

RCTs to Scale: Comprehensive Evidence from Two Nudge Units*

Stefano DellaVigna Elizabeth Linos
UC Berkeley and NBER UC Berkeley

January 2021

Abstract

Nudge interventions have quickly expanded from academic studies to larger implementation in so-called Nudge Units in governments. This provides an opportunity to compare interventions in research studies, versus at scale. We assemble a unique data set of 126 RCTs covering 23 million individuals, including all trials run by two of the largest Nudge Units in the United States. We compare these trials to a sample of nudge trials in academic journals from two recent meta-analyses. In the Academic Journals papers, the average impact of a nudge is very large—an 8.7 percentage point take-up effect, which is a 33.5% increase over the average control. In the Nudge Units sample, the average impact is still sizable and highly statistically significant, but smaller at 1.4 percentage points, an 8.1% increase. We document three dimensions which can account for the difference between these two estimates: (i) statistical power of the trials; (ii) characteristics of the interventions, such as topic area and behavioral channel; and (iii) selective publication. We estimate a meta-analysis model and find strong evidence of selective publication in academic journals. Correcting for it brings down the estimated impact of a nudge in academic journals to 3.2 percentage points, still larger than the Nudge Unit estimate, though not significantly so. Counterfactuals show that low statistical power in the Academic Journals sample exacerbates the impact of selective publication.

*We are very grateful to the Office of Evaluation Sciences and Behavioral Insights Team North America for supporting this project and for countless suggestions and feedback. We thank Johannes Abeler, Isaiah Andrews, Oriana Bandiera, Shlomo Benartzi, John Beshears, Abel Brodeur, David Card, Benjamin Enke, Etan Green, Johannes Hermle, Eric Johnson, Maximilian Kasy, David Laibson, George Loewenstein, Rachael Meager, Katherine Milkman, Gautam Rao, Adam Sacarny, Robert Sugden, Cass Sunstein, Richard Thaler, Eva Vivalt, Richard Zeckhauser and participants in seminars at APPAM 2020, BEAM 2020, the Behavioral Science and Regulation Conference, Behavioural Insights Team (global), CHIBE 2020, ideas42, Harvard University, the LSE, the NBER Summer Institute 2020 (Public Economics and Household Economics), the University of Chicago, the University of Michigan, the University of Pittsburgh, University of California, Berkeley, the University of Rotterdam, and the University of Zurich for helpful comments. We are grateful to Margaret Chen and Woojin Kim and a team of undergraduate research assistants at UC Berkeley for exceptional research assistance.

1 Introduction

Thaler and Sunstein (2008) define nudges as “*choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives.*” These light-touch behavioral interventions—including simplification, personalization, and social-norm comparison—have become common in the literature, spanning hundreds of papers in fields such as economics, political science, public health, decision-making, and marketing.

Soon after researchers embraced these interventions, nudges also went mainstream within governments in larger-scale applications. While behavioral interventions were already being used on a case-by-case basis within government, the launch of ideas42 in the US in 2008, the UK’s Behavioural Insights Team (BIT) in 2010 (see, e.g., Halpern, 2015), and the Office of Evaluation Sciences (OES) in 2015 spurred an explosion of government teams dedicated to using behavioral science to improve government services. As of last count, there are more than 200 such units globally (Figure A1, OECD, 2017).

The rapid expansion of behavioral interventions through Nudge Units offers a unique opportunity to compare the impact of interventions as implemented by researchers to the larger roll-out of similar interventions “at scale” (Muralidharan and Niehaus, 2017). Do nudges impact, for example, take up of vaccinations, contribution to retirement plans, or timely payment of fines similarly in interventions by academic researchers and in larger-scale implementations within governments? Understanding how RCTs scale is a key question, as researchers and policy-makers build on the results of smaller interventions to plan larger implementations. To the best of our knowledge, this comparison of Nudge Unit experiments to the papers in the literature has not been possible so far, given the lack of comprehensive data on Nudge Unit interventions.

In this paper, we present the results of a unique collaboration with two major “Nudge Units”: BIT North America, which conducts projects with multiple US local governments and OES which collaborates with multiple US Federal agencies. Both units kept a comprehensive record of all trials from inception in 2015. As of July 2019, they conducted a total of 165 trials testing 349 nudge treatments and affecting over 37 million participants. In a remarkable case of administrative transparency, each trial had a trial report, including in many cases a pre-analysis plan. The two units worked with us to retrieve the results of all trials, 87 percent of which have not been documented in working

papers or academic publications.

This evidence differs from a traditional meta-analysis in two ways: (i) the large majority of these findings have not previously appeared in academic journals; (ii) we document the entirety of trials run by these units, with no scope for selective publication.

We restrict our data set to RCTs (excluding 13 natural experiments and differences-in-differences) and require that the trials have a clear control group (excluding 15 trials), do not use financial incentives (3 trials excluded), and have a binary outcome as dependent variable (excluding 8 trials). The last restriction allows us to measure the impact with a common metric—the percentage point difference in outcome, relative to the control. Finally, we exclude interventions with default changes (2 nudges in 1 trial). This last restriction ensures that the nudge treatments are largely comparable, consisting typically of a combination of simplification, personalization, implementation intention prompts, reminders, and social norm comparisons introduced in administrative communication. Examples of such interventions include a letter encouraging service-members to re-enroll in their Roth Thrift Savings Plans, and a post-card from a city encouraging people to fix up their homes in order to meet code regulations. Our final sample includes 126 trials, involving 243 nudges and collectively impacting over 23 million participants.

We compare these trials to nudges in the literature, leaning on two recent papers summarizing over 100 published nudge RCTs across many settings (Benartzi et al., 2017 and Hummel and Maedche, 2019). We apply identical restrictions as in the Nudge Units sample, leaving a final sample of 26 RCTs, including 74 nudge treatments collectively affecting 505,337 participants. While this sample is fairly representative of the type of nudges in the literature, the features of these interventions do not perfectly match with the treatments implemented by the Nudge Units, a difference to which we return below.

What do we find? In the 26 papers in the Academic Journals sample, we compute the average (unweighted) impact of a nudge across the 74 nudge interventions. On average, a nudge intervention increases take up by 8.7 (s.e.=2.5) percentage points (pp.), a 33.5 percent increase over the average control group take up of 26.0 percent.

Turning to the 126 trials by Nudge Units, we estimate an unweighted impact of 1.4 pp. (s.e.=0.3), an 8.1 percent increase over an average control group take up of 17.2 percent. While this impact is highly statistically significantly different from 0 and sizable, it is about one sixth the size of the estimate in academic papers.

What accounts for this large difference? We document three key distinguishing

dimensions (although other characteristics also play a role): (i) statistical power of the trials; (ii) characteristics of the nudge interventions; and (iii) selective publication. We then propose a model that can account for differences in these three dimensions, and explore the relative role of each by simulating counterfactuals. While the model is aimed at our specific setting, it likely has implications for other instances in which experimental evidence is collected by researchers, and then rolled out “at scale”.

First, we document a large difference in the sample size and thus statistical power of treatment arms within the trials. The median nudge intervention in the Academic Journals sample has a treatment arm sample size of 484 participants and a minimum detectable effect size (MDE, the effect size that can be detected with 80% power) of 6.3 pp. In contrast, the interventions in the Nudge Units have a median treatment arm sample size of 10,006 participants and MDE of 0.8 pp. Thus, the statistical power in the Academic Journals sample is an order of magnitude smaller.

