analysis in two ways. First, it involves a different domain of decision-making, illustrating the breadth and portability of our methods. Second, it allows us to study how learning and feedback affect people’s perceptions of the value of BEs. Specifically, this experiment involved a series of tasks, first utilized by Ambuehl et al. (2020), in which participants were shown an image with a set of arithmetic equations that were either correct (e.g., $10 + 12 = 22$) or incorrect (e.g., $10 + 12 = 23$). Participants were asked to indicate whether an image contained more correct or more incorrect equations, and were rewarded for accuracy. All participants completed baseline tasks, in which the image had 100 equations and either 60 or 40 percent of equations were correct.

We randomly assigned participants to also do one of two less cognitively demanding versions of the task, the BEs in this experiment. In the *length* condition, the easier task involved only 10 equations. In the *discernibility* condition, the easier task had either 95 or 5 correct equations. Participants completed two blocks of seven tasks, with each block containing three baseline tasks, three easier tasks, and one task that might be affected by the participants’ preferences. Analogous to the second experiment, we varied incentives for successful task completion (i.e., accuracy in this case) and, prior to each block, we elicited individuals’ WTP to make the remaining task easier for the different incentive levels. Additionally, prior to the WTP elicitation in the second block, we randomly gave some participants feedback about their performance in the first block on the baseline and easier tasks.

We find that participants were 19 ($se = 0.647$) and 26 ($se = 0.607$) percentage points more likely to accurately answer the easier tasks in the *length* and *discernibility* conditions, respectively, and these differences were nearly identical across both blocks. Under the null of correct perceptions, participants’ WTP to decrease difficulty in the *length* and *discernibility* conditions should thus increase by approximately \$0.19 and \$0.26 with each extra dollar of accuracy incentives.⁴ However, we find that in the first block, the WTP increases are only \$0.10 ($se = 0.034$) and \$−0.01 ($se = 0.030$) in the *length* and *discernibility* conditions,

³This approximation is valid if an additional \$1 has negligible effects on behavior.

⁴See footnote 3.

respectively. Thus, as in the other experiments, participants undervalue the BEs.

Our data from this experiment provide two additional lessons that complement the insights from the first two experiments. First, in the first block, participants' valuations of the discernibility BE are significantly more biased than their valuations of the length BE. This fact illustrates the potential context-specificity of BE valuations, and the need for portable methods that can quantify them across contexts. Second, we study the effects of learning and feedback. We find that in the second block, participants' WTP to make the task easier increases by only \$0.03 ($se = 0.025$) and \$0.03 ($se = 0.024$) for each extra dollar of accuracy incentives in the length and discernibility conditions, respectively. In the length condition, this result reflects a significant *decrease* relative to the first block. We show that this decrease is concentrated among people who received feedback that they performed at least as well in the longer version of the task than the shorter version, which suggests that—on average—people overweighted experiences in which the length-decreasing BE did not improve performance. This pattern is consistent with recent work on misspecified learning (e.g., Heidhues et al., 2018; Gagnon-Bartsch et al., 2021), which suggests that experience and feedback do not necessarily eliminate mistakes.

Our results contribute to the literature in several ways. First, we build on the supply theory framework developed by Caplin et al. (2020) to develop methods for assessing whether individuals optimally choose their attention environment. Broadly speaking, models of rational inattention—particularly when applied to dynamic environments—assume that individuals know their attention cost functions and thus will optimally invest in BEs. Thus, our method allows researchers to test key assumptions in models of rational inattention. Despite the recent proliferation of work on rational inattention, surprisingly little work in the economics literature has been devoted to individuals' understanding of the limitations of their attention.⁵ Our experiments illustrate how our method can be applied both to the kinds of psychometric settings where rational inattention models have traditionally been tested, as well as to settings concerning behaviors such as education and health investments where the study of attention has been more reduced-form.

⁵There is more work in economics on the optimality of individuals' information acquisition strategies. See, e.g., Gabaix et al. (2006); Hanna et al. (2014); Bartoš et al. (2016); Martin (2016); Dean and Neligh (2018); Ambuehl et al. (2020); Caplin et al. (2020); Morrison and Taubinsky (forthcoming). There is also a large psychological literature on metacognition, including work on how individuals act on their environments to create external triggers (e.g., calendar events or reminders) for delayed intentions (Gilbert, 2015a,b). This work finds that whether individuals utilize such tools depends on what they perceive to be the required internal cognitive demands that would otherwise be necessary, as well as the expected value of achieving the goal (Shenhav et al., 2013). These evaluations may be erroneous (Gilbert et al., 2020), leading to suboptimal decisions. We provide a quantitative toolbox for exact quantitative tests of whether people value BEs optimally, which can help advance the more qualitative psychological work on metacognition.

Second, a large body of work looks at the impact of BEs, such as planning prompts and reminders, on behaviors such as medical compliance, educational attainment, savings, loan repayment, wage reporting, voting, and charitable donation.⁶ We advance this literature by studying individuals' demand for BEs. Our approach sheds light on whether provision of such BEs is efficient. If individuals valued these BEs optimally, then external provision of the BEs would be inefficient because the market already provides individuals with many opportunities to acquire reminder technologies and plan-making tools in the form of various smartphone and computer applications, online calendars, smart caps on pill bottles, and so on. As noted above, incomplete take-up of BEs does not by itself imply that people undervalue them, because in addition to any pecuniary costs, reminders and plan-making tools may carry private nuisance costs (Damgaard and Gravert, 2018), time costs, or detract scarce attention from other important tasks (Nafziger, 2020; Altmann et al., forthcoming).

Closest to our second experiment, Ericson (2011) and Tasoff and Letzler (2014) conduct lab experiments and find that individuals' willingness to pay for a rebate exceeds the expected returns. Their results suggest overestimation of future attention to the rebate and thus follow-through, although other biases, such as Tasoff and Letzler's (2014) proposed *weak cost-salience*, plausibly also play a role. Our approach and results from the second experiment complement Ericson (2011) and Tasoff and Letzler (2014) in a few ways. First, overconfidence about one's baseline level of attention need not imply under-appreciation of the incremental impact of BEs, and vice versa. Second, by directly estimating individuals' (mis)valuations of BEs, our method allows us to directly speak to how much take-up of BEs should be encouraged through subsidies or other interventions. Third, the richness of our second experiment provides new insights about variation in attention, such as our result that people are more attentive at higher stakes.

Last, our work relates to the broader literature that studies whether individuals' beliefs are well-calibrated. A common approach is to directly elicit individuals' beliefs. However, the beliefs that individuals state in an abstract elicitation are not necessarily the decision weights that individuals apply in all real-stakes decisions because of salience and context effects (see, e.g., Bernheim and Taubinsky, 2018 for a recent discussion). For example, many individuals know how large sales taxes are and what products they apply to, but still neglect to incorporate them into their decisions (Chetty et al., 2009; Taubinsky and Rees-Jones, 2018). Similarly, individuals might have an abstract understanding of how BEs affect

⁶See, e.g., Nickerson and Rogers (2010); Milkman et al. (2011); Altmann and Traxler (2014); Castleman and Page (2016); Bronchetti et al. (2015); Karlan et al. (2016a); Calzolari and Nardotto (2017); Damgaard and Gravert (2018); Marx and Turner (2019); Zhang et al. (2021). See also Carrera et al. (2018) and Oreopoulos et al. (forthcoming) for examples of null effects.

behavior yet still undervalue them in real-stakes decisions. Our approach is thus a useful complement to this other work.

Methodologically, this paper also contributes to recent work that tests for behavioral biases by measuring consumer surplus using two different approaches: (i) inferring it from how behavior responds to incentives and (ii) directly eliciting consumers' WTP to engage in the behavior at different incentives. In different domains where bounded self-control rather than bounded rationality is plausibly implicated, Allcott et al. (2022) and Carrera et al. (2022) follow this strategy to estimate the degree of time inconsistency.⁷

The paper proceeds as follows. Section 1 presents our theoretical framework. Sections 2-4 present the designs and results from our three experiments. Section 5 concludes.

1 Theoretical Framework

1.1 Attention Strategies and Payoffs

We consider individuals who are faced with a task, the successful completion of which requires both attentional inputs and possibly other auxiliary inputs. For example, in our first two experiments, individuals first choose attention strategies that determine their likelihood of being attentive to the task in the future, and, conditional on being attentive to the task, choose whether or not to provide auxiliary inputs to complete the task. In our third experiment, individuals' choice of attention strategy affects their likelihood of correctly solving a cognitively demanding task, and there are no auxiliary actions.

Formally, let S_a denote the set of possible attention strategies, with generic element s_a , and let S_o denote the set of strategies over auxiliary actions, with generic element s_o . Let $z = 1$ indicate that the task is completed successfully, with $z = 0$ otherwise. The likelihood of $z = 1$ is given by $Q(s_a, s_o, \omega)$ where $\omega \in \Omega$ is a state of the world drawn from a prior μ . Individual i 's utility function is given by $U_i = rz - K_i(s_a, s_o, \omega)$, where r is the financial reward for successfully completing the task, and K_i is the net utility cost of choosing (s_a, s_o) in state ω .⁸

To ease exposition, we assume that the state is given by $\omega = (\omega_a, \omega_o) \in \Omega_a \times \Omega_o$ and that $Q(s_a, s_o, \omega) = Q_a(s_a, \omega_a) \cdot Q_o(s_o, \omega_o)$. For example, we conceptualize our first two ex-

⁷See also DellaVigna and Malmendier (2004) and Acland and Levy (2015) for early precursors to this strategy in the time-inconsistency context. See Strack and Taubinsky (2022) for formal results about the robustness of this strategy for measuring limited self-control.

⁸Note that any non-pecuniary benefits from completing the task can be incorporated into the net utility cost function K .

periments’ settings where individuals take actions to maintain attention to the task, with Q_a giving the probability that individuals are attentive to the task. If individuals are not attentive to the task, they cannot complete it. If they are attentive to the task, they can take auxiliary actions to complete the task. Thus, $Q(s_a, s_o)$ is multiplicatively separable. We also assume that K_i is additively separable, given by $K_i = K_{ai}(s_a) + K_{oi}(s_o, \omega_o)$, where it is without loss of generality to assume that K_{ai} does not depend on the state. The general framework covers a variety of settings where attention is implicated, as the examples below illustrate.

Dynamic decisions and sustained attention Suppose that in periods $t = 1, \dots, T$, individuals realize an attention outcome α_t , and choose an attention action a_t . Let $h_t = (\alpha_1, \dots, \alpha_{t-1}, a_1, \dots, a_{t-1})$ denote the period- t history, and let $A_t(h_t, \alpha_t)$ denote the period- t set of available actions, allowing for the possibility that α_t is itself endogenous to h_t (and to an underlying state ω_a). Then, strategies s_a are feasible plans for a choice of a_t after each realized history. For example, let $\alpha_t \in \{0, 1\}$ encode whether an individual is attentive to the task at the beginning of period t , so that $Q_a = Pr(\alpha_T = 1)$. At each point in time t , suppose that $Pr(\alpha_t = 1)$ is a function of the history h_t , and suppose that A_t is non-empty and time-invariant if $\alpha_t = 1$, and $A_t = \emptyset$ otherwise. That is, as in Ericson (2017) and Taubinsky (2014), individuals can take actions to affect their future attention to the task if they are presently attentive to it, but if they forget about the task, then they cannot take action. The attention-sustaining actions might involve setting reminders for oneself, asking others for reminders, or engaging in internal “rehearsal” (e.g., Mullainathan, 2002).

If the individual is inattentive to the task in period T , then she does not complete it; otherwise, an auxiliary action s_o determines whether the individual completes the task. Suppose that the cost of completing the task is given by ω_o , which is observed in period T (if the individual is attentive), so that an individual’s utility from completing the task is $r - \omega_o$. Strategies s_o are then functions $s_o : \omega_o \rightarrow \{0, 1\}$, where 1 is an indicator for choosing to complete the task conditional on being attentive. The individual’s optimal strategy s_o^* is then a cutoff rule, where $s_o^*(\omega_o) = 1$ if and only if $r \geq \omega_o$.

Rational inattention in cognitively demanding tasks Consider a cognitively demanding task, like the kind employed in experiments testing rational inattention (Dean and Neligh, 2018; Caplin et al., 2020; Caplin, 2021), where the person must allocate mental effort to identify the state $\omega \in \{1, \dots, N\}$. The possible actions are messages $m \in M = \{1, \dots, N\}$, and $z = 1$ if and only if $m = \omega$. The decision maker can receive signals γ from a set of cardinality at least N , and an attention strategy is any joint distribution s over signals and

states that are consistent with the prior μ , so that $\int_{\gamma} s(\gamma, \omega) d\gamma = \mu(\omega)$. The cost of the information strategy is $K_a(s)$, which for tractability is sometimes assumed to be proportional to the expected mutual information between the state and the signal. The message m sent by the decision maker is the most likely state given the realized posterior.

As shown by Matějka and McKay (2015) and others, such rational inattention models can be equivalently reformulated such that a feasible strategy is any joint distribution over actions and states that is consistent with the prior. That is, S_a is the set of all probability distributions s over $\Omega \times M$ such that $\sum_m s(m, \omega) = \mu(\omega)$ for each ω . Under this definition, $Q(s_a, \omega) = s(m, \omega) / \mu(\omega)$.

1.2 A Simplifying Restatement

Building on Caplin et al. (2020), we perform a change-of-variable to reduce the dimensionality of the individual's optimization problem, which leads to a particularly straightforward interpretation of attention costs.

In the framework we have presented thus far, individual i solves the optimization problem

$$\max_{(s_a, s_o) \in S_a \times S_o} \mathbb{E} [rQ_a(s_a, \omega_a)Q_o(s_o, \omega_o) - K_{ai}(s_a) - K_{oi}(s_o, \omega_o)] \quad (1)$$

where the expectation is taken with respect to the prior μ . Define $\underline{q}_a := \inf_{s_a} \mathbb{E}Q_a(s_a, \omega_a)$ and $\bar{q}_a := \sup_{s_o} \mathbb{E}Q_a(s_a, \omega_a)$, and define \underline{q}_o and \bar{q}_o analogously. Define $\bar{K}_{ai}(q_a) = \inf_{s_a} \{K_{ai}(s_a) | \mathbb{E}Q_a(s_a, \omega_a) \geq q_a\}$ and $\bar{K}_{oi}(q_o) = \inf_{s_o} \{\mathbb{E}K_{oi}(s_o, \omega_o) | \mathbb{E}Q_o(s_o, \omega_o) \geq q_o\}$. Then the optimization problem in (1) is equivalent to

$$\max_{(q_a, q_o) \in [\underline{q}_a, \bar{q}_a] \times [\underline{q}_o, \bar{q}_o]} [rq_aq_o - \bar{K}_{ai}(q_a) - \bar{K}_{oi}(q_o)]. \quad (2)$$

We formalize the notion of equivalence in Lemma 1 below.

Lemma 1. *Suppose that (s_a^*, s_o^*) is a solution to (1). Then $(q_a^* = \mathbb{E}Q(s_a^*, \omega_a), q_o^* = \mathbb{E}Q(s_o^*, \omega_o))$ is a solution to (2). Conversely, if (q_a^*, q_o^*) is a solution to (2) then there exists (s_a^*, s_o^*) that is a solution to (1), with $q_a^* = \mathbb{E}Q(s_a^*, \omega_a)$ and $q_o^* = \mathbb{E}Q(s_o^*, \omega_o)$.*

In other words, we can reformulate the individual's decisions as a two-dimensional optimization problem, with a one-dimensional attentional input q_a and a one-dimensional auxiliary input q_o , at costs $\bar{K}_{ai}(q_a)$ and $\bar{K}_{oi}(q_o)$, respectively. Lemma 1 shows that the functions $\bar{K}_{ai}, \bar{K}_{oi}$ are sufficient statistics for an individual's surplus: there are many different economic

environments that can generate the same aggregate cost functions \bar{K}_{ai} and \bar{K}_{oi} , and an individual's surplus depends only on these aggregate cost functions. This result allows us to focus all discussion below on \bar{K}_{ai} and \bar{K}_{oi} , and to omit other details of the attention process.