We propose one way to interpret this difference. In economics, 80% statistical power calculations are commonplace and constitute, for example, the large majority of formal power calculations on the AEA Registry. Holding this criterion constant, the differences in sample size may arise optimally if academic researchers expect much larger effect sizes and therefore are comfortable with a larger MDE *ex ante*. While we do not observe whether the academic trials in our sample were designed based on power calculations, we were able to collect forecasts on expected nudge effect sizes from a survey of academic researchers and nudge practitioners. The median nudge practitioner expects an average impact of 1.95 pp. for the Nudge Unit trials, remarkably in line with our findings and with the MDE for that sample. Conversely, the median academic researcher expects a larger effect size of 4.0 pp. for the Nudge Units sample and of 7.0 pp. for the Academic Journals sample. This suggests that academics’ optimistic expectations about nudge effect sizes may shape their trial design, and thus the differences in statistical power. Further, we document that in the Academic Journals sample researchers run more treatment arms per trial than in the Nudge Units sample, further diluting the (smaller) initial sample. Thus the smaller sample per arm in the Academic Journals sample does not merely reflect a capacity constraint.

Second, the nudge interventions in the two samples have different features, and some of these differences may account for the effect size discrepancy. We consider differences in (i) the degree of academic involvement, (ii) the types of behavioral mechanism used,

such as Academic Journals nudges having more in-person contact and choice design, (iii) the policy area, with the Academic Journal studies being more likely to tackle, for example, environmental policy questions; and (iv) the characteristics of the trials, such as institutional constraints. We document differences in all the above dimensions, but some of the differences either are small, or, such as academic involvement, do not predict differences in effect sizes. Differences in the type of behavioral mechanism used and in the policy area can account for some share of the difference in effect sizes.

Third, we consider selective publication as a function of statistical significance. In the Academic Journals sample, there are over 4 times as many studies with a t statistic between 1.96 and 2.96 for the most significant arm in a paper, versus studies where the most significant arm has a t between 0.96 and 1.96. Therefore, part of the different effect sizes in the two samples may come from the censoring of statistically insignificant trials in published papers. We stress that with “selective publication” we include not only whether a journal would publish a paper, but also whether a researcher would write up a study (the “file drawer” problem, e.g, Franco, Malhotra, and Simonovits, 2014). In the Nudge Units sample, all these selective steps are removed.

Building on this evidence, we estimate a model based on Andrews and Kasy (2019) which allows for selective publication of the Academic Journals trials, as well as for a different effect sizes in the two samples. This model also takes as inputs differences in statistical power. We find strong evidence of publication bias: we estimate that trials with no significant results are ten times less likely to be written up and published than trials with a significant result. We estimate that if there were no publication bias, the average effect of a nudge in the Academic Journals sample would be 3.2 pp. (s.e.=1.9). While this estimate is much smaller than the uncorrected estimate of 8.7 pp., it is consistent with the Academic Journals nudges having a somewhat larger effect size than the Nudge Unit interventions (though the difference is not statistically significant). We use these estimates to produce simulated counterfactuals of the average effect size of a nudge if one were to alter each of the three dimensions separately and in combination. Selective publication has the largest impact on the point estimate, but statistical power also plays a role, in that underpowered studies exacerbate the impact of selective publication.

Overall, selective publication and (to a smaller extent) differences in features of nudge interventions account for the gap in effect size between the Nudge Units and Academic Journals. These results also suggest that the 1.4pp. estimate for the Nudge Unit trials is

a reasonable estimate for the average impact of a nudge at scale on government services. While a cost-benefit analysis is not the focus of this paper (see Benartzi et al., 2017), we stress that this impact comes with a marginal cost that is typically zero or close to zero, thus suggesting a sizable return on investment.

This paper is related to a growing literature on effectiveness of nudges (e.g., Laibson, 2020; Milkman et al., 2020). We contribute what, to our knowledge, is the first comprehensive estimate of the effect of RCTs from a Nudge Unit. The 1.4 pp. estimate likely is a lower bound of the impact of behavioral science for three reasons. First, the Nudge Units face institutional constraints which, for example, largely rule out default changes, that tend to have larger impacts (Jachimowicz et al., 2019). Second, the trials we consider typically have multiple arms; while we estimate the average impact of each arm, organizations can adopt the most successful nudge in the trial. Third, researchers can build on the most successful results in the design of later interventions.

This paper is also related to the literature on publication bias (e.g., Simonsohn, Nelson, and Simmons, 2014; Brodeur et al., 2016) and research transparency (Miguel et al., 2014; Christensen and Miguel, 2018). We show encouraging evidence of best-practice transparency in Nudge Units, which ran appropriately powered trials and kept track of all the results, thus enabling a comprehensive evaluation. In comparison, we document a large role of selective publication in published papers. We also apply the publication-bias correction of Andrews and Kasy (2019) and show that the normality assumption traditionally used in meta-analyses is too restrictive and would lead to biased estimates.

Bringing these two literatures together, a key question is the extent to which selective publication leads to bias in the estimate of the impact of behavioral science. On the one hand, it leads to the publication of results with large effect sizes due to luck or p-hacking, especially given the many under-powered interventions. These results are unlikely to replicate at the same effect size, thus inducing bias. Indeed, replications (in other settings) typically yield smaller point estimates than the original results, e.g., for laboratory experiments (Camerer et al., 2016) or TV advertising impacts (Shapiro, Hitsch, and Tuchman, 2020). On the other hand, selective publication may also highlight the interventions that turn out to be truly successful at inducing a behavior; these “good ideas” would presumably replicate. Our results cannot measure the magnitude of the two forces, given that the Nudge Unit interventions are not exact replications of the Academic Journal nudges.

Finally, the paper is related to the literature on scaling RCT evidence (Banerjee and Duflo, 2009; Allcott, 2015; Bold et al., 2018; Dehejia, Pop-Eleches, and Samii, 2019; Meager, 2019a; Vivald, 2020). In our case, “scaling” nudges did not entail the examination of, for example, general-equilibrium effects (e.g., Muralidharan and Niehaus, 2017) which are important in other contexts. Rather, the key aspects of scaling in our setting are the ability to conduct adequately powered interventions, within the institutional constraints that are more likely to arise at scale. The latter aspect echoes the findings on scaling in Bold et al. (2018) and Vivald (2020).

2 Setting and Data

2.1 Trials by Nudge Units

Nudge Units. We analyze the trials conducted by two large “Nudge Units” operating in the US: the Office of Evaluation Sciences (OES), which collaborates with federal government agencies; and the Behavioral Insights Team’s North America office (BIT-NA), which worked primarily with local government agencies during the period in question.

OES was launched in 2015 under the Obama Administration as the core of the White House Social and Behavioral Sciences Team (SBST). The formal launch was coupled with a Presidential Executive Order in 2015, which directed all government agencies to “develop strategies for applying behavioral science insights to programs and, where possible, rigorously test and evaluate the impact of these insights.” OES staff work with federal agencies to scope, design, implement, and test a behavioral intervention. Also in 2015, the UK-based Behavioural Insights Team (BIT) opened its North American office (BIT-NA), aimed at supporting local governments to use behavioral science. Mainly through the What Works Cities initiative, BIT-NA has collaborated with over 50 U.S. cities to implement behavioral experiments within local government agencies.

The two units have shared goals: to use behavioral science to improve the delivery of government services through rigorous RCTs, and to build the capacity of government agencies to use RCTs. In interviews with the leadership of both units, both teams noted that their primary goal is to measure what changes will make a measurable difference on key policy outcomes. The vast majority of their projects are similar in scope and methodology. They are almost exclusively RCTs, with randomization at the individual

level; they often involve a low-cost nudge using a mode of communication that does not require in-person interaction (such as a letter or email); and they aim to either increase or reduce a behavioral variable, such as increasing take-up of a vaccine, or reducing missed appointments. Furthermore, the two units embrace practices of good trial design and research transparency. All trial protocols, including power calculations, and results are documented in internal registries irrespective of the results. All data analyses go through multiple rounds of code review. Moreover, OES has taken the additional step of making all trial results public, and recently, posting pre-analysis plans for every project.

These units are central to the process of taking nudge RCTs to scale in a meaningful way. In this case, scaling means two things. First, “scaling” occurs in the numerical sense, because government agencies often have access to larger samples than the typical academic study, and so the process of scaling nudge interventions tells us how an intervention fares when the sample is an order of magnitude larger than the original academic trial. Second, the selection of trials that Nudge Units conduct also tells us something about which academic interventions are politically, socially, and financially feasible for a government agency to implement—“scalable” in the practical sense.