We refer to \bar{K}_{ai} as the attention production function, which parallels standard models of competitive supply (Caplin et al., 2020). For example, if the aggregate cost functions are differentiable, then the optimal choice of (q_a, q_o) equates marginal benefits and marginal costs, so that $\bar{K}'_{ai}(q_a) = rq_o$ and $\bar{K}'_{oi}(q_o) = rq_a$.

We make one regularity assumption, which is that there exists a solution to (1), and therefore (2). We do not make additional assumptions about differentiability, continuity, or convexity, but Lemma 2 below shows that the aggregate cost functions \bar{K}_{ai} and \bar{K}_{oi} will be convex whenever the individual's objective function in (1) is concave. For example, in the special case of rational inattention in cognitively demanding tasks, the individual's optimization problem is concave under the common assumption that attention costs are proportional to mutual information (e.g., Sims, 2003; Matějka and McKay, 2015). In Appendix Lemma A.1, we show that convexity of the cost functions is sufficient to ensure differentiability of many statistics of interest, which facilitates the first-order conditions in Theorem 1 in the next subsection.

Lemma 2. *Suppose that $S_a \times S_o$ is a convex subset of \mathbb{R}^n . $\bar{K}_{ai}(q_a)$ is strictly convex in q_a if $\mathbb{E}Q_a(s_a, \omega_a)$ is concave in s_a and $K_{ai}(s_a)$ is convex in s_a , with at least one of these strictly so. $\bar{K}_{oi}(q_o)$ is strictly convex in q_o if $\mathbb{E}Q_o(s_o, \omega_o)$ is concave in s_o and $\mathbb{E}K_{oi}(s_o)$ is convex in s_o , with at least one of these strictly so.*

1.3 Choice of Attention Technology

We now introduce an initial choice of whether to simplify the attentional demands required to successfully complete the task. In our first two experiments, this choice involves planning prompts and reminders, respectively. In our third experiment, this is a choice of whether or not to make the task less cognitively demanding. In addition to our specific experimental settings, this formalism can also apply to settings where people exert mental effort under a piece-rate incentive scheme (e.g., Dean, 2019; Kaur et al., 2021; Bessone et al., forthcoming), and may choose task difficulty, decision aids, or the level of distraction in the environment.

Formally, individual i first makes a choice $j \in \{0, 1\}$ between attention cost functions \bar{K}_{ai}^0 and \bar{K}_{ai}^1 .⁹ We think of \bar{K}_{ai}^1 as constituting a bandwidth enhancement (BE) over \bar{K}_{ai}^0 . We let p

⁹In principle, we could formalize this choice as part of the attention strategy s_a . However, because this initial choice is the main choice that is observable to the analyst, we formally distinguish it from the other

denote the incremental cost of choosing $j = 1$ over $j = 0$, and we assume—consistent with our experiments—that it is incurred at the same time as the variable reward r . In our first experiment, $-p$ corresponds to the incentives we create for choosing our plan-making tool, while in our second and third experiments, p is the price of reminders or making the task easier.¹⁰ We think of the BEs in our experiments as increasing the likelihood of success for a given attention cost, which is equivalent to decreasing the marginal cost of attentional inputs: $\bar{K}_{ai}^1(q'_a) - \bar{K}_{ai}^1(q_a) < \bar{K}_{ai}^0(q'_a) - \bar{K}_{ai}^0(q_a)$ for all $q'_a > q_a$. The decrease in marginal costs does not preclude the possibility that the BEs may carry nuisance costs, formalized as $\bar{K}_{ai}^1(0) > \bar{K}_{ai}^0(0)$.

Define

$$V_i^j(r) := \max_{(q_a, q_o) \in [\underline{q}_a, \bar{q}_a] \times [\underline{q}_o, \bar{q}_o]} \left[r q_a q_o - \bar{K}_{ai}^j(q_a) - \bar{K}_{oi}(q_o) \right]$$

as the indirect utility functions, given attention production function \bar{K}_{ai}^j and incentives r . It is optimal for i to choose attention technology $j = 1$ if and only if $V_i^1(r) - p \geq V_i^0(r)$. Our main result characterizes testable restrictions of optimal choice using measurable statistics of aggregate behavior. The first statistic is the willingness to pay (WTP) for technology $j = 1$; that is, the highest p at which $j = 1$ is preferred to $j = 0$. If the nuisance cost of $j = 1$ is sufficiently high, this statistic can be negative, even if $j = 1$ lowers the marginal cost of attention and thus increases successful task-completion probability. Average WTP is given by

$$\bar{W}(r) := \mathbb{E}_i [V_i^1(r) - V_i^0(r)].$$

The other key statistics, which are also at the population level, are $Pr(j = 1 | p, r)$, the probability of individuals choosing technology $j = 1$ given the financial incentive r and the incremental cost p of choosing $j = 1$; $Pr(z = 1 | p, r)$, the probability of individuals successfully completing the task given the financial incentive r and the incremental cost p of choosing $j = 1$; and $Pr(z = 1 | j, r)$, the probability of successfully completing the task given financial incentive r and when exogenously assigned attention technology j .

Lemma A.1 in Appendix A.1 shows that the statistics defined above are (almost everywhere) differentiable under mild assumptions. Differentiability of the cost functions \bar{K}^j is not necessary for these statistics to be differentiable.

attention-affecting choices allowed by our general model.

¹⁰Individual differences in K_i^0 and K_i^1 could result from individual differences in baseline attentiveness such as differences in how well-suited the BE is to an individual's needs; differences in the nuisance costs of reminders, and the personal and social costs of failing to execute a plan that one creates; or (in reduced-form) differences in the indirect costs of having one's attention to other activities reduced.

Theorem 1. *Assume that individuals choose attention strategies optimally. Define*

$$D(z = 1|r) := Pr(z = 1|j = 1, r) - Pr(z = 1|j = 0, r).$$

For any $r > 0$ and $\Delta > 0$, average WTP for the BE satisfies

$$\bar{W}(r + \Delta) - \bar{W}(r) = \int_{x=r}^{r+\Delta} D(z = 1|x)dx. \quad (3)$$

Moreover, $\bar{W}(r)$ is differentiable almost everywhere, differentiable at any point where $D(z = 1|r)$ is continuous, differentiable everywhere if the cost functions \bar{K}_{ai}^j , \bar{K}_{oi} are strictly convex, and satisfies

$$\frac{d}{dx}\bar{W}(x)|_{x=r} = D(z = 1|r) \quad (4)$$

at any point of differentiability. At any pair (p, r) where $Pr(z = 1|p, r)$ and $Pr(j = 1|p, r)$ are continuously differentiable, the likelihood of choosing technology $j = 1$ and the likelihood of completing the task satisfy

$$\frac{d}{dr}Pr(j = 1|p, r) = -\frac{d}{dp}Pr(z|p, r) \quad (5)$$

$$= -\frac{d}{dp}Pr(j = 1|p, r)D(z = 1|r). \quad (6)$$

Although we make minimal assumptions about the economic environment and individuals' utility functions, we obtain the sharp characterization of Theorem 1 by utilizing Milgrom and Segal's (2002) envelope theorems for general choice sets. Equation (4) of Theorem 1 states that, if technology $j = 1$ increases individuals' likelihood of choosing $z = 1$ by, e.g., 10 percentage points under incentive r , then a small increase dr in r should increase individuals' average willingness to pay for $j = 1$ by approximately $dr \times 0.1$. Equation (3) is an integral version of equation (4) that does not require differentiability. Appendix A.4 provides an instructive graphical illustration of this result, using standard concepts from competitive supply.

However, the condition in equation (4) requires rich data that is difficult to collect in some field settings, and that we do not have in our first experiment. Equation (5) builds on equation (4) by characterizing how the probability of choosing $j = 1$ and the probability of $z = 1$ are related. The condition in equation (5) formalizes the basic intuition that if attention is allocated optimally, then increasing the incentives for $z = 1$ should increase the

desire to adopt a technology that increases the likelihood of $z = 1$. But while the qualitative comparative static could be consistent with individuals under- or overvaluing the benefits of BEs, the quantitative condition clarifies exactly how much individuals should seek BEs that increase task completion. The condition in equation (6) is a restatement of the condition in (5) that reveals the connection to equation (4) by utilizing the function $D(z = 1|r)$.

The condition in (5) is a variation on the Slutsky symmetry condition that cross-price elasticities of compensated demand functions must be equal to each other and is analogous to the tests of sparse-max decision making derived in Gabaix (2014). Intuitively, $-\frac{d}{dp}Pr(z = 1|p, r)$ is an indication of how adoption of technology $j = 1$ affects the probability of choosing $z = 1$. In our online education experiment, this derivative is the average effect of our plan-making incentives on the likelihood of completing course modules. The higher this number is, the higher the effect of our plan-making tool on the likelihood of completing the course modules will be. Then, the higher the effect of the plan-making tool, the higher the effect of a small change in incentive r on the tool's value, as formalized in the first part of Theorem 1. This translates to a higher derivative $\frac{d}{dr}Pr(j = 1|p, r)$.¹¹

Condition (4) in Theorem 1 is a limit result in the sense that it applies to marginal changes in the task-completion incentive. Corollary 1 below clarifies how this condition can be used to evaluate “small” but not “vanishing” changes.

Corollary 1. *Assume that agents choose attention technologies optimally. Then*

$$\Delta \min_{x \in [r, r+\Delta]} D(z = 1|x) \leq \bar{W}(r + \Delta) - \bar{W}(r) \leq \Delta \max_{x \in [r, r+\Delta]} D(z = 1|x). \quad (7)$$

If $D(z = 1|x)$ is smooth on $[r, r + \Delta]$ and terms of order $\frac{d}{dr^2}D(z = 1|r)(\Delta^3)$ are negligible, then

$$\bar{W}(r + \Delta) - \bar{W}(r) \approx \frac{\Delta}{2} \cdot (D(z = 1|r) + D(z = 1|r + \Delta)). \quad (8)$$

Corollary 1 shows that if the change Δ in incentives is sufficiently small, then, roughly speaking, we can take a first-difference approximation to condition (4) by replacing $\frac{d}{dx}\bar{W}(x)|_{x=r}$

¹¹Note that any data set that can be used to test condition (4) can be used to test the Slutsky symmetry condition as well. To see this, first note that eliciting individuals' WTP for $j = 1$ at incentive r is equivalent to eliciting the demand curve for $j = 1$ at incentive r , which means that this data set identifies $Pr(j = 1|p, r)$ for all p and for each task-completion incentive r utilized in the experiment. Thus, $\frac{d}{dr}Pr(j = 1|p, r)$ and $\frac{d}{dp}Pr(j = 1|p, r)$ are identified in this data set. Second, the right-hand-side term of (4), $D(z = 1|r)$, is identified by assumption. Thus, all of the statistics necessary to test (6), and therefore also (5), are available. Note, however, that the Slutsky symmetry conditions use strictly less data than condition (4): these conditions consider the function $Pr(j = 1|p, r)$ only in a neighborhood around a single incentive level r , while condition (4) considers $Pr(j = 1|p, r)$ across all possible values r for which $Pr(j = 1|p, r) \in (0, 1)$.

with $(\bar{W}(r + \Delta) - \bar{W}(r)) / \Delta$. For example, if, as in our second and third experiments, the analyst finds that the change Δ negligibly affects the likelihood of successful task completion, then condition (7) implies that $\bar{W}(r + \Delta) - \bar{W}(r) \approx \Delta D(z = 1|r)$ for optimizing agents.

The heuristic approximation above can be further refined in several ways. First, under the assumption that $D(z = 1|x)$ is monotonic on $[r, r + \Delta]$, condition (7) shows that $\Delta \min\{D(z = 1|r), D(z = 1|r + \Delta)\}$ is a robust lower bound for $\bar{W}(r + \Delta) - \bar{W}(r)$, which is helpful for analysis attempting to rigorously show that $\bar{W}(r + \Delta) - \bar{W}(r)$ is too low relative to the optimizing benchmark. Second, under the assumption that $D(z = 1|x)$ is locally linear on $[r, r + \Delta]$ —which is a justifiable assumption whenever the impacts of Δ are “small”—condition (8) provides a better approximation to $\bar{W}(r + \Delta) - \bar{W}(r)$.¹²

1.3.1 Remarks and Qualifications

Differences in fixed costs, $\bar{K}_i^1(0) - \bar{K}_i^0(0)$, may result from the potential nuisance costs of attention-improving technologies, which is consistent with negative WTP for reminders by some individuals in our second experiment. Thus, the value of a BE cannot be equated with its impact on the change in expected earnings, $rPr(z = 1)$. Simply documenting that, for example, individuals’ valuations for a reminder that increases their chance of earning \$10 by 10% is smaller than \$1 is not a rejection of correct valuation of the reminder, because nuisance costs could decrease the value of the reminder. Our more robust test focuses instead on how individuals’ valuations of the BE change as the pecuniary incentives for being attentive change.

Second, note that condition (4) is a test of whether individuals correctly value the BE *on the margin*, i.e., whether individuals correctly perceive the effects of the BE at the current incentive level r . The statistic $\bar{W}(r)$, however, captures individuals’ perceptions of the difference in total costs, which includes $K_i^1(0) - K_i^0(0)$. The impact on $\bar{W}(r)$ of additional treatments like opportunities for learning and feedback—as in our third experiment—can provide some insight about biases in the perception of total costs, and how those might differ from biases about the effects of the BE on the margin.

Third, note that we have assumed that utility is quasilinear in the financial incentives. This is a plausible assumption for the small stakes featured in our experiment, as non-negligible deviations from this assumption would imply implausible levels of risk aversion for higher incentives (Rabin, 2000). Appendix E shows that incorporating standard estimates of risk aversion in our model negligibly impacts the quantitative implications of Theorem 1,

¹²Note that Δ need not necessarily be small. What is important is that $D(z = 1|r + \Delta) - D(z = 1|r)$ is small. For example, if this statistic is zero, then the condition in Corollary 1 is exact even for large Δ .

and thus does not confound our empirical conclusions.

1.4 Empirical Tests of Theorem 1 and Their Interpretation

A failure to verify the optimality conditions in Theorem 1 could result because of the following reasons, which we discuss in turn below.

- E1. Individuals mischaracterize (either due to incorrect beliefs or other forms of bounded rationality) their attention production process. That is, they mischaracterize $Q(s_a, s_o, \omega)$ or $K_i(s_a, s_o, \omega)$ on a positive measure of states ω .
- E2. Individuals have biased priors μ .
- E3. In a given experiment, the states ω are not realized independently across individuals. For example, there is uncertainty about a fixed parameter of the economic environment.¹³

Deviations due to E1 or E2 imply systematic misperceptions. Systematic biases have been documented in a variety of domains of decision making, including other types of costly effort decisions (e.g., DellaVigna and Pope, 2017; Hoffman and Burks, 2020). These explanations are also equivalent in the simplified representation (2), as both lead to biased beliefs about the aggregate cost functions \bar{K}^j . In experiments such as our first two, it is difficult to differentiate between E1 and E2 because the state space is not specified or observed by the analyst. In our third experiment, which follows standard protocols that test rational inattention models, E2 may be less likely as the experimenter specifies the state space, the prior distribution over states, and the mapping from states and actions to financial rewards. The only way in which E2 can apply to such experimental protocols is if the state space is, contrary to the experimenter's efforts, richer than the one specified by the experimenter, for reasons we discuss below in reference to E3.

Deviations due to E3 do not necessarily involve systematic biases. To formalize E3, let ξ denote an environmental parameter such that the effects of the BE are given by $D(z = 1|r, \xi)$. Let each individual i receive a signal ζ_i about $D(z = 1|r, \xi)$, which is affiliated with ξ according to some joint distribution $H(\zeta, \xi)$, with marginal distributions H_ζ and H_ξ . Suppose that individuals have an unbiased prior about the effects of the BE, centered around $\bar{D}(z = 1|r) := \int D(z = 1|r, \xi) dH_\xi(\xi)$, the average effect of the BE across the different possible environments. Then, if individuals are Bayesian, the martingale property of beliefs

¹³We thank an anonymous referee for pointing out this mechanism and motivating our discussion of it.

implies that individuals' posterior beliefs after receiving signal ζ_i , the perceived effect of the BE, $\tilde{D}_i(z = 1|r, \zeta_i)$, must also be unbiased on average: $\mathbb{E}_{i, \xi} [\tilde{D}_i(z = 1|r, \zeta_i)] = \bar{D}(z = 1|r)$.

For example, ξ could capture environmental features that determine the efficacy of a reminder. As shown in our second experiment, reminders have larger effects when the task is further off into the future. Individuals may have a correct prior about how effective reminders are on average but may not have learned how effective reminders are for a specific task-completion delay. In our third experiment, individuals may have a correct understanding of how generally difficult it is to identify whether there are more correct or incorrect arithmetic equations, based on similar tasks they have done in the past, but there may be correlated uncertainty about exactly how much more difficult it is to complete the task with one hundred rather than ten equations. Formalizing this type of uncertainty requires specifying a richer state space.

An analyst applying Theorem 1 to an experiment that features only one particular realization of ξ would fail to verify the condition in equation (4) of the Theorem, since individuals' decisions would satisfy $\frac{d}{dr}\bar{W}(r) = \mathbb{E}_i [\tilde{D}_i(z = 1|r, \zeta_i)|\xi]$ rather than $\frac{d}{dr}\bar{W}(r|\xi) = D(z = 1|r, \xi)$. Concretely, suppose that individuals have a normal prior with variance σ_0^2 , and that they receive a signal about $D(z = 1|r, \xi)$ that is normally distributed around $D(z = 1|r, \xi)$ with variance σ_1^2 . Then individuals' perceptions, and thus by Theorem A.1 their WTPs, satisfy

$$\begin{aligned} \frac{d}{dr}\bar{W}(r|\xi) &= \mathbb{E}_i [\tilde{D}_i(z = 1|r, \zeta_i)|\xi] \\ &= (1 - \theta)\bar{D}(z = 1|r) + \theta D(z = 1|r, \xi), \end{aligned} \quad (9)$$

where $1 - \theta = \sigma_1^2 / (\sigma_0^2 + \sigma_1^2)$ is the degree of Bayesian shrinkage toward the prior mean $\bar{D}(z = 1|r)$. This implies that individuals undervalue the BE when it is more effective than average, in the sense that $D(z = 1|r, \xi) > \bar{D}(z = 1|r)$. But individuals also overvalue the BE when it is less effective than average, in the sense that $D(z = 1|r, \xi) < \bar{D}(z = 1|r)$. Thus, showing that individuals overestimate or underestimate the effects of a BE in one particular decision environment does not imply that the miscalibration is systematic.

Importantly, if the assumption that $\bar{D}(z = 1|r)$ is an unbiased prior mean is relaxed, then equation (9) is mathematically equivalent to a “meta-inattention” model in the spirit of the attribute-misperception model of Gabaix (2014). In this interpretation, \bar{D} is some default perception that people “anchor” on. Thus, equation (9) is a convenient parametrization that can also capture systematic biases as in E1 and E2.

Estimating equation (9) can provide at least suggestive evidence for differentiating be-

tween E3 and systematically biased perceptions. In data sets where there is exogenous variation in conditions ξ that generate variation in the efficacy of the BE, equation (9) can be estimated simply through the linear regression model

$$\frac{d}{dr}\bar{W}(r|\xi) = \beta_0 + \beta_1 D(z = 1|r, \xi). \quad (10)$$

The coefficient β_1 identifies θ , which implies that $\bar{D}(z = 1|r) = \beta_0/(1 - \beta_1)$. That is, $\beta_0/(1 - \beta_1)$ is a sufficient statistic for the behavioral implications of the prior μ , including any possible biases in the prior. To illustrate how the coefficients can help differentiate between E3 and the systematic biases in E1 and E2, suppose that it is known that the effects of reminders are generally non-negative.¹⁴ Then, detecting strongly positive effects in some conditions suggests that $\bar{D}(z = 1|r) > 0$, and thus finding $\beta_0 = 0$ suggests that people systematically underestimate the effects of reminders because they anchor on the erroneous default perception of null effects.

Implications for Models of Costly Attention To summarize, failure to verify the conditions of Theorem 1 generates several possible implications. One is that individuals are systematically miscalibrated about their attention cost functions or have biased priors (E1 and E2). A different possibility, as captured by E3, is that even in highly controlled empirical studies where the experimenter attempts to specify the states and probabilities—such as our third experiment—the individuals’ subjective state space is richer than what has typically been assumed. Given that state-dependent stochastic choice (SDSC) data is a key empirical object for testing and estimating rational inattention models (Caplin and Dean, 2015; Caplin, 2016), this possibility raises intriguing challenges for this agenda.

Direct Versus Indirect Reasons for Misperceptions of BE Effects Individuals might misperceive the value of BEs for direct reasons, such as misperceiving $\bar{K}_a^1 - \bar{K}_a^0$, or due to indirect reasons, such as misperceiving \bar{K}_o in economic environments that involve auxiliary actions. The latter possibility is not implicated in our third experiment but is in principle possible in our first two experiments. If individuals underestimate \bar{K}_o , then they will overestimate their optimal choice of q_o , and thus *overestimate* the returns to higher attention and thus to the BE. The converse holds if individuals overestimate \bar{K}_o . However, because most plausible biases about the costs of auxiliary actions in our first two experiments—such as underestimating how busy one is in the future (e.g., the *planning fallacy* articulated in

¹⁴I.e., an unbiased prior puts little weight on negative effects.

Kahneman and Tversky, 1982)—lead people to underestimate \bar{K}_o , under-valuation of BEs in these experiments cannot be plausibly explained without direct misperceptions of $\bar{K}_a^1 - \bar{K}_a^0$.

2 Online Education Experiment

Our first experiment was designed around the Slutsky symmetry test in equation (5) of Theorem 1. It was run in the fall of 2018. We partnered with Code Avengers, an online platform for learning to code, to offer participants a free, eight-week course in three different programming languages (HTML/CSS, Javascript, and Web Dev).¹⁵ Screenshots of all experimental instructions are in Screenshots Appendix F.1.

2.1 Design and Implementation

2.1.1 Participant Pool

We recruited students and recent alumni from six Philadelphia-area colleges using an email campaign. Enrollees were eligible to be included in our study if they reported in the onboarding survey that they regularly used either Google Calendar or Apple’s iCal as an electronic calendar. Perhaps due to the relative youth of the participant pool, usage rates were high, at around 60 to 70 percent. Recruitment resulted in a pool of 1373 study-eligible participants.¹⁶

2.1.2 Implementation

Just before the 8-week course began, participants received an introductory email with information on their treatment assignment. This email also contained a recommendation that participants aim to complete three, 15-minute sessions of the coding course per week, a prompt to encourage participants to make a plan for when they would do the coding lessons, and a link to make plans for working on the coding lessons, which would be created in their electronic calendars. Participants who were eligible for financial rewards were informed that they would be paid their cumulative earnings in the form of an Amazon gift card at the end of the 8-week period.

Over 90 percent of participants opened the initial emails informing them of the incentives they faced (i.e., their treatment), giving us confidence that most were aware of the incentives

¹⁵These languages are commonly used tools for building modern web sites. See <http://www.codeavengers.com> for more details.

¹⁶Appendix Table A.1 presents characteristics of the participant pool. Females, first-years, and seniors were most likely to participate.

for which they were eligible. As expected from random assignment of treatment, email opening rates were very similar across treatments, ranging from 88 to 91 percent.

After the course had begun, all participants received a reminder email at the start of each week. The reminder email contained the same recommendation, planning prompt, and link to create plans as the initial email.

2.1.3 Experimental Design

The experiment consisted of a control group and five treatment arms, with varying levels of incentives for plan making and/or coding task completion. Participants assigned to the control group received the initial and reminder emails encouraging them to plan and complete the coding lessons and offering them the plan-making tool, but they were not eligible for financial rewards.

Those randomly assigned to the two *Pay-to-Plan* treatments received either \$1 or \$2 for making a plan for when to do their coding lessons that week (i.e., clicking the plan-making link within the weekly email). In the two *Pay-to-Code* treatments, participants received either \$2 or \$5 for completing three 15-minute sessions of the coding course during the week. Finally, participants in the *Combination* treatment arm were paid \$1 for making a plan plus \$2 if they completed three 15-minute sessions of the coding course during the week. Participants could earn these amounts each week, regardless of what they had done in previous weeks. In addition, making a plan did not restrict when a participant could do the coding lessons (i.e., participants in the *Pay-to-Code* and *Combination* treatments could complete the 15-minute sessions at any time during the week and still earn their coding-task incentives, regardless of whether or not they made a plan or when they had scheduled the three 15-minute sessions).

To measure plan making, we tracked whether a participant clicked on the provided plan-making link to create calendar events for when they planned to complete the 15-minute coding sessions.¹⁷ Consistent with our theoretical framework, this observable plan-making is not the only available bandwidth enhancement (BE), or even the only available plan-making opportunity. For example, some participants might have other means of making plans or might directly edit their calendars without using our link. However, nearly 40 percent of the control group clicked to make a plan in the first week, despite receiving no financial rewards for doing so, and participants with higher incentives for completing the coding task were more likely to use the plan-making tool, implying that our plan-making tool was not

¹⁷When participants clicked on the plan-making link, they were given three default times, which they could change. This default ensured that as long as a participant clicked on the link, a calendar event would be created.

a perfect substitute for the plan making individuals would do otherwise.¹⁸ This may be because the act of making a plan by using our link generates an internal cue, as theorized in the implementation intentions literature (Gollwitzer and Sheeran, 2006).

To measure completion of the coding coursework, we received real-time, backend data from Code Avengers on the number of minutes participants spent actively working on their coding coursework each day. The session timer stopped running after approximately 30 seconds of inactivity within the course. Once they had completed 15 minutes of active work, participants were notified with a pop-up that congratulated them but did not prevent or discourage them from continuing.

2.2 Results

2.2.1 Empirical Framework

Our primary analysis focuses on measuring the effect of plan-making and coding-task incentives on plan making and coding task completion. We estimate treatment effects using regressions of the form:

$$y_{ict} = \beta T_{ict} + \alpha_c + \alpha_t + \gamma X_i + \varepsilon_{ict}, \quad (11)$$

where y_{ict} measures either plan making or completing at least $\tau \in \{0, 10, 20, 30, 40, 45, 50, 60\}$ minutes in week t for participant i at campus c . We include fixed effects α_c for campus interacted with student status (i.e., current student or alumni), which was the level at which we randomized. We also control for course week α_t and a vector of participant characteristics X_i , but random assignment implies that these additional controls should not affect our estimated treatment effects. Our preferred measure of treatment T_{ict} is value in dollars of the participant’s incentive, which assumes a linear relationship between the incentive and behavior. We also consider a specification with indicators for different incentive sizes. We estimate regressions separately for the Pay-to-Plan sample, which includes the control group and the two *Pay-to-Plan* treatments, and the Pay-to-Code sample, which includes the control group and the two *Pay-to-Code* treatments.

¹⁸Our theoretical framework only requires that the plan-making tool we offer is not a perfect substitute to other forms of planning individuals already undertake. Heterogeneity in attention cost functions accommodates the possibility that some participants who use our plan-making tool simply substituted from creating their own calendar reminders while others who use our plan-making tool would not have created a plan themselves.

2.2.2 Plan-making Incentives

In Table 1, we estimate the impacts of plan-making incentives on plan making and on coding task completion. In the context of our model, these specifications measure $\frac{d}{dp}Pr(z = 1|p, r)$ and $\frac{d}{dp}Pr(j = 1|p, r)$. The analysis sample includes 705 participants and eight pooled weekly observations per participant. In Panel A, we estimate the effect of plan-making incentives on the propensity to plan in week 1, weeks 1 to 4, and weeks 1 to 8. Multiple-week outcomes average the indicator for whether a participant made a plan (or completed the coding task) in each week. In Panel B, we estimate the effect of plan-making incentives on the propensity to complete at least 20 minutes or at least 45 minutes of coding during week 1, weeks 1 to 4, and weeks 1 to 8, respectively. Although our financial incentives were specifically for completing at least 45 minutes of the coding task (i.e., the three 15-minute sessions), we also include the 20-minute benchmark in the main tables and text to show robustness. Appendix Tables A.2 and A.3 consider other time thresholds: 0, 10, 30, 40, 50, and 60 minutes per week. Our interpretation of the results is consistent with the evidence from these alternative thresholds.

The results indicate strong impacts of plan-making incentives on plan making, and modest impacts of plan-making incentives on coding task completion. For each \$1 of plan-making incentive, participants increase their plan making by 11.6 percentage points ($se = 1.3$) on average over the eight weeks of the study, an increase of 140% relative to the control group mean of 8.2 percentage points. Plan-making effects are 18.0 percentage points ($se = 2.0$) in week 1, and 14.2 percentage points ($se = 1.4$) on average over weeks 1 to 4, which suggests an attenuated response over the course of the study. However, the control mean falls even more quickly, from 38.1% in week 1, to 15.0% in the first four weeks, to 8.2% over the full study, such that the relative impact of plan-making incentives increases over time. Panel A of Appendix Table A.4 shows the effects of the \$1 and \$2 plan-making incentives separately.