Figure 1a-b shows an intervention from OES aimed to increase service-member savings plan re-enrollment. The control group received the status-quo email (Figure 1a), while the treatment group received a simplified, personalized reminder email with loss framing and clear action steps (Figure 1b). The outcome is measured as the rate of savings plan re-enrollment. Figure A2a presents two additional examples of OES interventions, focused respectively on increasing vaccine uptake among veterans and improving employment services for UI claimants in Oregon. Figure A2b presents an example of a nudge intervention run by BIT-NA encouraging utilities customers in the treatment group to enroll in AutoPay and e-bill using inserts in their utility bill, while the control group received the status quo utility bill. The outcome is measured as the rate of enrollment in either AutoPay or e-bills.

Sample of Trials. As Figure 2a shows, from the universe of 165 trials conducted by the units, we limit our sample to projects with a randomized controlled trial in the field, removing 13 trials. We then remove 15 trials without a clear “control” group, such as horse races between two behaviorally-informed interventions. We then remove 3 trials with monetary incentives, and limit the scope further to trials with a primary outcome that is binary, removing 8 trials. We also remove trials where the “treatment” is

changing the default, since they are the rare exception among Nudge Unit interventions in our sample (only two treatment arms of one trial), and substantively different.

Our final sample consists of 126 randomized trials that include 243 nudges and involve 23.5 million participants. For each trial, we observe the sample size in the control and treatment groups and the take-up of the outcome variable in each group, e.g., the vaccination rate or enrollment in a savings plan. Whenever there are multiple dependent variables specified in the pre-analysis or trial report, we take the primary binary variable specified. We do not observe the individual-level micro data though, arguably, given the 0-1 dependent variable this does not lead to much loss of information. We discuss additional details in Online Appendix A.1. To our knowledge, only 16 of these trials (listed in Table A1a) have been written or published as academic papers; we discuss the results for this subset in Online Appendix A.2.

2.2 Trials in Academic Journals

Sample of Trials. We aim to find broadly comparable published nudge studies, without hand-picking individual papers. In a recent meta-analysis, Hummel and Maedche (2019) select 100 papers screened out of over 2,000 initial papers identified as having “nudge” or “nudging” in the title, abstract, or keyword. We report their selection criteria in the Online Appendix A.3 The papers cover a number of disciplinary fields, including economics, public health, decision-making, and marketing. A second meta-analysis that covers several areas is Benartzi et al. (2017), which does a cost-benefit comparison of a few behavioral interventions to traditional incentive-based interventions. Hummel and Maedche (2019) review 9 other meta-analyses, which we do not include because they focus exclusively on one policy area or topic. We thus combine the behavioral trials in Hummel and Maedche (2019) and in Benartzi et al. (2017), for a total of 102 trials.¹

We apply parallel restrictions as for the Nudge Units sample, as Figure 2b shows. First, we exclude lab experiments, survey experiments with hypothetical choices, and non-RCTs (e.g., changes in a cafeteria menu over time, with no randomization), for a remaining total of 52 studies. Second, we exclude treatments with financial incentives, removing 3 trials. Third, we require binary dependent variables, dropping 21 trials.

¹This sample omits some influential published nudge RCTs, such as Bhargava and Manoli (2015) and Hallsworth et al. (2017). We did not add any such papers to avoid subjective paper additions.

Finally, we exclude default interventions, dropping just 2 trials. This leaves a final sample of 26 RCTs, listed in Table A1b, including 74 nudge treatments with 505,337 participants. For each paper, we code the sample sizes and the outcomes in the control and the nudge treatment groups, as well as features of the interventions.

2.3 Comparison of Two Samples and Author Survey

Categories of Nudges. We categorize each nudge by its policy area, communication channel, and behavioral mechanisms, as highlighted in Tables 1 and A2, and Figure A3.

The policy areas are best described through examples. A typical “revenue & debt” trial involves nudging people to pay fines after being delinquent on a utility payment, while an example of a “benefits & programs” trial encourages individuals to enroll in a government program, such as pre- and post-natal care for Medicaid-eligible mothers. The “workforce and education” category includes prompting job-seekers to plan their job search strategy. One “health” intervention urges people to get vaccinated or sign up for a doctor’s appointment. A “registration” nudge asks business owners to register their business online as opposed to in person, and a “community engagement” intervention motivates community members to attend a local town hall meeting. Compared to the Nudge Units sample, the Academic Journals sample has a larger share of trials about health outcomes and environmental choices and fewer about revenue and debt, benefits, and workforce and education.

The medium of communication with the treatment group tends to be through email, letter, or postcard in the Nudge Units sample, as opposed to in-person interactions, which are common in the Academic Journals sample. In 61% of the Nudge Unit trials and in 43% of the Academic Journal trials, the researchers do not send the control group any communication within the field experiment (although the control group may still be receiving communication about the specific program through other means).

We also code the behavioral mechanisms, with details in Online Appendix A.4. In the Nudge Units sample, the most frequent mechanisms include: simplification of a letter or notice; drawing on personal motivation such as personalizing the communication or using loss aversion to motivate action; using implementation intentions or planning prompts; exploiting social cues or building social norms into the communication; adjusting framing or formatting of existing communication; and nudging people towards an active choice

or making some choices more salient. In the Academic Journals sample, there are fewer cases that explicitly feature simplification and information as one of the main levers, and more cases that feature personal motivation and social cues, changes in framing and formatting, or choice re-design (e.g., active choice).

Features of Trials. While Table 1 categorizes the nudges on several descriptive dimensions, Table 2 summarizes broader features of the trial, such as the degree of academic involvement, the difficulty of changing the selected behavior, trial design decisions, and features of implementation. We draw on a combination of information from the papers and trial reports, as well as from a short survey of authors and of the Nudge Unit leadership teams. We present details of the survey in the Online Appendix A.5.

The first feature is the degree of academic involvement. While all studies in the Academic Journals sample (Column 1) are led by academics, there is significant heterogeneity in the Nudge Units sample (Column 2). BIT North America employs behavioral scientists and other researchers directly, and so BIT-NA trials (Column 3) are designed by internal staff in collaboration with government partners. In comparison, OES interventions (Column 4) are often designed in coordination with academic fellows, PhD students, academics on sabbaticals at OES, and university faculty who collaborate on individual trials. The 50% of OES trials that record a full-time faculty member as lead or affiliate (Column 5) are therefore more similar in this respect to the Academic Journal trials, and we will consider them separately.

Second, we consider the difficulty of moving a behavioral outcome. Differences in treatment effects could arise if trials in either the Nudge Units or the Academic Journals sample tackle behaviors that are harder to shift. While we do not observe this directly (e.g., through a measure of elasticity), we lean on two proxies. A first proxy is the control group take-up, given that outcomes with very low take up may be especially hard to shift. The average control group take-up is 26.0% in the Academic Journals sample, and 17.2% in the Nudge Units sample, a difference that is only marginally statistically significant, suggesting that the two samples are reasonably comparable in this dimension. As a second proxy, we measure the time horizon of the outcome variable, i.e., the number of days between the receipt of the intervention and the recorded behavior. For example, if the outcome is whether the recipient clicks a link in the email on the day it is sent, we record a 1-day time frame, and if the outcome is re-enrollment in college six months after the receipt of a letter, we record 180 days. Short-run responses are presumably

easier to affect with a nudge. The OES interventions have a longer time frame than the Academic Journals trials, though the difference is not statistically significant

Third, we consider differences in trial design. Trials in the Academic Journals sample have fewer behavioral mechanisms per treatment arm: an average of 1.5, compared to an average of 2.2 in Nudge Unit trials ($p < 0.01$). At the same time, the average trial in the Academic Journals sample evaluates more treatment arms, 2.8 versus 1.9. The typical treatment arm in the Academic Journals sample is also less statistically powered with a larger MDE, a point we revisit below. Trial design may also be affected by institutional constraints, where the implemented design may differ from the ideal design initially envisioned. We therefore asked survey respondents to indicate on a Likert scale from 1 (vastly different) to 5 (exactly the same) how close the final intervention was to the ideal one they had hoped to implement. We find a clear difference: while most Academic Journal RCTs have a rating of 4 or 5, the BIT or OES interventions are typically rated 3, indicating a stronger impact of institutional constraints.