The treatment effect of plan-making incentives on coding task completion is more modest but still meaningful. Focusing on course completion of at least 45 minutes a week, we find that \$1 of plan-making incentive increases coding task completion by 3.8 percentage points ($se = 1.8$) in week 1, an increase of 22% relative to the control group mean of 17.4 percentage points. However, the effect declines to a marginally significant 1.7 percentage points ($se = 1.2$) over weeks 1 to 4, and to a statistically insignificant 0.6 percentage points ($se = 0.9$) over weeks 1 to 8. In Panel C, we combine the plan making and coding task completion estimates in an instrumental variables estimation of the effect of plan making on

Table 1: The Effect of Plan-Making Incentives on Plan Making and Task Completion

A. The Effect on Plan Making (First Stage)						
	(1)	(2)	(3)			
	Week 1	Weeks 1-4	Weeks 1-8			
Plan Incentive	0.180*** (0.020)	0.142*** (0.014)	0.116*** (0.013)			
Obs.	705	705	705			
R ²	0.137	0.163	0.131			
Control Mean	0.381	0.150	0.082			
Controls	Yes	Yes	Yes			
Campus FE	Yes	Yes	Yes			

B. The Effect on Coding Task Completion (Reduced Form)						
	(1)	(2)	(3)	(4)	(5)	(6)
	>20 (1)	>20 (1-4)	>20 (1-8)	>45 (1)	>45 (1-4)	>45 (1-8)
Plan Incentive	0.040** (0.020)	0.028** (0.013)	0.013 (0.011)	0.038** (0.018)	0.017 (0.012)	0.006 (0.009)
Obs.	705	705	705	705	705	705
R ²	0.057	0.049	0.051	0.036	0.035	0.041
Control Mean	0.280	0.212	0.158	0.174	0.156	0.116
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Campus FE	Yes	Yes	Yes	Yes	Yes	Yes

C. The Effect of Plan Making on Coding Task Completion (IV)						
	(1)	(2)	(3)	(4)	(5)	(6)
	>20 (1)	>20 (1-4)	>20 (1-8)	>45 (1)	>45 (1-4)	>45 (1-8)
Plan Making	0.221** (0.105)	0.194** (0.087)	0.114 (0.086)	0.213** (0.096)	0.118 (0.076)	0.053 (0.074)
Obs.	705	705	705	705	705	705
R ²	0.147	0.174	0.133	0.092	0.120	0.085
Control Mean	0.280	0.212	0.158	0.174	0.156	0.116
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Campus FE	Yes	Yes	Yes	Yes	Yes	Yes

This table estimates the effect of plan-making incentives (“Plan Incentive”) on plan making and coding task completion. Panel A shows the effect of plan-making incentives (in dollars) on whether participants made a plan. Column (1) shows the effect of plan-making incentives in week 1 of the experiment. Column (2) shows the average effect for the weeks 1–4. Column (3) shows the average effect for all weeks. Panel B shows the effect of plan-making incentives on coding task completion. Columns (1-3) show the effect on an indicator variable for whether or not the participant worked on the task for more than 20 minutes: Column (1) estimates the effect over week 1, Column (2) over weeks 1–4, and Column (3) over all weeks. Columns (4-6) show analogous estimates for an indicator for whether or not the participant worked on the task for more than 45 minutes. Panel C shows the 2SLS estimates instrumenting for whether participants made a plan using the plan-making incentive as an instrument. The dependent variables are the same as in Panel B. Standard errors, clustered at the participant level, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

coding task completion. Making a plan increases the probability of coding task completion by 21 to 22 percentage points in week 1, an 81% to 124% increase relative to control group means. This large effect is precisely estimated for week 1 and weeks 1 to 4 but diminishes over the full experimental period. Overall, the results point to the value of plan making for people who have some intrinsic motivation to complete the coding sessions. Panel B of Appendix Table A.4 shows the effects of the \$1 and \$2 plan-making incentives separately.

The decrease in treatment effects over time is not surprising, as many participants appear to attrit out of the coding course. Appendix Figure A.2 plots control group means for plan-making and coding task completion over the weeks of the experiment. Engagement in the first two weeks of the study is relatively high in the absence of monetary incentives—control group participation hovers between 20 and 30 percent. However, many participants disengage from both the plan-making tool, which falls close to zero by week 3, and from continuing the coding course, which falls to 10% participation by week 5, suggesting that motivation for the coding course diminished over time.¹⁹

2.2.3 Coding-task Incentives

Table 2 estimates the impacts of coding-task incentives on plan making and coding task completion. In the context of our model, these specifications measure $\frac{d}{dr}Pr(z = 1|p, r)$ and $\frac{d}{dr}Pr(j = 1|p, r)$. The analysis sample includes 714 participants and eight pooled weekly observations per participant. Following the structure of Table 1, in Panel A we estimate the effect of coding-task incentives on the propensity to plan in week 1, weeks 1 to 4, and weeks 1 to 8. In Panel B, we estimate the effect of coding-task incentives on the propensity to complete at least 20 minutes or at least 45 minutes of coding during week 1, weeks 1 to 4, and weeks 1 to 8, respectively.

Coding-task incentives have substantial effects on coding task completion, as shown in

¹⁹Note that this by itself does not imply a deviation from optimal Bayesian decision making. As a simple illustration, suppose that for each participant, the beliefs about the benefits of course completion, denoted $B(r)$, take the form of a Bernoulli random variable that takes on the values $\bar{B} > 0$ with probability 0.2 and $\underline{B} < 0$ with probability 0.8, such that $0.2\bar{B} + 0.8\underline{B} > 0$. Then participants would initially sign up given the positive expectation of $B(r)$. But if the realizations of B are independently distributed across participants, 80 percent of them would attrit after discovering that $B = \underline{B}$. And if the realizations of B are positively correlated across participants, then in “bad” states the number of participants attriting could be much larger than what participants initially expected. For example, if the realizations are perfectly correlated, and the prior is that $Pr(B = \bar{B}) = 0.8$, then in the state $B = \underline{B}$ the participants who attrit will have ex-ante expected to complete the course with 80 percent chance. This does not pose a threat to our theoretical results about tests of optimal valuation of BEs because the experiment involves weekly measures of engagement with the planning tool. Thus, even if participants initially overestimated their enthusiasm about the course, they had the opportunity to adjust those expectations before the next planning decision.

Panel B. We estimate that each \$1 of coding-task incentive increases completion rates for 45 minutes in week 1 by 3.5 percentage points ($se = 0.8$), an increase of 20% relative to the control group mean of 17.4 percentage points. For the \$2-incentive and \$5-incentive groups, this coefficient implies an increase in the probability of coding task completion of 7 and

Table 2: The Effect of Coding-Task Incentives on Plan Making and Task Completion

A. The Effect on Plan Making						
	(1)	(2)	(3)			
	Week 1	Weeks 1-4	Weeks 1-8			
Task Incentive	0.025*** (0.009)	0.010** (0.004)	0.007** (0.003)			
Obs.	714	714	714			
R ²	0.050	0.058	0.049			
Control Mean	0.381	0.150	0.082			
Controls	Yes	Yes	Yes			
Campus FE	Yes	Yes	Yes			

B. The Effect on Coding Task Completion						
	(1)	(2)	(3)	(4)	(5)	(6)
	>20 (1)	>20 (1-4)	>20 (1-8)	>45 (1)	>45 (1-4)	>45 (1-8)
Task Incentive	0.038*** (0.009)	0.031*** (0.006)	0.025*** (0.005)	0.035*** (0.008)	0.024*** (0.006)	0.020*** (0.005)
Obs.	714	714	714	714	714	714
R ²	0.043	0.059	0.069	0.041	0.057	0.075
Control Mean	0.280	0.212	0.158	0.174	0.156	0.116
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Campus FE	Yes	Yes	Yes	Yes	Yes	Yes

This table estimates the effect of coding-task incentives (“Task Incentive”) on plan making and coding task completion. Panel A shows estimates of the effect of coding-task incentives (in dollars) on whether or not participants made a plan. Column (1) shows the effect of coding-task incentives week 1 of the experiment. Column (2) shows the average effect over weeks 1–4. Column (3) shows the effect over all weeks. Panel B shows the effect of coding-task incentives on coding task completion. Columns (1-3) show the effect on an indicator variable for whether or not the participant worked on the task for more than 20 minutes: Column (1) estimates the effect over week 1, Column (2) over weeks 1–4, and Column (3) over all weeks. Columns (4-6) show analogous estimates for an indicator for whether or not the participant worked on the task for more than 45 minutes. Standard errors, clustered at the participant level, are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

17.5 percentage points, respectively, or 40% and 101% relative to the control mean of 17.4 percentage points. Again, the treatment effects diminish over time to 2.4 percentage points ($se = 0.6$) per \$1 over weeks 1 to 4, and to 2.0 percentage points ($se = 0.5$) per \$1 over the eight weeks of the study.²⁰

A more novel result is that coding-task incentives also increase the probability of plan making, as shown in Panel A. Column 1 shows that for each \$1 of coding-task incentive, participants increase their plan making by 2.5 percentage points ($se = 0.9$) in week 1, by 1.0 percentage point ($se = 0.4$) in weeks 1 to 4, and by 0.7 percentage points ($se = 0.3$) over the eight weeks of the study. Relative to the control group means of 38, 15, and 8 percentage points, these correspond to plan making increases of 6.6%, 6.7%, and 8.5% per \$1 of plan-making incentive.

2.2.4 Symmetry Test

Participants clearly recognize the potential value of plan making in helping them achieve their coding course participation. But do they value plan making enough? To answer this question, we compare the cross-price elasticities estimated in the Pay-to-Plan and Pay-to-Code samples, implementing the test in equation (5) of Theorem 1. The coefficients for \$1 of plan-making incentives on coding task completion are 0.039, 0.017, and 0.006 in week 1, weeks 1 to 4, and weeks 1 to 8, respectively. The analogous coefficients for \$1 of coding-task incentives on plan making are 0.025, 0.010, and 0.006. The difference in coefficients provides our first test of under-planning, delivering estimates of 0.014 ($se = 0.019$), 0.007 ($se = 0.012$), and -0.0004 ($se = 0.009$), respectively.²¹ The positive sign of the differences, particularly in the early weeks of the study, hints at the possibility that participants might undervalue plan making. However, the standard errors are too wide to draw strong conclusions from this data about whether participants plan optimally.

Figure 1 plots week-by-week coefficients for plan-making and coding-task incentives to illustrate how the effect of incentives evolves over the course of the experiment. The effect of

²⁰We exclude the *Combination* treatment from our main analysis and separately evaluate whether this treatment exhibits complementarity effects (i.e., whether combining a \$1 plan-making incentive with a \$2 coding-task incentive induces plan making or coding effects that are significantly different from the \$1 *Pay-to-Plan* or \$2 *Pay-to-Code* treatments in isolation). For weeks 1 to 8, the *Combination* treatment effect on plan making is 26.7 percentage points ($se = 2.6$) compared to 23.9 percentage points ($se = 2.7$) for the \$1 *Pay-to-Plan* treatment (p-value of difference = 0.31). The *Combination* treatment effect on average course completion is 3.8 percentage points ($se = 2.4$) compared to 4.6 percentage points ($se = 2.1$) for the \$2 *Pay-to-Code* treatment (p-value of difference = 0.72). Thus, we find no statistically significant complementarity effect of the *Combination* treatment.

²¹Standard errors for coefficient differences are estimated via seemingly unrelated regression.

coding-task incentives on plan making is consistently close to zero (after week 1) and tightly estimated. In contrast, the effect of plan-making incentives on coding task completion is positive for the first half of the study and then decays toward zero, with relatively wider confidence intervals.²² This provides suggestive evidence of under-planning.

Findings from Experiment 1. *Take-up of our plan-making tool increased with incentives for completing coding lessons, but the ratio of cross-price effects for plan-making and task completion suggests that participants undervalued plan-making. The impact of task-completion incentives on planning-prompt demand was 74% as large as the ex-post optimal benchmark implies.*

3 Online Survey-completion Experiment

Complementing our first experiment, we ran a survey-completion experiment on Amazon’s Mechanical Turk platform (MTurk). The experiment is tightly tied to the test in equation (4) of Theorem 1, described in Section 1. The test states that for individuals who optimally invest in bandwidth enhancements (BEs), a \$1 increase in the incentive for task completion must increase willingness to pay for such a technology by \$1 times its efficacy (i.e., by the change in the probability of task completion due to the BE). Screenshots of all experimental instructions are in Screenshots Appendix F.2.

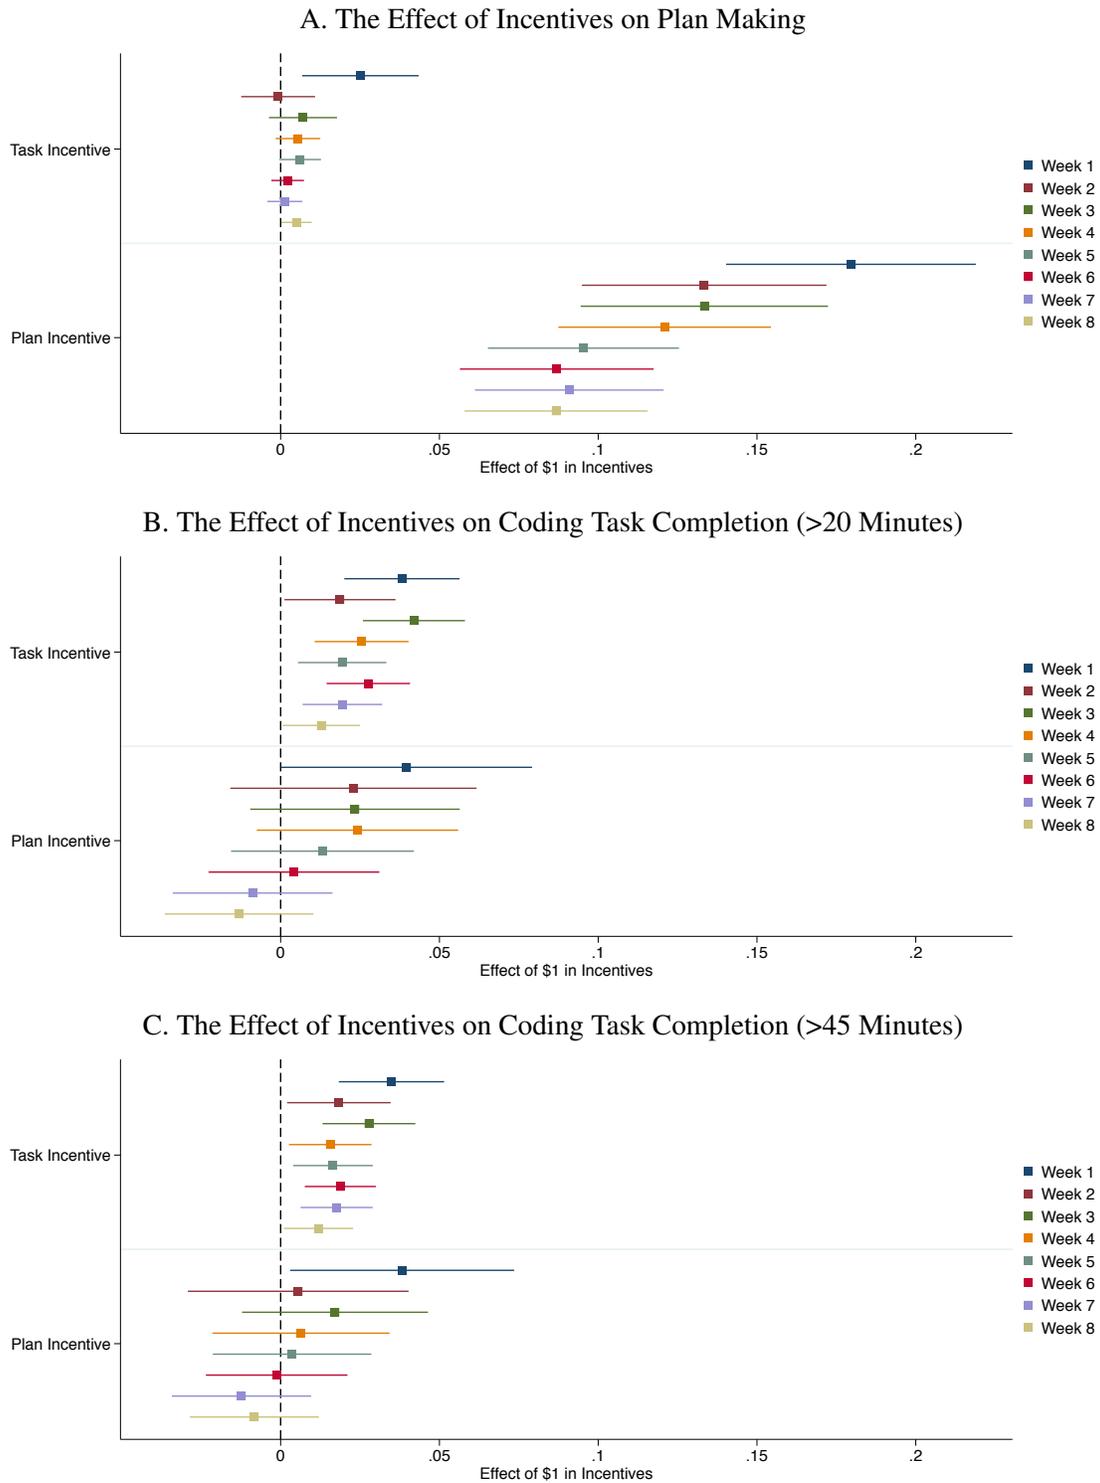
3.1 Design and Implementation

Participants were recruited to complete part 1 of the study each weekday between September 7 and September 24 of 2021. Our recruitment material informed potential participants that part 1 of the study would require 15 minutes of time immediately (for which participants were paid a guaranteed \$2.50 and had the possibility of earning a bonus), and that they would be invited to participate in part 2 of the study at a later date for additional compensation by accessing a website provided to them in part 1 of the study.

When participants clicked to begin the study, they were told that part 2 of the study—a survey that needed to be completed in one sitting of approximately 20 minutes—would only be available starting on some day in the future to be randomly determined during part 1. Participants were told that they would have a one-week window to complete it.

²²The difference in standard errors across treatments is due to higher variance in coding-task incentives (\$0, \$2, and \$5) relative to plan-making incentives (\$0, \$1, and \$2).

Figure 1: The Effect of Incentives on Plan Making and Coding Task Completion



This figure shows estimates for the effect of incentives on plan making and coding task completion for each week of the study. Panel A shows estimates of the effect of incentives on whether or not participants made a plan. Panels B and C show the effect of incentives on completing at least 20 minutes and at least 45 minutes, respectively, of coding during the week. Whiskers report 95% confidence intervals around each estimate.

The first part of the study elicited participants' willingness to pay (WTP) for a set of three reminder emails (i.e., the BE in this experiment) that would come during the one-week window in which participants would be able to complete the survey. The goal was to generate data that would allow us to directly measure how much more participants were willing to pay for reminder emails as the incentive to complete the survey increased.

To ensure that participants understood the specific details of the reminder emails, we explained that the emails would come at 12 p.m. ET on the first, middle, and final days of the one-week window in which they could complete the survey.²³ Participants were told that emails would be sent using the MTurk email system—which MTurk uses for communicating with workers on its platform—so participants did not have to provide an email address and so the reminder emails would be unlikely to go to spam. Participants were told that the link to the survey would be included in the reminder emails so that initiating the survey would be as easy as clicking a link in the email. Participants were also explicitly told that they would not receive any reminders to complete the survey unless they were selected to receive these three reminder emails. We also clarified what the part-2 survey would look like (i.e., answering 40 hypothetical questions about gambles), and provided two example questions, in order to reduce ambiguity about the future tasks.²⁴

Participants were informed that the part-2 survey would only be available starting in either 2 days, 1 week, 3 weeks, or 6 weeks, and that each delay was equally likely to be selected. Participants also learned that their incentive for completing the part-2 survey would be either \$3, \$4, \$11, or \$12. For ease of exposition, we refer to \$3 and \$4 as *low incentives* and \$11 and \$12 as *high incentives*. For each of the 16 combinations of the four possible incentive amounts and four possible delays, participants faced an incentivized multiple price list (MPL) that traded off part-1 bonus payments (up to \$4, in 25-cent increments, for the low-incentive MPLs and up to \$12, in 75-cent increments, for the high-incentive MPLs) against being sent the three reminder emails to complete the survey. Participants were informed that all possible bonus rewards, including part-1 and part-2 bonus payments, would be paid at the same point in time, after the one-week window to complete the survey ended.²⁵ Participants

²³For example, for participants who completed part 1 of the study on September 7, the 2-day-delay window was open from September 9–15. To any participants in the 2-day-delay group who were selected to receive reminder emails, we sent the emails at 12 p.m. ET on September 9, 12, and 15.

²⁴In order to participate in the study, participants needed to correctly answer questions demonstrating their understanding of the compensation structure, the tasks in part 2 of the study, and the conditions for receiving reminders (i.e., they had to answer “True” to the statement: “You will not receive any reminders to complete part 2 of the study unless you are selected to get them in this part of the study.”). Participants were also shown an MPL attention check screen that was used to remove participants who might click through the MPLs without reading the instructions.

²⁵Part-1 bonus payments were paid out at the same time as any part-2 bonus payments, three days after

were randomly selected to either first respond to the eight low-incentive MPLs (the “low-incentive block”) or to the eight high-incentive MPLs (the “high-incentive block”). Within each block, the eight MPLs were shown to participants in a random order.

Because nuisance costs can lead participants to have negative WTP for the reminders, the MPL allowed participants to report both positive and negative willingness to pay for the reminder emails.²⁶ Participants were told that whichever incentive amount was randomly selected for them (\$3, \$4, \$11, or \$12) would be the bonus they would receive for completing the part-2 survey. In addition, they were told that for the randomly selected incentive amount, there was a 10% chance that one of the rows of that MPL would be randomly selected (each with equal probability) and that whatever the participant chose in that row would be implemented (i.e., they would receive whatever part 1 bonus payment was indicated in their choice, and they would receive the reminder emails if they chose the option on the left). Because testing the optimality conditions in Theorem 1 also requires estimating the effect of the reminder emails on completing the survey, we did not guarantee that one of the MPL rows would be selected. Instead, we randomized 45% of participants to receive the reminder emails and 45% of participants not to receive the reminder emails, regardless of their MPL choices. We use this random variation to estimate the effect of reminder emails on completing the survey. We randomly assign reminder emails in this way, and estimate the effect of reminders using this sample, in order to avoid potential selection bias that might arise if there were a correlation between WTP for reminders and the rate at which individuals completed the survey.

the end of the one-week window to complete the survey, mitigating concerns that part-1 bonuses would be viewed as being paid immediately, which might have made them particularly valuable from the perspective of a quasi-hyperbolic discounter.

²⁶Consistency on an MPL requires a participant to always choose the option on the left, always choose the option on the right, or switch from choosing the option on the left to choosing the option on the right in one row of the MPL. Our MPL was implemented to allow participants to choose a single cross-over point, thus enforcing consistency in choices. Use of single-cross-over MPLs is common in the experimental literature as they make the decision faster and easier for participants. The main concern is failing to identify participants who are clicking randomly through the study (i.e., those who would likely be identified as inconsistent on the MPL if required to make a selection in each row). This concern is mitigated in our setting because of our extensive attention checks.

3.2 Results

3.2.1 Sample

A total of 2743 individuals fully completed the first part of the study.²⁷ Additionally, we make the conservative sample restriction to limit all of our analysis to individuals whose WTPs were never top-coded at the smaller incentive (i.e., \$3 or \$11) or bottom-coded at the larger incentive (i.e., \$4 or \$12).²⁸ Mechanically, these top-coded and bottom-coded individuals cannot increase WTP when the task-completion incentive rises, which could lead to an attenuation bias in our estimates of how WTP for reminders changes with task-completion incentives. Given the wide range of values offered in the MPL, only 8.37% and 8.12% of responses were top-coded on the low-incentive and high-incentive MPLs, respectively, and 0.80% and 0.54% of responses were bottom-coded on the low-incentive and high-incentive MPLs, respectively. In what follows, we report on data from the remaining 2306 participants. Our restriction is conservative because it can only increase our estimates of how WTP for reminders varies with incentives; indeed, without this restriction, the point estimates are slightly lower.

3.2.2 Impact of Reminders on Survey Completion

As described in Section 3.1, we randomized 90% of participants to either get or not get the reminder emails, regardless of their reported WTP. This randomization allows us to generate an estimate of the effect of the reminders on survey completion at each delay. In addition, since we independently randomized the incentive level for completing the survey, we can estimate the effect of reminders at low and high incentive levels.

Figure 2 presents this data. Panel A shows the rate at which participants complete the survey at each delay, and by whether participants receive reminders. Panel B summarizes the treatment effect of receiving reminders at each delay and incentive level. Without reminders, completion rates decrease with delay at both high and low incentives. With re-

²⁷This number does not include the 1854 participants who were automatically screened out of the study (and prevented from participating further) because they failed attention checks, ensuring our pool of participants understood the instructions in our experiment. It also excludes 36 individuals who were excluded for having an invalid MTurk ID or the 36 individuals who had technical issues in the display of MPL screens or recording of the data (e.g., being shown the wrong combination of incentives and delays or not receiving a link to the part-2 survey).

²⁸We define *top-coded* participants as those who chose the option on the right in each row, indicating a WTP for reminders of more than \$4 (on the low-incentive MPLs) or \$12 (on the high-incentive MPLs). *Bottom-coded* participants chose the option on the left in each row, indicating a WTP for reminders of less than -\$4 or -\$12.

mindings, however, the impact of delay on completion rates is much smaller. This translates into an increasing impact of reminders as delay increases, as shown in Panel B.

Table 3 quantifies these results. Column (1) shows that receiving the reminders increases the likelihood that participants complete the survey by 23 percentage points. The estimate on High Incentive shows that participants who receive high incentives to complete the survey are 7 percentage points more likely to complete the survey than those who receive low incentives. The coefficient on $\ln(\text{P2 Delay})$ implies that participants are less likely to complete the survey as the delay increases. Column (2) shows that reminders have a significantly smaller effect at high incentives, but a significantly larger effect at longer delays. Column (3) shows that because reminders have very small effects in the short-delay / high-incentive conditions (Panel B of Figure 2), the impact of delay on the effects of reminders is particularly large in the high-incentive condition.²⁹ Columns (4) and (5) show that neither the point estimates nor the standard errors of the column (2) and (3) models change when (i) including fixed effects for when participants begin the study or (ii) doing two-way clustering by when participants started part 1 and by when part 2 was available.³⁰ Column (6) shows that a \$1 change in the incentives is small enough to not significantly affect behavior, which allows us to utilize Corollary 1.

As formalized in Appendix A.5, the negative interaction between reminders and high incentives suggests that the effect of the high incentives on survey completion was at least in part due to individuals choosing a higher level of attention in the absence of reminders. In our model, higher incentives increase the likelihood that individuals complete the task conditional on being attentive. However, since reminders increase the likelihood of being attentive, there would then be a *positive* interaction between reminders and incentives. Instead, if higher incentives increase individuals' effort to be attentive even in the absence of reminders, then there is less need for reminders, leading to a negative interaction.

3.2.3 How WTP Changes with the Incentive to Complete the Survey

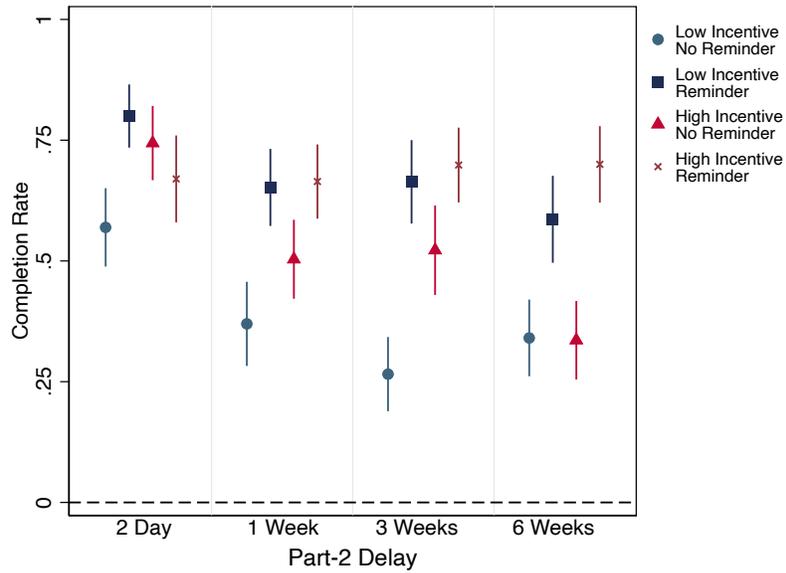
Figure 3 presents the average WTP for reminder emails for each part-2 incentive level at each of the four delays. Participants are willing to pay significantly more for reminders at the high incentives (i.e., \$11 and \$12, shown on the right of each panel) than at the low incentives

²⁹Appendix Table A.5 presents a less parametric regression analysis that separately estimates the effect of each of the possible delays, as well as its interaction with reminders, on survey completion.

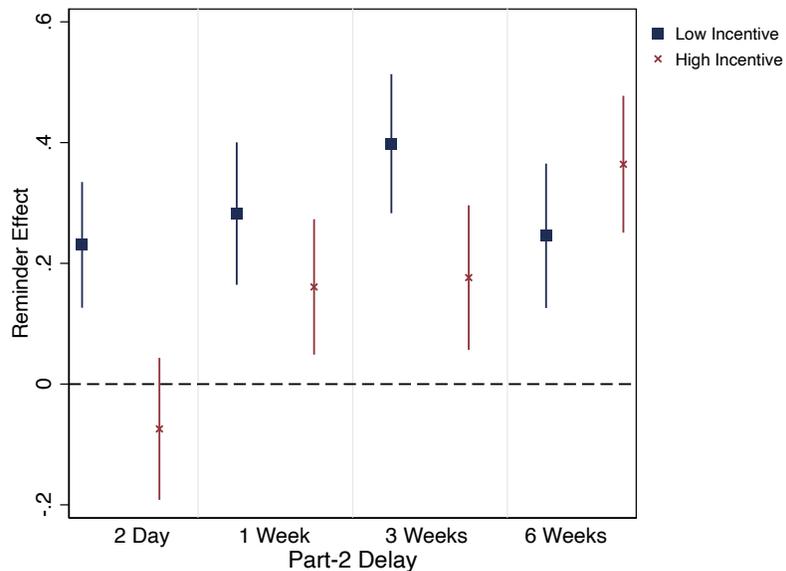
³⁰We do not include fixed effects for part-2 survey availability because certain start dates are possible only in the 2-day condition, which means that the full set of fixed effects is not separately identified from the other covariates. However, because clustering by part-2 survey start date does not change standard errors, this is unlikely to matter.