A fourth trial feature is decision-making around trial planning and implementation. The Academic Journal interventions may involve a more extensive planning and design process, which may impact the effect size. We thus asked authors of the papers in the Academic Journals sample to indicate the approximate number of months of total duration of the RCT, as well as the months of planning, of intervention and data collection, and of data analysis and write-up. We also asked for the full-time staff or PI months spent on a project. We asked parallel questions to the BIT and OES staff, and we contacted the academic fellows for the Academic-Affiliated OES trials. As Table 2 shows, the answers are closer than one may have thought. The average duration of the planning and intervention periods is similar for the Academic Journals sample and the OES sample (11-13 months), and somewhat shorter for the BIT sample (around 7 months). The average personnel time is higher for the Academic Journals sample than for the Nudge Units sample (14.9 months vs. 5.8 months), but this difference is amplified by a couple outliers in the Academic Journals sample: the difference in the medians is quite modest (9 vs. 6 months). The data analysis and write-up period is shorter for the Nudge Unit interventions, given that most are not written up as academic papers.

Overall, we identified a few differences in trial features between the Academic Journals sample and the Nudge Units sample. The Nudge Unit interventions are less likely to be led by academics, tend to face higher institutional constraints, and involve fewer

personnel. These features (or at least the first two) seem typical of an “at scale” intervention. Nudge Unit trials also include fewer arms per trial, a larger sample and more behavioral mechanisms per arm. This may suggest a different objective function, where more emphasis is placed on moving the policy outcome and less on disentangling the exact mechanism. We return to these differences in Section 4.2.

3 Impact of Nudges

We present the unweighted impact of the nudges in the two samples.

Academic Journals. As Column 1 in Table 3 shows, the average treatment effect for the 74 nudges in 26 trials in the Academic Journals sample is 8.68 pp. (s.e.=2.47), a large increase relative to the average control group take-up of 25.97 percent.

Figure 3a shows the estimated nudge-by-nudge treatment effect together with 95% confidence intervals, plotted against the take-up in the control group. The figure shows that there is substantial heterogeneity in the estimated impact, but nearly all the estimated effects are positive, with some very large point estimates. The plot also shows suggestive evidence that the treatment effect seems to be highest in settings in which the control take-up is in the 20%-60% range.

Nudge Units. Column 2 in Table 3 shows the unweighted average impact of the 243 nudge treatments in the 126 Nudge Unit trials. The estimated impact is 1.38 pp. (s.e.=0.30), compared to an average control take-up of 17.20 pp.. This estimated treatment effect is still sizable and precisely estimated to be different from zero, but is one sixth the size of the point estimate in Column 1 for the Academic Journal papers.

Figure 3b shows the estimated treatment effect plotted against the control group take up. The treatment effects are mostly concentrated between -2pp. and +8pp., with a couple of outliers, both positive and negative. Among the positive outliers are treatments with reminders for a sewer bill payment and emails prompting online Auto Pay registration for city bills. One trial that produced a negative effect is a redesign of a website aimed to encourage applications to a city board.

The comparison between Figures 3a and 3b, which are set on the same x - and y -axis scale, visually demonstrates two key differences between published academic papers and Nudge Unit interventions. The first, which we already stressed, is the difference in estimated treatment effects, which are generally larger, and more dispersed, in the

Academic Journals sample. But another equally striking difference is the statistical precision of the estimates: the confidence intervals are much tighter for the Nudge Unit studies, which are typically run with a much larger sample.

Robustness. Tables A3a-b and A4a-b display additional information on the treatment effects. As Table A3a shows, the difference in treatment effects between the two samples is parallel in log odds terms (which can be approximately interpreted as percent effects): 0.50 log points (s.e.=0.11) for the Academic Journals sample, compared to 0.27 log points (s.e.=0.07) in the Nudge Units sample. The impact in log odds point is larger than the impact that one would have computed in percent terms from Table 3, given that treatment effects with lower control take-up are larger in log odds. Table A3b shows that the average treatment effect for the Nudge Unit trials that have been published in academic journals (see Online Section Appendix A.2 for details) is 0.97 pp. (s.e.=0.23), similar to the entire Nudge Units sample, albeit slightly smaller. Table A4a displays the number of treatments that are statistically significant, split by the sign of the effects.

Table A4b shows that the estimates in both samples are slightly larger if we include the default interventions, with the caveat that these are just 3 arms in the Academic Journals sample and 2 arms in the Nudge Units sample. Next, while we cannot fully capture the “importance” of the outcome variable in each nudge, in Table A4b we consider the subset of nudges with “high-priority” outcomes, as rated by a team of undergraduates, which aim to capture variables closer to the policy of interest (for example, measuring actual vaccination rates as opposed to appointments for a vaccination). The estimated nudge impact for this subset is somewhat lower for the Academic Journal nudges at 6.5 pp., but at least as high for the Nudge Unit ones, at 1.6 pp.. We then consider the subset of Nudge Unit interventions that are low-cost, that is, either relying on email contact or on existing means of communication with the control group. We replicate the same effect size. Finally, estimates weighted by citations for the Academic Journals sample yield slightly lower point estimates.

4 Nudge Units Vs. Academic Journal Nudges

We sketch a model of decision-making around nudge experimentation, highlighting features of experimental design (as in, e.g., Frankel and Kasy, 2020 and Azevedo et al., forthcoming), as well as selective publication (as in Andrews and Kasy, 2019).

We assume that researchers and nudge units alike design experiments aiming to provide evidence on a particular treatment, in our case a nudge. Our model makes three key sets of assumptions, capturing the trial design, underlying effect size, and selective publication. First, both academic researchers and Nudge Units design an experiment to detect an effect size d with 0.80 statistical power. Second, there is a true effect size of the nudge intervention β distributed with a random effect. Third, results that are not statistically significant are published by academic researchers with some probability $\gamma < 1$, while results that are statistically significant are published with probability 1.

We propose that the observed differences between trials in our two samples can be largely explained by differences in each of those three components, d , β and γ . First, the samples differ in the expected effect size d , with $d_{NU} < d_{AJ}$, resulting in differences in statistical power and number of treatment arms. Second, they differ in the average effectiveness of the nudges, with $\beta_{NU} < \beta_{AJ}$. Third, they differ in selective publication, that is $\gamma_{AJ} < 1$, but not for (this sample of) Nudge Units. We discuss these differences in turn below, as well as a few possible alternative explanations of the findings.

4.1 Experimental Design

Models of optimal experimental design typically center on the goal of collecting evidence on the effectiveness of an intervention. Frankel and Kasy (2020), for example, model a researcher that decides the optimal sample size for a treatment (and whether to run an experiment) as a function of the prior, the cost of collecting evidence, and other factors. Azevedo et al. (forthcoming) discuss the trade-off for a given sample between running fewer treatment arms with a larger sample, or more treatment arms with less power per arm, as a function of the fatness of the tails in the distribution of treatment effects.

For simplicity, we assume a simple, rule-of-thumb experimentation rule, based on the Cohen (1965) convention, that researchers aim for statistical power of 0.80, given an expected effect size d for a treatment arm. We recognize that this is a descriptive model and we do *not* see it as normative, given the need to take into account the cost of collecting observations, the priors, etc. Still, aiming for 0.80 power is widespread. We collected the power calculations for all pre-registrations on the AEA Registry, the largest data set we know with systematic power calculations. Among the 267 that were registered prior to the start of their intervention and provided a minimum detectable

effect with a targeted level of power, 240 use a power target of 0.80 (Table A5).

In our setting, given the binary dependent variable, the implicit MDE d for 80 percent power can be computed using just the control take-up and the sample sizes in the control and treatment groups, all of which we observe.² As Figure 4a shows, the Academic Journals sample has a median MDE d_{AJ} of 6.30 pp., and an average MDE of 8.18 pp.; most of these studies are powered to only detect really quite large treatment effects. In contrast, the Nudge Units sample has a median MDE d_{NU} of 0.78 pp. and an average MDE of 1.72 pp. Thus, the statistical power is an order of magnitude larger in the Nudge Units sample than in the Academic Journals sample. Figure A4a shows the corresponding difference in sample size per treatment arm: a median of 484 in the Academic Journals sample versus 10,006 in the Nudge Units sample.

Is it plausible that the Nudge Units were expecting effect sizes of, on average, around 1 percentage point, while academics were expecting effect sizes closer to 8 pp.? While we did not ask this question exactly, we surveyed academic researchers and nudge practitioners about their expectation for the findings of our study, as in DellaVigna and Pope (2018) and along lines outlined by DellaVigna, Pope, and Vivalt (2019). Specifically, we created a 10-minute survey eliciting forecasts using a convenience sample through email lists and Twitter ($n=237$). The survey explained the methodology of our analysis, described the two samples, showed participants three nudge interventions randomly drawn out of 14 exemplars, and asked for predictions (in percentage point units) of: (a) the average effect size for the Nudge Units sample; and (b) the average effect size for the Academic Journals sample. Among the respondents, 28 self-identify as Nudge Unit practitioners, and 66 as academic researchers. We focus on the predictions of these two samples, with more detail in Online Appendix A.6.

In Figure 4b the blue continuous line indicates the distribution of forecasts by (self-identified) nudge practitioners about the effect size finding for the Nudge Unit interventions. The median nudge practitioner expects an average impact of 1.95 pp., which is remarkably in line not only with our findings, but also with the calculation of the projected effect size d_{NU} implied by the observed sample size.

Do researchers expect a larger effect size? Figure 4b shows two predictions by academics: for the Nudge Units sample, and for the Academic Journals sample. The

²As far as we can tell, none of the papers in the Academic Journals sample were pre-registered, so we do the power calculation ourselves.

researchers significantly overestimate the findings for the Nudge Units sample: the median academic expects an effect size of 4.0 pp. As for the Academic Journals sample, the median academic expects an effect size of 7.0 pp., close to the observed average effect size in this sample and also close to the effect size d_{AJ} implied by the power calculations.

Put differently, when asked about the same sample, with the same information, academic researchers expect a larger effect size than nudge practitioners. We take these results as suggestive of the fact that academics' expectations about the effect size of their own trial may shape their trial design, and can at least partially explain differences in statistical power between the two samples.³ We acknowledge that, while we attribute this difference in MDE between the two samples to different expected effect sizes, it may also be due to different views of what is a policy-relevant or publishable effect size.

A reasonable objection is that sample sizes may, to a large extent, be fixed, leaving experimenters with little choice (other than deciding whether to run the experiment altogether). One decision, though, that they clearly have is how many treatment arms to run: the more arms, the lower the statistical power per arm. If researchers actively aim for a higher MDE d (at the cost of lower statistical power), we would observe it in this dimension. As Figure A4b and Table 2 show, the number of arms is significantly *larger* in the Academic Journals sample than in the Nudge Units sample, despite academic researchers having (on average) a smaller sample size to start with.

Figure A4c and Table 2 document an additional trial design feature. Academic researchers have fewer behavioral mechanisms per arm (as coded in our categorization). Plausibly, they are more concerned with establishing a mechanism for the potential findings, thus requiring multiple arms, each with a different mechanism. Nudge Units, instead, may be mostly focused on studying combinations of mechanisms to yield higher policy impact. Nudge Unit leadership confirmed that this was their primary concern when designing trials, and as one interviewee noted, "You only get to add more treatment arms if you can show with certainty that you are powered well enough to detect an effect between one treatment and control."

Overall, we find a major difference in trial design: much larger sample sizes per treatment arm in the Nudge Units sample. We propose a simple explanation: academic

³In part, this difference depends on experience running trials: 54% of nudge practitioners have conducted 5 or more field experiments, versus only 18% of academic researchers. Among the academic researchers, those who have run 5 or more trials predict an average effect size of 6.4 pp. for the Academic Journals sample compared to a prediction of 8.4 pp. among academics who have run 1 or fewer trials.

researchers expect or target larger effect sizes and therefore are comfortable with a larger MDE. While we cannot prove this directly, we document that this difference in sample size is in line with differences in forecasts of effect sizes in survey predictions. Below, in Section 4.4, we consider the implications of these differences in statistical power.

4.2 Differences in Nudge and Trial Features

A second potential difference between the two samples is in the average effect size β . We now examine a number of observable features of the trials and of the nudges that may explain why the two samples could have a different β .

4.2.1 Academic Involvement

As we documented in Table 2, while BIT trials are typically designed by internal staff, the OES interventions are typically designed in collaboration with academic fellows. The two sets of trials also differ in other dimensions: the OES trials have a longer planning and intervention duration and higher personnel FTE involvement.

Thus, in Table 3 we revisit the effect size separately for the two Nudge Units. The average effect size for BIT interventions (1.67 pp, s.e.=0.52, Column 3) is similar to, and in fact slightly larger than, the effect size for the OES interventions (1.02 pp, s.e.=0.21, Column 4). Further for the 24 OES trials with explicit academic involvement, the point estimate is essentially the same as for the overall OES sample (0.98 pp., s.e.=0.41, Column 5). Thus, differences in academic involvement and in the set up of the two Nudge Units per se do not appear to explain our findings.

4.2.2 Categories of Nudges

Next, we separate the impact of nudge treatments by category. As Table 1 shows, the average effect size (ATE) varies substantially across interventions: for example, in-person interventions or nudges on the environment policy area have larger effect sizes. Both types of interventions are more common in the Academic Journals sample, and could thus contribute to the different effect sizes.

In Column 2 of Table 4 we thus include in the effect size specification the nudge features controls in Table 1, as well two additional variables: a quadratic of the average take-up in the control group, which could proxy for the difficulty in affecting a behavior

(e.g., the persuasion rate), and the outcome time frame, which could capture harder-to-affect longer-run outcomes.⁴ The point estimate is larger for studies focused on the environment, for cases with no previous communication, and for cases in which the contact takes place in person, as opposed to via email or mail; also simplification and especially choice design appear to have the largest effects.⁵

Importantly, these controls reduce the difference in point estimate between the samples from 7.3 pp. (Column 1) to 2.4 pp. (Column 2). Thus, while these controls do not fully explain the difference, they bridge two thirds of the gap.⁶ In Column 3 of Table A7 we present an alternative procedure to account for these features, reweighting according to a propensity score. Reweighting does not affect the Nudge Unit point estimate much, but it sizably reduces the Academic Journals estimate, thus shrinking the gap by half.

4.2.3 Features of Trials

While so far we controlled for the type of nudge, next we control for general features of the trial described in Table 2. In Columns 5-7 of Table 4 we hold academic involvement constant and consider only the Nudge Unit trials with an academic affiliate, as well as the Academic Journal trials. In this subsample we replicate the large difference in effect size, with a 7.7 pp. larger effect size in the Academic Journals sample (Column 5).

Adding controls for the features of trials in Column 6, we find that the Likert rating for how close the intervention was to the planned one has a positive impact, but is not significant. The measure of personnel involvement also has a positive, but not statistically significant, impact. Altogether, these features have only modest explanatory power for the effect size difference between the two samples.

Overall, we can explain some of the gap between our two samples as a difference in β , where β is determined by nudge features including the policy outcome, the behavioral mechanism used, and the mode of communication. We recognize that these nudge

⁴We exclude the indicators of early vs. late years, which are not comparable across the samples.