Figure 2: Completion Rates and Treatment Effects of Reminders

A. Completion Rates by Incentive Level, Delay Type, and Reminders



B. Treatment Effect of Reminders by Incentive Level and Delay



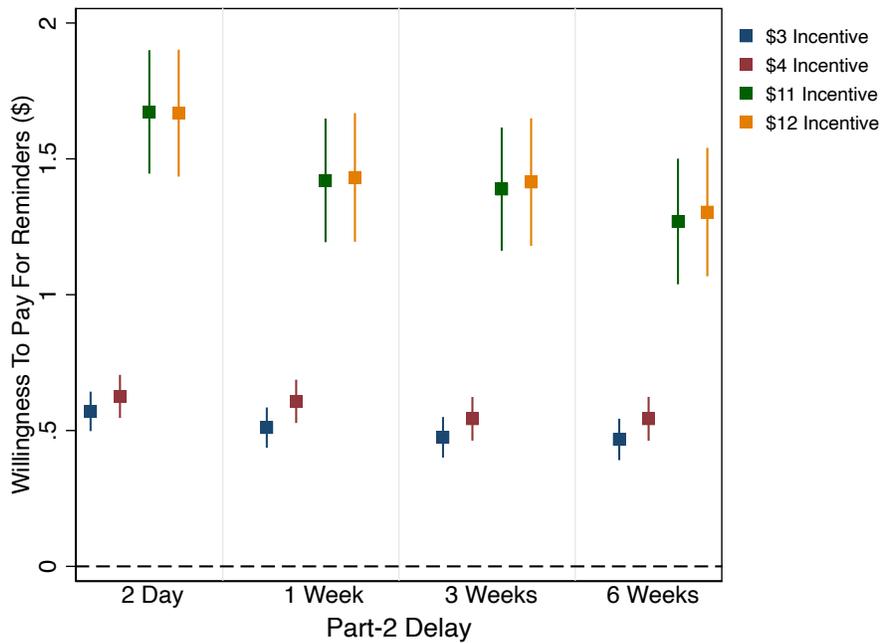
Panel A shows the part-2 survey completion rate and how it varies with the incentive for completing the survey, the amount of time after which part 2 of the survey became available (i.e., the part-2 delay), and whether the participant received reminders. Panel B shows the point estimates from a regression of the part-2 survey completion rate on whether the participant received a reminder. Both panels only include participants who were randomly assigned to receive or not receive reminders. The lines represent 95% confidence intervals.

Table 3: The Effect of Incentive, Delay, and Reminders on Part-2 Survey Completion

	Completed Part-2 Survey					
	(1)	(2)	(3)	(4)	(5)	(6)
Received Reminder	0.23*** (0.021)	0.12** (0.051)	0.23*** (0.064)	0.12** (0.055)	0.22*** (0.058)	0.23*** (0.021)
High Incentive	0.07*** (0.021)	0.13*** (0.029)	0.21*** (0.066)	0.13*** (0.031)	0.20*** (0.063)	0.07*** (0.021)
Ln(P2 Delay)	-0.07*** (0.009)	-0.10*** (0.013)	-0.08*** (0.017)	-0.10*** (0.020)	-0.08*** (0.022)	-0.07*** (0.009)
Received Reminder \times High Incentive		-0.13*** (0.042)	-0.38*** (0.094)	-0.11*** (0.039)	-0.32*** (0.070)	
Received Reminder \times Ln(P2 Delay)		0.07*** (0.018)	0.02 (0.025)	0.07*** (0.020)	0.02 (0.025)	
High Incentive \times Ln(P2 Delay)			-0.03 (0.025)		-0.03 (0.025)	
Received Reminder \times High Incentive \times Ln(P2 Delay)			0.11*** (0.036)		0.09*** (0.031)	
Extra \$1						0.00 (0.021)
Constant	0.57*** (0.028)	0.62*** (0.036)	0.59*** (0.045)	0.62*** (0.054)	0.59*** (0.057)	0.57*** (0.029)
Observations	2,076	2,076	2,076	2,076	2,076	2,076
Number of Participants	2,076	2,076	2,076	2,076	2,076	2,076
S.E. Clustered by P1 & P2 Date				X	X	
P1 Date FE				X	X	

This table estimates how survey completion varies with reminders, the natural log of delay (in days), and whether participants are offered high incentives (i.e., \$11 or \$12) or low incentives (i.e., \$3 or \$4) to complete the survey. This table only includes participants who were randomly assigned to receive or not receive reminders. Columns (4) and (5) reproduce Columns (2) and (3) with fixed effects for the date that part 1 of the study was taken and with standard errors clustered for the date the participant completed part 1 and the date part 2 was made available to them. Column (6) reproduces Column (1) but estimates the impact of an extra \$1 of incentive for completing the survey (i.e., the incentive being \$4 or \$12). Standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 3: Willingness to Pay for Reminders



This figure shows mean willingness to pay for the reminders across different experimental conditions. The lines represent 95% confidence intervals with standard errors clustered at the participant level.

(i.e., \$3 and \$4, shown on the left of each panel). Additionally, WTP for reminders is higher at the \$4 incentive than at the \$3 incentive.

We formalize the results from Figure 3 in Table 4, combining data from all four incentive levels and all four delays to estimate how average WTP changes with the incentive to complete the survey.³¹ The coefficient on Extra \$1 is the impact on willingness to pay for reminders of increasing the survey-completion incentive from \$3 to \$4. The coefficient on High Incentive compares average WTP at \$3 to average WTP at \$11. The coefficient on Extra \$1 \times High Incentive compares the impact on WTP of increasing the incentive by \$1 when the incentives are high (i.e., going from \$11 to \$12) to the impact when incentives are low (i.e., going from \$3 to \$4).

Column (1) shows that as the incentive to complete the survey increases from \$3 to \$4, participants are on average willing to pay around 7 cents more for the reminders. However, the interaction on Extra \$1 \times High Incentive is negative and similarly sized, implying that when incentives are high, the extra dollar of incentive does not lead to an increase of WTP.

³¹Appendix Table A.6 replicates Table 4 for the 90% of participants who either receive or do not receive the reminder emails based on random assignment. As one would expect from the fact that this 90% is randomly selected, estimates are nearly identical to those in Table 4.

That the coefficient on High Incentive is large and positive shows that WTP is on average much higher when incentives are high.³²

Column (2) includes a covariate for the delay until the part-2 survey, $\text{Ln}(\text{P2 Delay})$, as well as the interaction $\text{Ln}(\text{P2 Delay}) \times \text{Extra } \1 . The coefficient on $\text{Ln}(\text{P2 Delay})$ reveals that participants are willing to pay less for reminders when the task is further out in the future, although the coefficient on $\text{Ln}(\text{P2 Delay}) \times \text{Extra } \1 is directionally positive. Column (3) shows that these effects are amplified at the high incentive level. Column (4) returns to the specification in Column (1) but shows that controlling for the day participants completed part 1 of the study does not impact our estimates. Finally, Columns (5) and (6) replicate the specifications in (1) and (2) using a Tobit model to account for participants who were top-coded at the higher incentive levels within each group (\$4 or \$12). The results are quantitatively and qualitatively very similar in columns (5) and (6).³³

3.2.4 Do Participants Invest in Reminders Optimally?

Taken together, the results show that, for a \$1 increase in the incentive for completing the survey, participants are on average willing to pay around 7 cents more for the reminders at low incentive levels but only 2 cents more at high incentive levels.

Part 1 of Theorem 1 states that if participants are optimally investing in the BE, then a small increase dr in incentives for completing the survey should increase WTP for the BE by approximately dr times the increase in the likelihood of survey completion due to the BE. As Column (6) of Table 3 shows, a \$1 change in incentives does not have a large effect on task completion, which implies that we can apply the test in part 1 of Corollary 1 to a \$1 change in incentives. On average, reminders had a 29 and a 16 percentage point effect on survey completion in the low-incentive and high-incentive conditions, respectively. This is significantly larger than the respective \$0.07 and \$0.02 changes in WTP with respect to an

³²While not as natural a test of the theory since it spans a much larger increase in incentives, the 93 cent increase in WTP reflects an \$8 increase in the incentive level, or $\frac{93}{8} = 11.63$ cents per dollar, which is not that much larger than the 7 cent increase identified above. An additional difference that confounds this particular analysis, however, is that the MPL we use to elicit WTP for the high incentive levels was different (i.e., contained 33 rows where WTP increased in 75-cent increments) than the MPL for the low incentive levels (i.e., where WTP increased in 25-cent increments). For this reason as well, the estimate on Extra \$1 is the more natural test of the theory.

³³Appendix Table A.7 shows that the results about WTP for reminders do not differ (at conventional levels of statistical significance) when we restrict to the first 2, 4, 6, or 8 MPL screens that participants encountered. This suggests that the within-subject design did not introduce demand or anchoring effects that altered our estimates. Appendix Table A.8 presents a less-parametric regression analysis that separately estimates the effect of each of the possible delays, as well as its interaction with the Extra \$1 covariate, on WTP for reminders.

Table 4: The Effect of Incentive and Delay on Willingness to Pay for Reminders

	WTP for Reminders (\$)					
	(1)	(2)	(3)	(4)	(5)	(6)
Extra \$1	0.07*** (0.017)	0.05 (0.051)	0.06* (0.036)	0.07*** (0.017)	0.10*** (0.018)	0.07 (0.052)
High Incentive	0.93*** (0.077)	0.93*** (0.077)	1.14*** (0.112)	0.93*** (0.077)	0.93*** (0.077)	0.93*** (0.077)
Extra \$1 × High Incentive	-0.06 (0.048)	-0.06 (0.048)	-0.08 (0.109)	-0.06 (0.048)	-0.07 (0.048)	-0.07 (0.048)
Ln(P2 Delay)		-0.08*** (0.022)	-0.03*** (0.012)			-0.08*** (0.022)
Extra \$1 × Ln(P2 Delay)		0.01 (0.021)	0.00 (0.014)			0.01 (0.021)
High Incentive × Ln(P2 Delay)			-0.09** (0.035)			
Extra \$1 × Ln(P2 Delay) × High Incentive			0.01 (0.041)			
Constant	0.51*** (0.032)	0.69*** (0.057)	0.59*** (0.041)	0.51*** (0.032)	0.50*** (0.032)	0.69*** (0.057)
Observations	36,896	36,896	36,896	36,896	36,896	36,896
Number of Participants	2,306	2,306	2,306	2,306	2,306	2,306
Specification	OLS	OLS	OLS	OLS	Tobit	Tobit
P1 Date FE				X		

This table estimates the effect of incentive, the natural log of delay (in days), and having an incentive in the high-incentive pair on the willingness to pay for reminders. The extra \$1 variable is an indicator for an incentive of \$4 or \$12. Participants in the high-incentive pair had a completion incentive of \$11 or \$12, compared to the low-incentive pair of \$3 or \$4. The “High Incentive” variable is an indicator for whether participants were in the high-incentive pair group. Column (1) shows OLS estimates for incentive, having an incentive in the high-incentive pair, and incentive interacted with having an incentive in the high-incentive pair; Column (2) maintains the specification in Column (1) and adds the natural log of delay; Column (3) adds an interaction between incentive, the natural log of delay, and an indicator for the high-incentive pair, as well as an interaction between the natural log of delay and the high-incentive pair, and an interaction between the natural log of delay and incentive; Column (4) shows Column (1) with fixed effects for the date that part-1 of the survey was taken; Column (5) uses the same variates as in Column (1) and shows Tobit estimates with censors at -\$4 and \$4 for the low-incentive group and censors at -\$12 and \$12 for the high-incentive group; Column (6) uses the same variates as in Column (2) and shows Tobit estimates with censors at -\$4 and \$4 for the low-incentive group and censors at -\$12 and \$12 for the high-incentive group. Standard errors, clustered at the participant level, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

additional \$1 incentive in these conditions (Wald test $p < 0.01$).³⁴

We can also apply approximation (8) of Corollary 1 across each of the eight different incentive-delay pairs (i.e., $\{low\ and\ high\} \times \{2\ days,\ 1\ week,\ 3\ weeks,\ and\ 6\ weeks\}$) generated in our experiment to estimate equation (10) from Section 1.4. A trivial application of the corollary is that under those same assumptions, equation (10) can be rewritten as

$$\bar{W}(r + \Delta|\xi) - \bar{W}(r|\xi) = \beta_0 + \beta_1 \frac{D(z = 1|r, \xi) + D(z = 1|r + \Delta, \xi)}{2}. \quad (12)$$

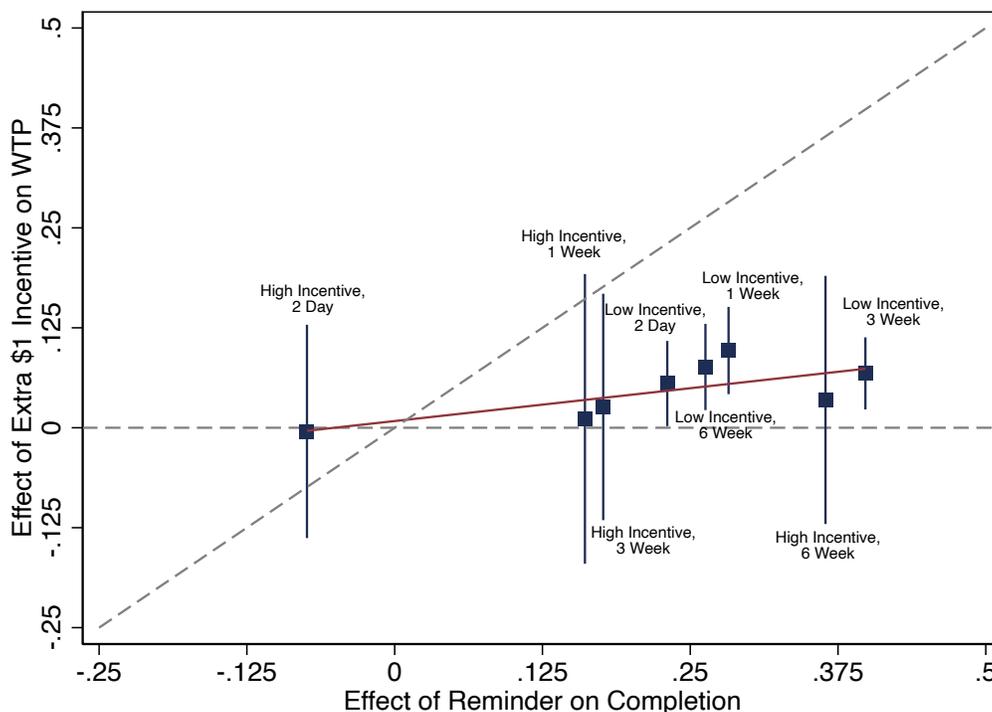
We estimate this equation by treating the eight different experimental conditions as variation in ξ . We then regress the change in WTP with respect to a \$1 change in incentives on the estimated effect of reminders, $D(z = 1|r, \xi)$, in each condition. Formally, this procedure is equivalent to a two-stage least squares (2SLS) estimator, where the eight different experimental conditions are instruments for the effects of reminders, and where the dependent variable is the change in WTP with respect to a \$1 change in incentives.

Figure 4 provides a visualization of the second stage of this 2SLS estimator. On the x -axis, this figure shows the estimated effect of reminders on survey completion rates as reported in Panel B of Figure 2. On the y -axis, this figure shows the estimated increases in willingness to pay for the reminders as the incentive increases by \$1, together with the 95% confidence intervals of the estimates. If participants were optimally valuing the reminder technology, the WTP for the reminder would be on the 45-degree line (e.g., such that when reminders increase survey completion by 25 percentage points, the willingness to pay for the reminders increases by \$0.25 with a \$1 increase in incentives). Instead, our estimates are far from the 45-degree line. Estimated WTP is below the 45-degree line for seven of the eight estimates, and the 95% confidence intervals exclude the 45-degree line in six of the estimates. An estimate of equation (12) yields $\beta_0 = 0.01$ and $\beta_1 = 0.16$. That is, perceptions of the effects of reminders are attenuated toward a prior mean of approximately 0, by a factor of 84 percent. In line with the discussion in Section 1.4, this evidence may be more consistent with a model in which people’s prior (or “default,” in the sense of Gabaix, 2014) perceptions of the effects of the BE in this setting are systematically biased. For a prior mean of approximately 0 to be an unbiased prior, it would have to be that reminders often have no effect, and sometimes even have negative effects on task completion.