⁵We can compare these findings to the ones in Hummel and Maedche (2019). While the categories differ from our coding, a commonality is that the average nudge in the “Environment” policy area is highly effective. As for the behavioral mechanisms, Hummel and Maedche (2019) find nudges in the “Default” category, which often overlaps with “Choice design” in our scheme, to be highly effective, which aligns with our findings. We caution though against a causal interpretation of these heterogeneity results. The differences in trial characteristics and in treatment effects may reflect feasibility constraints.

⁶In Table A6a-b we present similar regressions run separately for the Academic Journals sample and the Nudge Units sample. The nudge features have somewhat similar estimates.

features are not exogenously selected and that at least some of these differences in observables, such as lack of in-person interventions, may be part of going to scale. We return to differences in effect size below when modeling jointly with selective publication.

4.3 Selective Publication

The third component of our model is selective publication. Following the literature (e.g., Andrews and Kasy, 2019), we include any channel leading to selective publication out of the sample of all studies run, including decisions by journal editors on which papers to publish, but also by researchers of which studies to write up and submit for publication (the file drawer effect). We expect some publication bias in the Academic Journals sample, but not in the Nudge Units sample where we access all trials.

4.3.1 Graphical Evidence on Publication Bias

As a first test, following Card and Krueger (1995), in Figure 5a we plot each point estimate for the nudges in the Academic Journals sample as a function of the statistical precision of the estimate, in our case measured with the statistical power (MDE). The plot shows evidence of two phenomena. First, there is a fanning out of the estimates: the less-powered studies (studies with larger MDE) have a larger variance of the point estimates, just as one would expect even without any selective publication. Second, the less-powered studies also have a larger point estimate for the nudge. Indeed, a simple linear regression estimate displayed on the figure documents a strong positive relationship: $y = 0.116(s.e. = 1.935) + 1.047(s.e. = 0.303)MDE$. This second pattern is consistent with publication bias: to the extent that only statistically significant results are published, less imprecise studies will lead to a (biased) inference of larger treatment effects. We observe similar patterns when we plot the treatment effect against the standard error, another measure of precision, as shown in Figure A5.

As a second test, following Brodeur et al. (2016), in Figure 5b we plot the distribution of t statistics around the standard 5% significant threshold ($t=1.96$) for the nudge treatments in the Academic Journals sample. We detect no bunching in t statistics to the right of the $t=1.96$ threshold. Behavioral studies, however, often employ multiple treatment arms in one trial, compared to a control group, often in a horse race of alternative behavioral levers. In such a setting, arguably, for publication what matters is

that at least *one* nudge or treatment arm be statistically significant, not all of them.

In Figure 5c, thus, we plot the distribution of the most significant t -statistic across the different nudge treatments in a trial. There are 9 papers with a (max) t statistic between 1.96 and 2.96, but only 2 papers with (max) t statistic between 0.96 and 1.96. This suggests that the probability of publication for papers with no statistically significant results is only a fraction of the probability of publication for studies with at least one significant result.⁷ Zooming in closer around the threshold, there is only 1 study with a max t statistic between 1.46 and 1.96, versus 6 between 1.96 and 2.46.

In Figure 6 we produce the same plots for the Nudge Unit trials. The contrast of Figure 6a with Figure 5a is striking: in the Nudge Units sample there is no evidence that the less-powered studies have larger point estimates. Indeed, a linear regression of point estimate on MDE returns $y = 1.012(s.e. = 0.339) + 0.210(s.e. = 0.246)MDE$, providing no evidence of a positive slope. Further, Figures 6b and 6c show there is no discontinuity in the distribution of the t statistic, nor in the max of the t -statistic by trial. This is consistent with the fact that we observe the universe of completed trials.

As a further piece of evidence, in Figure A5c-f we present funnel plots as in Andrews and Kasy (2019), plotting the point estimate and the standard errors, with bars indicating the results that are statistically significant. These plots indicate an apparent missing mass for the Academic Journal papers when considering the max t statistics (Figure A5d), and no evidence of a missing mass for the Nudge Unit trials (Figure A5f).

This evidence thus points to selective publication in the nudge experiments run by academic researchers. This evidence adds to the publication bias literature in two ways. First, it suggests that the maximal t statistic may play an even larger role in determining publication than all individual t statistics. Second, our result appears to differ from the findings of Brodeur, Cook, and Heyes (2020) who do not find statistically significant evidence of p -value manipulation for experimental studies (as opposed to in difference-in-differences or instrumental variable studies) using the universe of papers in top-25 economics journals in 2015 and 2018. In Figure A6, we re-consider the data in Brodeur, Cook, and Heyes (2020) comparing the evidence of manipulation when we consider each t -stat on its own, as opposed to the max t -stat in a paper. We focus on t statistics from

⁷A binomial test indicates a probability of 9 or more significant results out of 11 (assuming a null of 0.5) of $p = 0.0327$. This may even understate the extent of publication bias. Among the three Academic Journals trials with statistically insignificant results (see Table A1b), two actually emphasize statistically significant results, either on a subsample or on a different outcome.

the main table in each paper. Figure A6a replicates the finding of bunching around $t=1.65$ ($p=0.10$) and $t=1.96$ ($p=0.05$) for individual t statistics. When considering the maximal t -statistic, Figure A6b shows a sizable jump in the distribution around $t=1.65$ ($p=0.10$): there are 10 studies with maximal t -stat just above this threshold, but only 2 just below it. There is, however, no obvious jump at $t=1.96$. When restricting the sample to only experimental studies, the evidence is suggestive given the much smaller sample, but still there is an apparent gap in the distribution of the maximal t -statistic below $t=1.65$, as opposed to above, qualitatively consistent with the findings in our sample of published nudge interventions (Figures A6c-d).⁸

4.3.2 Reduced-form Evidence on Impact of Publication Bias

Before we turn to our full model of selective publication, we consider reduced-form evidence in the spirit of Egger’s test. Returning to Table 4, in Column 3 we control for statistical power (MDE) in both the Nudge Units sample and in the Academic Journals sample. The idea of this test is to obtain the predicted effect size for experiments with a very large sample size (and thus no role for sampling error or publication bias). The nudge effect size is strongly increasing with the MDE in the Academic Journals sample, but not in the Nudge Units sample, consistent with the pattern in Figure 5a and 6a. Adding these controls can explain the *entire* difference in effect size: for trials with, hypothetically, zero MDE, the effect size is indistinguishable in the two samples, and is 1 percentage point in the Nudge Units sample. This implies that publication bias can explain the entire difference in point estimates. We replicate this result in Table A7 with alternative specifications for this test, in Column 1 using an exact Egger’s test, with standard errors as regressors, and inverse-variance weights, and in Column 2 instead of regression controls, we present results weighted by $1/\text{MDE}$.

4.4 Meta-Analysis Model with Publication Bias Correction

Bringing these components together, we now turn to our meta-analytic model of treatment effect sizes with publication bias based on Andrews and Kasy (2019). The model allows for all three key dimensions: (i) (*statistical power*) the model takes as input the

⁸We are very grateful to Abel Brodeur for promptly sharing the data for this analysis. Figure A5g-j reproduces the evidence from the nudge experiments with the smaller bin size employed in Figure A6.

precision of the estimates implied by the differences in statistical power; (ii) (*effect size*) the model allows for different effect sizes β for academic researchers and Nudge Unit interventions; (iii) (*publication bias*) papers with no statistically significant results in the Academic Journals sample have a probability of publication $\gamma_{AJ} \leq 1$.

The Andrews and Kasy (2019) model builds on traditional random-effects meta-analysis models, and adds selective publication. In a meta-analysis, the researcher collects a sample of studies (indexed by i), each with an observed effect size $\hat{\beta}_i$ that estimates the study’s true effect size β_i , and an observed standard error $\hat{\sigma}_i$. A *random-effects model* allows the true effect β_i to vary around the grand true average effect $\bar{\beta}$ with some variance τ^2 . The parameter τ may represent differences in context, populations, design features, etc. In our setting, there are multiple treatment arms in nearly each study. Thus, we introduce a *within-trial* random effect variance. This allows for different nudges within the same trial (i.e. study) to have more similar results than nudges across different studies, since they share a setting, a behavioral outcome, and basic design. Formally, the trial-level base effect β_i is drawn from $N(\bar{\beta}, \tau_{BT}^2)$, and the treatment-level true effect β_{ij} is drawn from $N(\beta_i, \tau_{WI}^2)$.

To start with, in Panel A of Table 5, we present maximum likelihood estimates from such traditional meta-analysis (other than allowing for a separate within-trial variance).⁹ The estimated effect sizes $\bar{\beta}$ are very close to the unweighted estimates in Table 3, at 8.58 pp. for the Academic Journals sample and 1.49 pp. for the Nudge Units sample.

Figures 7a shows the distribution of treatment effect for the Academic Journals sample and the fit of this model (blue dotted line). This normal-based estimator provides a poor fit, given the nearly bi-modal distribution of the underlying data: most treatment effects are in the range between 0 and 10 pp., but there is also a right tail above 10 pp., with no corresponding left tail. The substantial right skew, which a normal distribution cannot fit, leads to an upward bias in the estimate for $\bar{\beta}$ and a very large estimate for τ^2 ; this implies that the meta-analysis estimate is very close to the unweighted average.

Figure 7b displays the distribution of treatment effects for the Nudge Unit trials. Once again, this normal-based model (blue dotted line) does not fit the data well: there are more effect sizes in the right tail than under the estimated distribution.

⁹Table A8 presents alternative traditional meta-analysis estimators explained in Online Appendix A.8. Some of these estimators shrink the effect size for the Academic Journals sample sizably, and especially the fixed-effect estimator. The estimates for the Nudge Unit trials vary in a more limited range between 0.9 and 1.4 pp.

We extend this meta-analysis method in two dimensions, further discussed in Online Appendix A.7. First, recognizing the skewness of treatment effects in Figures 7a-b, we allow for the trial-level effects to be drawn from a mixture of two normals, each with its own between- and within-trial variance.¹⁰ Second, for the Academic Journals sample, we allow for publication bias as in Andrews and Kasy (2019). We assume that studies with no significant results are γ_{AJ} times as likely to be published as studies with a significant intervention. Selective publication in favor of significant results would imply that γ_{AJ} is less than 1. We assume that publication bias occurs at the level of the most significant nudge j within a paper i , consistent with the evidence from Figures 5b-c, that is,

$$Pr(\text{Publish}_i) = \begin{cases} 1 & \text{if } \max_j(\hat{\beta}_{ij}/\sigma_{ij}) \geq 1.96 \\ \gamma_{AJ} & \text{otherwise} \end{cases}$$

The estimates for the Nudge Units sample, in Panel B, have a drastically improved log likelihood. We estimate that the treatment effects comes from two distributions, one centered at 0.34 pp., a second one centered at 5.10 pp., with 78% of trials drawing from the first distribution. The overall estimated treatment effect, at 1.38 pp. (the weighted average of the means from the two normal distributions), is very similar to the estimate from the traditional meta-analysis, but now, as the continuous red line in Figure 7b shows, we have a much better fit of the distribution of treatment effects.

For the Academic Journals sample, we estimate a very significant (and quite precisely estimated) degree of publication bias: papers with no statistically significant results only have one tenth the probability of being published as studies with significant results ($\hat{\gamma}_{AJ} = 0.10$, s.e.=0.10). This parallels the non-parametric estimate from the t -statistics distribution in Figure 5c of of $\gamma = 2/9$. Accounting for publication bias has a vast impact on the estimated average impact of the nudges, which falls to 3.16 pp., quite a bit lower than the unweighted estimate of 8.7 pp, but still higher (though not significantly so) than the estimated Nudge Unit ATE. As Figure 7a shows (continuous red line), this model fits the distribution of treatment effects much better.

Robustness. Allowing for a flexible distribution of treatment effects is critical. Assuming a normal distribution even with a correction for selective publication (Panel

¹⁰The mixture of two normals has been suggested as a more flexible assumption for meta-analyses as early as Bohning, Dietz, and Schlattmann (1998) and van Houwelingen, Arends, and Stijnen (2002).

A of Table A9) would lead to a biased estimate of the parameters, as apparent from the poor fit displayed in Figure A7b. Conversely, a further non-parametric extension allowing for a mixture of three normals (Panel C of Table A9) leads to similar results, confirming that a mixture of two normals provides enough parametric flexibility.

As an alternative approach, in Online Appendix A.7 we assume that the Academic Journals trials and the Nudge Unit trials come from the same underlying distribution of results, modeled with a mixture of three normals, but the two sets of trials are drawn with different probabilities from the normals. We similarly estimate a high degree of selective publication and an ATE for the Academic Journals sample at 2.72 pp. (Tables A9b-c), again suggesting a somewhat larger impact than the Nudge Unit trials.

Summary. For the Nudge Units trials, the meta-analytic estimate is consistent with the unweighted estimate of 1.4 pp. For the Academic Journals trials, accounting for selective publication shrinks the estimated effect size from an unweighted average of 8.7 pp. to 3.16 pp. This model provides clear evidence of selective publication: we can reject the null of no publication bias in the Academic Journals sample at all traditional levels of significance. The model also provides suggestive evidence that, after accounting for selective publication, the Academic Journal nudges still have a larger effect size than the Nudge Units ones; however, this difference is not statistically significant.

4.5 Counterfactuals

The model in Panel B of Table 5 allows us to present counterfactuals to illustrate the role of each of the three dimensions in our model on the average effect size. In Row 1 of Table 6, we simulate the average treatment effect (ATE) for the Academic Journals sample given the estimated distribution of effect size, the observed statistical power, and the estimated degree of selective publication ($\gamma_{AJ} = 0.10$) from Panel B of Table 5. We reproduce the observed large impact of nudges, at 7.85 pp. (s.e.=1.35).

In the next series of counterfactuals, we change each component, and report the implied effect size. In Row 2 we only change the statistical power of the estimates: if the Academic Journal nudges had the same distribution of effect size and were still subject to publication bias, but had more precise standard errors (due to larger sample sizes) as in the Nudge Units sample, would the observed point estimate change? The point estimate would indeed be lower, at 6.64 pp., as the more precise estimate would blunt

to some extent the impact of selective publication, albeit not by much.

In Row 3, we show the impact of instead removing publication bias: the impact in this case is very sizable, since it would reveal the true average effect size, which is 3.16 pp. (as in Table 5, Panel B). With no publication bias, improving the statistical power would not have a further impact (Row 4).

In the next counterfactuals, we consider the Nudge Unit trials, but add features of the Academic Journals sample. With both low statistical power and publication bias (Row 5), the observed effect size would be 3.35 pp., more than twice the observed effect size (reproduced in Row 8). With publication bias but high statistical power (Row 6), the observed effect size would still be biased upward, but less so, at 2.39 pp. This further clarifies the potential of selective publication to bias average effect sizes, especially when the studies are not well-powered. Finally, in the absence of publication bias (Row 7), having lower statistical power would not impact the average point estimate.

These counterfactuals point to the critical role of publication bias, compounded by low statistical power, in explaining the large effect size in the Academic Journals sample.

4.6 Additional Differences

While we have focused on three main components—statistical power, effect size, and selective publication—there are other differences in the nudge experimentation model between the two samples that could affect the effect sizes. We consider here some of the possible alternatives, weaving in anecdotal evidence on how Nudge Units make decisions.