³⁴Appendix Table A.17 explores whether risk aversion can explain the deviations in experiment 2 between willingness to pay and the effect of reminders on survey completion. We find no difference in willingness to pay for reminders when comparing participants with relatively high versus low risk aversion, measured in terms of the number of risky choices they select in gambles in the part 2 survey. These results hold both for the low incentive and high incentive groups. Thus, risk aversion does not appear to drive our results.

Findings from Experiment 2. Willingness to pay for reminders increased with the size of the bonus for survey completion, but the increase was too small relative to the null of correct valuation of attention costs. Specifically, the bias parameter θ from the parametric model in (9) is estimated to be 0.16; that is, the responsiveness of valuations for reminders was 16% as large as the ex-post optimal benchmark implies.

Figure 4: Effect of Reminders on Completion vs. Effect of Extra \$1 on WTP



This figure displays estimates and standard errors for the effect of reminders on whether part 2 of the survey was completed, and the effect of an increase of \$1 in the part-2 survey incentive on the willingness to pay for reminders. The figure only includes participants who were randomly assigned to receiving or not receiving reminders. The lines represent 95% confidence intervals, which are computed from the standard errors clustered by the date part 1 of the study was taken.

4 Learning Experiment

Results from our first two experiments reveal that individuals appear to undervalue bandwidth enhancements (BEs). A key question is whether individuals can learn to accurately value BEs through experience with them. We examine this in our third experiment. In addition, we illustrate the versatility of our methods by deploying them in a very different

setting: a design that builds on the Ambuehl et al. (2020) paradigm and resembles other psychometric tasks that have been used to generate state-dependent stochastic choice data for testing rational inattention theories (Dean and Neligh, 2018; Caplin et al., 2020; Caplin, 2021). Experimental instructions are in Screenshots Appendix F.3.

4.1 Design and Implementation

We ran the learning experiment on Prolific Academic in October and November of 2021.³⁵ The study involved a series of tasks. In each task, participants were presented with an image that showed a series of equations that were either correct (e.g., $10 + 12 = 22$) or incorrect (e.g., $10 + 12 = 23$). Participants were asked to indicate whether the image contained more correct or more incorrect equations with an incentive for accuracy. The *baseline task* included an image with 100 equations about which participants were told that either 60% of the equations were correct (and 40% incorrect) or 40% of the equations were correct (and 60% incorrect).³⁶ Participants completed two blocks of seven tasks each, and at least three tasks in each block were baseline tasks. One task of the 14 tasks the participant completed was randomly selected, and participants were paid according to their accuracy in that task as described below.

Before completing each block of seven tasks, participants were asked for their willingness to pay to make the task easier (i.e., to take advantage of a BE). In the *length* arm, participants could shorten the task so that there were only 10 equations in the image, rather than 100 equations. In the *discernibility* arm, participants could make the fraction of correct equations either 95% correct (and 5% incorrect) or 5% correct (and 95% incorrect), rather than 60% or 40%. Of the seven tasks in the block, three were baseline tasks and three were easy tasks (i.e., shorter tasks in the *length* arm and more discernible tasks in the *discernibility* arm). If the remaining task in the block was randomly selected for payment, then the participant's choices (i.e., responses on a multiple price list, as described below) would determine

³⁵In the recruitment materials, potential participants were informed that the study would require 20 minutes of their time, for which they would receive a guaranteed \$2.50. They were also informed that they would have the possibility of completing a bonus, and that the study had to be completed on a desktop or laptop computer using Chrome or Firefox as their web browser, which was necessary to ensure participants could see the tasks that were part of this study.

³⁶The images were automatically constructed with randomly generated equations following the protocol in Ambuehl et al. (2020). The computer selected the number of equations (e.g., 100 in the baseline tasks) and randomly selected one of the possible fractions to be the percentage of correct equations (e.g., either 60% or 40% in the baseline tasks). The two numbers on the left side of the equation were each randomly selected from the range 1 to 99. For the correct equations, the true result appeared on the right side. For the incorrect equations, a number was randomly drawn from the range 1 to 5, which was either added to or subtracted from the true result at random. The equations were then shuffled for display.

whether that task was a baseline task or an easier task.³⁷

Following the design of our survey-completion experiment, we elicited willingness to pay to reduce attention costs for different incentives for accuracy in the task. Participants were informed that the computer would randomly and independently select an accuracy bonus—paid to a participant for providing an accurate answer in the task—of either \$2, \$3, or \$4 for each block of seven tasks.³⁸ This procedure allowed us to elicit a participant’s WTP to make the task easier for three different accuracy bonus levels (i.e., before they knew which accuracy bonus would be relevant for that block). In particular, we elicited participants’ willingness to pay to make the task easier using a set of three multiple price lists (MPLs). The interface of the MPLs was similar to that in the survey experiment.

A key design feature of this experiment is the opportunity to learn about the value of making the task easier. We do this by having participants experience both baseline and easier tasks in the first block of seven tasks (the baseline and easier tasks were presented in a random order). Participants then repeat the exercise—providing willingness to pay to make the task easier at each of the three accuracy incentive levels—for the second block of tasks.

The amount of feedback participants receive about their performance in the first block varies by treatment. In the *control* treatment, participants were not provided with information about their performance in the first block. In the *feedback* treatment, by contrast, participants were told the fraction of baseline and easier tasks they answered accurately. Moreover, participants were told how those accuracy levels translated to expected earnings at the three different levels of incentives.

In all MPLs, the options for willingness to pay ranged from $-\$4$ to $\$4$. This range is analogous to the range in the survey-completion experiment, where the highest possible MPL amount corresponded to the highest reward for task completion. Analogous to the survey-completion experiment, we make the same conservative sample restriction to limit to individuals who were never top-coded at the smallest incentive (i.e., $\$2$) or bottom-coded at the largest incentive (i.e., $\$4$). The logic behind this restriction is the same as in the previous experiment: since these individuals cannot increase their willingness to pay as the accuracy bonus increases, it is possible that including these participants could lead to an attenuation bias in estimates of the effect of the accuracy bonus on willingness to pay for a BE.

³⁷If the task was not randomly selected for payment, then the difficulty of the task would be chosen at random. This protocol ensured that the multiple price list responses only affected the task and generated bonus payments in the case when participants were going to be paid based on their accuracy on that particular task.

³⁸At the start of the block, participants were told which accuracy bonus had been randomly selected to apply for the block of seven tasks. Each task screen also reminded participants of the accuracy bonus. This bonus was paid if the participant provided an accurate answer on the task randomly selected for payment.

4.2 Results

Our analysis involves the 1465 participants who completed the study, and who were not top- or bottom-coded as described above.³⁹

4.2.1 Impact of Task Difficulty on Accuracy

To analyze the causal effect of task difficulty on performance, we exclude data from tasks that were potentially endogenously determined by participants' WTP. That is, we exclude cases in which the "remaining task" in the block was selected to be the task that counts, and thus was affected by participants' preferences.

Table 5 presents OLS regressions that quantify the impacts of task difficulty on the likelihood of correctly identifying whether there are more correct or incorrect equations in the picture. The first two columns analyze performance in the first block, while the latter two columns analyze performance in the second block. Columns (1) and (3) show that performance did not differ significantly between the two blocks. In both blocks, participants correctly completed the baseline task approximately 70 percent of the time, were approximately 26 percentage points more likely to complete it correctly when the task was more discernible (i.e., had either 95% or 5% correct equations), and were approximately 19 percentage points more likely to complete it correctly when the task was shorter (i.e., had only 10 equations). On average, the discernibility effect was moderately larger than the length effect: by 6.5 percentage points in block 1 (Chi-square test $p = 0.00$) and by 4.9 percentage points in block 2 (Chi-square test $p = 0.00$). Columns (2) and (4) show that variation in our incentives did not have a significant effect on performance. This is consistent with the survey-completion experiment, where a \$1 change in incentives was too small to have a significant effect, despite behavior being overall elastic (and thus responsive to large changes). This result implies that the incentive changes we analyze are sufficiently small to utilize Corollary 1.

Data on decision times is consistent with the baseline task being more difficult. On average, participants spent 87 seconds, 44 seconds, and 37 seconds on the baseline, shorter, and more discernible tasks, respectively. Appendix Figure A.4 presents the CDFs of response times across the three types of tasks. Appendix Figure A.3 shows that participants who spent more time on the tasks were more likely to answer them accurately, particularly in the more

³⁹520 participants were excluded from the analysis because in at least one set of their MPL elicitation they were either top-coded at the lowest incentive (\$2) or bottom-coded at the highest incentive (\$4). Additionally, 125 participants were automatically screened out of the study (and prevented from participating further) because they failed attention checks, ensuring our pool of participants understood the instructions in our experiment. 13 individuals who encountered technical glitches were also not included.

Table 5: The Effect of Length and Discernibility on Getting a Task Correct

	Answered Task Correctly			
	(1)	(2)	(3)	(4)
More Discernible	0.26*** (0.010)	0.25*** (0.008)	0.26*** (0.010)	0.27*** (0.008)
Shorter	0.19*** (0.011)	0.20*** (0.009)	0.21*** (0.012)	0.19*** (0.010)
Length Arm	0.01 (0.015)		-0.04** (0.015)	
Incentive (\$)		-0.00 (0.006)		0.01 (0.006)
Constant	0.71*** (0.010)	0.72*** (0.018)	0.71*** (0.010)	0.67*** (0.018)
Observations	10,143	10,143	10,157	10,157
Number of Participants	1,465	1,465	1,465	1,465
Block	1	1	2	2

This table estimates the effect of shorter length (i.e., 10 equations) and increased discernibility (i.e., 95% or 5% correct) on getting a task correct in block 1 and block 2. Tasks that had their difficulty determined by a participant’s MPL choices have been excluded. The columns correspond to different regression specifications and blocks: Column (1) shows OLS estimates in block 1, Column (2) shows OLS estimates including the incentive level in dollars in block 1, and Columns (3) and (4) show analogous specifications for block 2. Standard errors, clustered at the participant level, are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

difficult, baseline task.

4.2.2 Willingness to Pay to Simplify Tasks

Table 6 presents OLS regressions that estimate how participants’ WTP to make the tasks easier varies with accuracy incentives and other experimental manipulations. Column (1) shows that participants in the *length* arm increased their WTP to make the task shorter by approximately \$0.10 for every dollar of extra accuracy incentive in block 1. However, this effect goes down to approximately \$0.03 in block 2, which is a significant difference of approximately \$0.07 (Chi-square $p = 0.06$). In the *discernibility* arm, each \$1 of accuracy incentive increases WTP by \$−0.01 and \$0.03 cents in blocks 1 and 2, respectively. Neither of these is statistically significantly different from zero, nor are they different from each other. In block 1, the difference between the coefficients on incentive in the length and

discernibility arms is \$0.11, which is statistically significant ($se = 0.045$, Chi-square $p = 0.017$).

Given that simplifying the task increased accuracy by approximately 19 to 26 percentage points in both arms, Theorem 1 implies that WTP for simplifying the task should increase by approximately \$0.19 to \$0.26. This is higher than the effects reported in Table 6. In each of the four block-arm pairs, the WTP increase is significantly smaller than the theoretical benchmark (Wald test $p < 0.01$ in all arms).

Columns (1) and (2) thus reveal three key insights. First, people in this experiment undervalue BEs, as in our first two experiments. Second, participants initially undervalue discernibility improvements more than length improvements. An ex-post rationalization is that decreasing the length of a task is a simple and relatively common form of simplification that most people are familiar with, whereas increasing discernibility is a more abstract and less-common form of simplification. Third, experience does not bring people's decisions more in line with the correct-perceptions benchmark formalized in Theorem 1. Directionally, people value discernibility improvements more in block 2, but this is not statistically significant at conventional levels. In fact, there is stronger evidence that in block 2 of the length arm, there is *more* deviation from the Theorem 1 benchmark than in block 1. The difference between the coefficients on Incentive \times Block 2 in columns (1) and (2) is -0.11 ($se = 0.052$, Chi-square $p = 0.04$).

Columns (3) and (4) provide insight into why the deviation increases with experience in the length arm. As shown in column (3), the deviation increases primarily among the participants who received the feedback treatment. For participants in the control treatment, the coefficient on incentive increases by an insignificant 0.01 ($se = 0.050$). However, as the coefficient on the interaction Incentive \times Block 2 \times Feedback shows, the impact of experience is a statistically significant -0.16 ($se = 0.074$, $p = 0.034$) for participants in the feedback treatment. By contrast, feedback has no effect on participants in the discernibility arm, suggesting that participants in that arm have a very strongly held prior that discernibility would not affect their performance.

Columns (5) and (6) further explore the negative effect of feedback in the length arm. Column (5) restricts to participants who did not perform better on the shorter tasks than on the baseline tasks in block 1; column (6) restricts to participants who did perform better. The coefficient on Incentive \times Block 2 is nearly identical in those two columns, implying that these two groups of participants were not differentially affected by experience in the control condition. However, the negative coefficient on Incentive \times Block 2 \times Feedback is twice as large in magnitude in column (5) than in column (6). This result suggests that a reason for

Table 6: The Effect of Incentive, Block, and Feedback on Willingness to Pay for Easier Tasks

	Willingness to Pay (\$)					
	(1)	(2)	(3)	(4)	(5)	(6)
Incentive (\$)	0.10*** (0.034)	-0.01 (0.030)	0.04 (0.045)	0.02 (0.042)	0.07 (0.049)	0.02 (0.072)
Incentive (\$) × Block 2	-0.07* (0.037)	0.04 (0.037)	0.01 (0.050)	0.03 (0.052)	0.00 (0.063)	0.01 (0.074)
Incentive (\$) × Feedback			0.11 (0.067)	-0.07 (0.060)	0.12 (0.089)	0.10 (0.099)
Incentive (\$) × Block 2 × Feedback			-0.16** (0.074)	0.01 (0.074)	-0.22** (0.100)	-0.11 (0.108)
Block 2	0.15 (0.125)	-0.18 (0.123)	-0.09 (0.177)	-0.20 (0.176)	-0.16 (0.235)	-0.04 (0.259)
Block 2 × Feedback			0.49* (0.249)	0.05 (0.247)	0.87** (0.341)	0.17 (0.355)
Feedback			-0.22 (0.248)	0.22 (0.213)	-0.36 (0.342)	-0.11 (0.353)
Constant	0.20 (0.124)	0.53*** (0.107)	0.31* (0.172)	0.42*** (0.147)	0.32 (0.212)	0.31 (0.260)
Observations	3,996	4,794	3,996	4,794	1,788	2,208
Number of Participants	666	799	666	799	298	368
Block 1 Acc. Diff.	All	All	All	All	≤ 0	> 0
Arm	Length	Discernibility	Length	Discernibility	Length	Length

This table estimates the effect of accuracy incentives, block order, and whether the participant received performance feedback on willingness to pay for an easier task (i.e., a shorter task in the length arm and a more discernible task in the discernibility arm). Columns (5) and (6) restrict participants by their block-1 accuracy difference between the baseline and easier tasks, which equals the difference between the percentage of easier tasks and baseline tasks answered correctly in block 1. The mean block-1 accuracy differences for participants in Columns (5) and (6) are -0.06 and 0.40 , respectively. Column (1) shows OLS estimates for incentive, block, and the interaction of incentive and block order for participants in the length arm; Column (2) repeats this analysis in the discernibility arm; Column (3) maintains the specification in Column (1) and the restriction to participants in the length arm while adding whether feedback was received and the interactions between feedback, block, and incentive; Column (4) shows the OLS estimates in Column (3) for participants in the discernibility arm; Column (5) shows the OLS estimates in Column (3) for participants in the length arm who had a block-1 accuracy difference of less than or equal to 0; Column (6) shows the OLS estimates in Column (3) for participants in the length arm who had a block-1 accuracy difference of greater than 0. Standard errors, clustered at the participant level, are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

the negative effect of the feedback treatment is participants being disappointed to learn that their block-1 performance was not higher on shorter tasks.⁴⁰

The results about WTP in columns (3)–(6) of Table 6 cannot be explained by differences in block-2 performance among the different groups analyzed in those columns. Columns (1) and (2) of Appendix Table A.11 show that neither experience nor feedback altered participants’ performance on the baseline versus easy tasks. Columns (3) through (5) of Appendix Table A.11 restrict to the length arm, and show that: (i) consistent with mean reversion, participants who did not perform better on the shorter tasks than on the baseline tasks in block 1 *improved* their relative performance on the easy tasks in block 2, and (ii) feedback did not reduce the difference in block-2 performance between the easy and baseline tasks among these participants. The overall negative effect of the feedback treatment suggests participants updated their beliefs in a quasi-Bayesian manner by overweighting disappointing experiences.

Interestingly, columns (3) and (5) of Table 6 also suggest that, while experience and feedback led participants to underestimate the effect of task simplification on their performance at the incentives in the experiment, it did increase their overall WTP to simplify the task. This result illustrates the theoretical discussion in Section 1.3.1 about how accurate perceptions of total costs (including fixed costs) are not characterized in Theorem 1—perceptions of the effects of the BE at the current incentive level r are not a sufficient statistic for perceptions of total costs. Appendix Table A.12 shows that participants who did not perform better on the shorter tasks than on the baseline tasks in block 1 spent almost two minutes longer on the baseline versus easier tasks in block 1. Subsequently, they spent approximately 30 to 45 seconds less time on the baseline versus easier tasks in block 2, relative to the participants who did perform better on the shorter tasks in block 1. This suggests that the participants who did not perform better on the shorter tasks in block 1 also incurred significantly larger total costs on the baseline tasks. The feedback treatment may have helped prime this realization by inducing participants to further reflect on the differences between the baseline and shorter tasks.⁴¹ This illustrates the Section 1.3.1 discussion about how additional treatments such as our feedback treatment can provide additional insights into perceptions about total costs, complementing the tests in Theorem 1.⁴²

⁴⁰Alternatively, it could be consistent with the suggestive evidence from column (6) of Appendix Table A.11 that feedback may have decreased performance for these participants—which suggests that these participants overall chose to adopt attention strategies that would decrease relative performance in the shorter task.

⁴¹Appendix Table A.12 also shows that feedback treatment had a small negative effect on the time taken on the baseline tasks in block 2 for participants who performed better on the shorter tasks than on the baseline tasks in block 1. There is no effect on participants who did not perform better on the shorter tasks in block 1.

⁴²Appendix Tables A.9 and A.10 replicate Table 6 using Tobit models and dropping participants with the 10

Findings from Experiment 3. *The responsiveness of willingness to pay for the length BE was 50% as large as the ex-post optimal benchmark implies. Willingness to pay for the discernibility BE did not respond to incentive size at all, implying a responsiveness 0% as large as the ex-post optimal benchmark implies. Experience via feedback on past performance did not bring people’s decisions more in line with the correct-valuations benchmark.*

5 Conclusion

While a large and growing literature shows that attention-increasing interventions such as reminders and plan-making tools can have significant effects on economically important behaviors, this literature rarely asks the question of whether individuals value and deploy these tools optimally. This paper addresses this question with three theory-driven, quantitative tests. We find that individuals’ demand for attention-increasing tools is *qualitatively* consistent with the predictions of optimal management of limited attention, but it is quantitatively inconsistent with fully optimizing choice of attention. This suggests that individuals are uncertain and/or systematically biased about their attention cost functions. While this under-valuation of bandwidth enhancements may be context dependent, our methods can be applied more broadly to explore how individuals value attention-increasing technologies across various domains.

Our methods are immediately portable to other settings where the impact of reminders and planning prompts has already been documented, such as in medical compliance, savings, loan repayment, and voting (see footnote 6 for references). In addition, as exemplified by our third experiment, our methodology can be used to test whether people understand their production functions for attention-consuming tasks in field settings such as those of Dean (2019), Kaur et al. (2021), or Bessone et al. (forthcoming). More generally, our tests could be applied to any setting that involves domains of behavior that feature “intermediate” actions. For example, our methods could be used to quantify whether students fully understand the relationship between studying and test performance, whether individuals understand the link between education and earnings, or whether individuals properly invest in “good habits.”

Finally, our results that people do not select their attention environments optimally are consistent with the hypothesis that people might misoptimize their attention strategies in other ways. For example, people might misoptimize their choice of decision boundaries in sequential information acquisition problems (Reshidi et al., 2022) or might not choose their

percent fastest average task times by the length and discernibility arms separately. The results are quantitatively and qualitatively similar.

signal structure optimally in complex games with asymmetric payoffs (as in, e.g., Suen, 2004). These and other possibilities suggest an exciting research agenda on the question of whether attention is produced optimally.

References

- ABALUCK, J. AND A. ADAMS-PRASSL (2021): “What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses,” *The Quarterly Journal of Economics*, 136, 1611–1663.
- ACLAND, D. AND M. R. LEVY (2015): “Naiveté, projection bias, and habit formation in gym attendance,” *Management Science*, 61, 146–160.
- ALLCOTT, H., J. KIM, D. TAUBINSKY, AND J. ZINMAN (2022): “Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending,” *Review of Economic Studies*, 89, 1041–1084.
- ALTMANN, S., A. GRUNEWALD, AND J. RADBRUCH (forthcoming): “Interventions and Cognitive Spillovers,” *Review of Economic Studies*.
- ALTMANN, S. AND C. TRAXLER (2014): “Nudges at the Dentist,” *European Economic Review*, 72, 19–38.
- AMBUEHL, S., A. OCKENFELS, AND C. STEWART (2020): “Who Opt In?” *Working Paper no. 7091, CESifo*.
- BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): “Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition,” *American Economic Review*, 106, 1437–75.
- BERNHEIM, B. D. AND D. TAUBINSKY (2018): *Behavioral Public Economics*, New York: Elsevier, vol. 1, 381–516.
- BESSONE, P., G. RAO, F. SCHILBACH, H. SCOFIELD, AND M. TOMA (forthcoming): “The Economic Consequences of Increasing Sleep among the Urban Poor,” *The Quarterly Journal of Economics*.
- BRONCHETTI, E. T., D. B. HUFFMAN, AND E. MAGENHEIM (2015): “Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination,” *Journal of Economic Behavior and Organization*, 116, 270–291.
- CALZOLARI, G. AND M. NARDOTTO (2017): “Effective Reminders,” *Management Science*, 63, 2915–2932.
- CAPLIN, A. (2016): “Measuring and Modeling Attention,” *Annual Review of Economics*, 8, 379–403.
- (2021): “Economic Data Engineering,” *Working Paper no. 29378, National Bureau of Economic Research*.
- CAPLIN, A., D. CSABA, J. LEAHY, AND O. NOV (2020): “Rational Inattention, Competitive Supply, and Psychometrics,” *Quarterly Journal of Economics*, 135, 1681–1724.
- (2015): “Revealed Preference, Rational Inattention, and Costly Information Acquisition,” *American Economic Review*, 105, 2183–2203.
- CARRERA, M., H. ROYER, M. STEHR, J. SYDNOR, AND D. TAUBINSKY (2018): “The Limits of Simple Implementation Intentions: Evidence from a Field Experiment on Making Plans to Exercise,” *Journal of Health Economics*, 62, 95–104.

- (2022): “Who Chooses Commitment? Evidence and Welfare Implications,” *Review of Economic Studies*, 89, 1205–1244.
- CASTLEMAN, B. L. AND L. C. PAGE (2016): “Freshman Year Financial Aid Nudges: An Experiment to Increase FAFSA Renewal and College Persistence,” *Journal of Human Resources*, 51, 389–415.
- CHETTY, R., A. LOONEY, AND K. KROFT (2009): “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 99, 1145–1177.
- DAMGAARD, M. T. AND C. GRAVERT (2018): “The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising,” *Journal of Public Economics*, 157, 15–26.
- DE OLIVEIRA, H., T. DENTI, M. MIHM, AND K. OZBEK (2017): “Rationally Inattentive Preferences and Hidden Information Costs,” *Theoretical Economics*, 12.
- DEAN, J. T. (2019): “Noise, Cognitive Function and Worker Productivity,” *Working Paper*.
- DEAN, M. AND N. NELIGH (2018): “Experimental Tests of Rational Inattention,” *Working Paper*.
- DELLAVIGNA, S. AND U. MALMENDIER (2004): “Contract Design and Self-Control: Theory and Evidence*,” *The Quarterly Journal of Economics*, 119, 353–402.
- DELLAVIGNA, S. AND D. POPE (2017): “What Motivates Effort? Evidence and Expert Forecasts,” *The Review of Economic Studies*, 85, 1029–1069.
- ERICSON, K. (2011): “Forgetting We Forget: Overconfidence and Memory,” *Journal of the European Economic Association*, 9, 43–60.
- (2017): “On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation,” *Journal of the European Economic Association*, 15, 692–719.
- GABAIX, X. (2014): “A Sparsity-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 129, 1661–1710.
- (2019): “Behavioral Inattention,” in *Handbook of Behavioral Economics*, ed. by D. Bernheim, S. DellaVigna, and D. Laibson, Elsevier, vol. 2, 261–343.
- GABAIX, X., D. LAIBSON, G. MOLOCHE, AND S. WEINBERG (2006): “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96, 1043–1068.
- GAGNON-BARTSCH, T., M. RABIN, AND J. SCHWARTZSTEIN (2021): “Channeled Attention and Stable Errors,” *Working Paper*.
- GILBERT, S. J. (2015a): “Strategic offloading of delayed intentions into the external environment,” *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 68, 971–992.
- (2015b): “Strategic use of reminders: Influence of both domain-general and task-specific metacognitive confidence, independent of objective memory ability,” *Consciousness and Cognition*, 33, 245–260.
- GILBERT, S. J., A. BIRD, J. M. CARPENTER, S. M. FLEMING, C. SACHDEVA, AND P.-C. TSAI (2020): “Optimal use of reminders: Metacognition, effort, and cognitive offloading,” *Journal of Experimental Psychology: General*, 149, 501–517.
- GOLLWITZER, P. M. AND P. SHEERAN (2006): “Implementation Intentions and Goal Achievement: A Meta-analysis of Effects and Processes,” *Advances in Experimental Social Psychology*, 38, 69–119.
- HANNA, R., S. MULLAINATHAN, AND J. SCHWARTZSTEIN (2014): “Learning Through Noticing:

- Theory and Evidence from a Field Experiment,” *Quarterly Journal of Economics*, 129, 1311–1353.
- HEIDHUES, P., B. KÓSZEGI, AND P. STRACK (2018): “Unrealistic Expectations and Misguided Learning,” *Econometrica*, 86, 1159–1214.
- HOFFMAN, M. AND S. V. BURKS (2020): “Worker overconfidence: Field evidence and implications for employee turnover and firm profits,” *Quantitative Economics*, 11, 315–348.
- KAHNEMAN, D. AND A. TVERSKY (1982): “Intuitive Prediction: Biases and Corrective Procedures,” in *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, chap. Intuitive Prediction: Biases and Corrective Procedures, 414–421.
- KARLAN, D., M. MCCONNELL, S. MULLAINATHAN, AND J. ZINMAN (2016a): “Getting to the Top of Mind: How Reminders Increase Saving,” *Management Science*, 62, 3393–3411.
- KAUR, S., S. MULLAINATHAN, S. OH, AND F. SCHILBACH (2021): “Do Financial Concerns Make Workers Less Productive?” *Working Paper no. 28338, National Bureau of Economic Research*.
- MACKOWIAK, B., F. MATEJKA, AND M. WIEDRHOLT (2018): “Dynamic rational inattention: Analytical results,” *Journal of Economic Theory*, 176, 650–692.
- MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (forthcoming): “Rational Inattention: A Review,” *Journal of Economic Literature*.
- MARTIN, D. (2016): “Rational Inattention in Games: Experimental Evidence,” *Working Paper*.
- MARX, B. M. AND L. J. TURNER (2019): “Student Loan Nudges: Experimental Evidence on Borrowing and Educational Attainment,” *American Economic Journal: Economic Policy*, 11, 108–141.
- MATEJKA, P., J. STEINER, AND C. STEWART (2017): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” *Econometrica*, 85, 521–553.
- MATĚJKA, F. AND A. MCKAY (2015): “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 105, 272–298.
- MILGROM, P. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice sets,” *Econometrica*, 70, 583–601.
- MILKMAN, K. L., J. BESHEARS, J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2011): “Using implementation intentions prompts to enhance influenza vaccination rates,” *Proceedings of the National Academy of Sciences*, 108, 10415–10420.
- MORRISON, W. AND D. TAUBINSKY (forthcoming): “Rules of Thumb and Attention Elasticities: Evidence from Under- and Overreaction to Taxes,” *Review of Economics and Statistics*.
- MULLAINATHAN, S. (2002): “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 117, 735–774.
- NAFZIGER, J. (2020): “Spillover Effects of Nudges,” *Economics Letters*, 190.
- NICKERSON, D. W. AND T. ROGERS (2010): “Do You Have a Voting Plan? Implementation Intentions, Voter Turnout, and Organic Plan Making,” *Psychological Science*, 21, 194–199.
- OREOPOULOS, P., R. W. PATTERSON, U. PETRONIJEVIC, AND N. G. POPE (forthcoming): “Low-Touch Attempts to Improve Time Management among Traditional and Online College Students,” *Journal of Human Resources*.
- RABIN, M. (2000): “Risk aversion and expected-utility theory: A calibration theorem,” *Econometrica*, 68, 1281–1292.
- RESHIDI, P., A. LIZZERI, L. YARIV, J. CHAN, AND W. SUEN (2022): “Individual and Collective Information Acquisition: An Experimental Study,” *working paper*.

- SHENHAV, A., M. BOTVINICK, AND J. D. COHEN (2013): “The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function,” *Neuron*, 79, 217–240.
- SHENHAV, A., S. MUSSLICK, F. LIEDER, W. KOOL, T. L. GRIFFITHS, J. D. COHEN, AND M. M. BOTVINICK (2017): “Toward a Rational and Mechanistic Account of Mental Effort,” *Annual Review of Neuroscience*, 40, 99–124, PMID: 28375769.
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50, 665–690.
- STRACK, P. AND D. TAUBINSKY (2022): “Dynamic Preference ‘Reversals’ and Time Inconsistency,” *working paper*.
- SUEN, W. (2004): “The Self-Perpetuation of Biased Beliefs,” *Economic Journal*, 114, 377–396.
- TASOFF, J. AND R. LETZLER (2014): “Everyone Believes in Redemption: Nudges and Overoptimism in Costly Task Completion,” *Journal of Economic Behavior and Organization*, 107, 107–122.
- TAUBINSKY, D. (2014): “From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice,” *working paper*.
- TAUBINSKY, D. AND A. REES-JONES (2018): “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” *The Review of Economic Studies*, 85, 2462–2496.
- ZHANG, C. Y., J. HEMMETER, J. B. KESSLER, R. D. METCALFE, AND R. WEATHERS (2021): “Nudging Timely Wage Reporting: Field Experimental Evidence from the United States Social Supplementary Income Program,” *Working Paper no. 27875, National Bureau of Economic Research*.