First, it is possible that Nudge Units, more than researchers, can do repeated experiments with one partner, and thus finesse their design (leading to larger effect sizes over time), or vice versa, run out of creative nudge solutions to a problem (leading to smaller effect sizes over time). Series of collaborations with one government partner, however, are not too common. In 62% of cases, there are at most two experiments with a given government partner (in our sample). Furthermore, in cases in which there are at least two experiments with the same partner, we do not see differences in nudge effect sizes between experiments run earlier versus later, as Figure A8a shows.

Another possibility is that the extent of follow-on experiments with one government unit may depend on the effect size in the first intervention. For example, if government units persist experimenting in easier-to-affect settings in which the initial nudge

effect size is large, this may potentially upward bias the point estimate. Figure A8b, though, shows that there is no systematic relationship between the effect size in the first intervention and the total number of collaborations.

These findings match anecdotal evidence from interviews with the leadership of both Nudge Unit teams. The likelihood of working on an additional trial with a given government partner seems to be largely orthogonal to the results of any given trial, and is often decided before final results on one trial are completed. For example, at BIT, contracts on how many trials to run per city partner were set in advance of conducting those trials. Both Nudge Units noted that the decision to run a trial depends more on feasibility, such as availability of administrative data, technical feasibility to nudge, a large enough sample size, as well as funding constraints. This would not appear to bias the estimate of the difference in effect sizes in an obvious way.

5 Discussion and Conclusion

An ongoing question in both policy circles and in academia asks: what would it look like if governments began using the “gold standard of evaluation”—RCTs—more consistently to inform policy decisions? While this has not yet happened at scale, nudge interventions have been used frequently and consistently through Nudge Units in governments, thus creating an opportunity to measure what taking nudges to scale might look like.

By studying the universe of trials run across two large Nudge Units in the U.S., covering over 23 million people, and comparing our results to published meta-analyses, we make progress on this question. We find that, on average, nudge interventions have a meaningful and statistically significant impact on the outcome of 1.4 pp. This estimated effect is significantly smaller than in academic journal articles (8.7 pp.). Using a meta-analysis model, we decompose this difference. We show that the largest source of the discrepancy is selective publication in the Academic Journals sample, exacerbated by low statistical power in that sample. After accounting for selective publication, we estimate an average nudge impact in the published papers of 3.2 pp.. This estimate is larger than the Nudge Unit estimate, consistent with some different features of the nudges used at scale (though the difference is not significant).

We hope that future research will expand on our work in a number of directions. First, we do not observe the micro-data for each trial, limiting the ability to, for example,

estimate which outcome is more elastic to different interventions. Second, it will be interesting to see whether other Nudge Units achieve different effect sizes, especially if they use a different mix of interventions, such as changing defaults or changing the choice architecture. Third, it would be valuable to examine determinants of which government departments decide to select into working with Nudge Units, and why they do so. Fourth, we hope that future research will evaluate the extent to which the results of the Nudge Unit interventions are implemented by the government units.

References

- Allcott, Hunt. 2015. "Site Selection Bias in Program Evaluation." *Quarterly Journal of Economics* 130 (3): 1117–1165.
- Andrews, Isaiah and Maximilian Kasy. 2019. "Identification of and Correction for Publication Bias." *American Economic Review* 109 (8): 2766-94.
- Azevedo, Eduardo M., Alex Deng, José Luis Montiel Olea, Justin Rao, and E. Glen Weyl. Forthcoming. "A/B Testing with Fat Tails." *Journal of Political Economy*.
- Banerjee, Abhijit V. and Esther Duflo. 2009. "The Experimental Approach to Development Economics." *Annual Review of Economics* 1: 151-178.
- Benartzi, Shlomo, John Beshears, Katherine L. Milkman, Cass R. Sunstein, Richard H. Thaler, Maya Shankar, Will Tucker-Ray, William J. Congdon, and Steven Galing. 2017. "Should Governments Invest More in Nudging?" *Psychological Science*. 28 (8): 1041-1055.
- Bhargava, Saurabh and Daylan Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review*. 105 (11): 3489-3529.
- Böhning, D., Dietz, E. and Schlattmann, P. 1998. "Recent developments in computer-assisted analysis of mixtures." *Biometrics* 54 (2): 525-36.
- Bold, Tessa, Mwangi Kimenyi, Germano Mwabu, Alice Ng'ang'a, Justin Sandefur. 2018. "Experimental evidence on scaling up education reforms in Kenya." *Journal of Public Economics* 168: 1-20.
- Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg. 2016. "Star Wars: The Empirics Strike Back" *AEJ: Applied Economics* 8 (1): 1-32.

- Brodeur, Abel, Nikolai Cook, and Anthony Heyes. 2020. "Methods Matter: p-Hacking and Publication Bias in Causal Analysis in Economics." *American Economic Review* 110 (11): 3634-3660.
- Bronchetti, Erin Todd, Thomas S. Dee, David B. Huffman, and Ellen Magenheimer. 2013. "When a Nudge Isn't Enough: Defaults and Saving among Low-income Tax Filers." *National Tax Journal* 66 (3): 609-634.
- Camerer, Colin F., et al. 2016. "Evaluating Replicability of Laboratory Experiments in Economics." *Science* 351 (6280): 1433-1436.
- Card, David and Alan B. Krueger. 1995. "Time-Series Minimum-Wage Studies: A Meta-analysis." *American Economic Review, Papers and Proceedings* 85 (2): 238-243.
- Card, David, Jochen Kluve, and Andrea Weber. 2018. "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations." *Journal of the European Economic Association* 16 (3): 894-931.
- Christensen, Garrett and Edward Miguel. 2018. "Transparency, Reproducibility, and the Credibility of Economics Research." *Journal of Economic Literature* 56 (3): 920-980.
- Cohen, J. 1965. "Some statistical issues in psychological research." in B. B. Wolman (Ed.), *Handbook of clinical psychology*. New York: McGraw-Hill.
- Dehejia, Rajeev, Cristian Pop-Eleches, and Cyrus Samii. 2019. "From Local to Global: External Validity in a Fertility Natural Experiment." *Journal of Business and Economic Statistics*.
- DerSimonian, Rebecca and Nan Laird. 1986. "Meta-Analysis in Clinical Trials." *Controlled Clinical Trials* 7 (3): 177-88.
- DellaVigna, Stefano, and Devin Pope. 2018. "What Motivates Effort? Evidence and Expert Forecasts." *Review of Economic Studies* 85: 1029-1069.
- DellaVigna, Stefano, Devin Pope, and Eva Vivaldi. 2019. "Predict science to improve science." *Science* 366 (6464): 428-429.
- Franco, Annie, Neil Malhotra, Gabor Simonovits. 2014. "Publication bias in the social sciences: Unlocking the file drawer." *Science* 345 (6203): 1502-1505.
- Frankel, Alexander, and Maximilian Kasy. 2020. "What findings should be published?" Working paper.
- Hagmann, David, Emily Ho, and George Loewenstein. 2019. "Nudging out support for a carbon tax." *Nature Climate Change* 9: 484-489.

- Hallsworth, Michael, John A. List, Robert D. Metcalfe, and Ivo Vlaev. 2017. "The behavioralist as tax collector: Using natural field experiments to enhance tax compliance." *Journal of Public Economics* 148 (C): 14-31.
- Halpern D. 2015. *Inside the Nudge Unit: How Small Changes Can Make a Big Difference*. London, UK: WH Allen.
- Hummel, Denis and Alexander Maedche. 2019. "How Effective Is Nudging? A Quantitative Review on the Effect Sizes and Limits of Empirical Nudging Studies." *Journal of Behavioral and Experimental Economics* 80: 47-58.
- Jachimowicz, Jon M., Duncan, Shannon, Weber, Elke U., and Johnson, Eric. J. 2019. "When and why defaults influence decisions: a meta-analysis of default effects." *Behavioral Public Policy* 3 (2): 159-186.
- Johnson et al. 2012. "Beyond Nudges: Tools of a Choice Architecture." *Marketing Letters* 23: 487-504.
- Laibson, David. 2020. "Nudges are Not Enough: The Case for Price-Based Paternalism". <https://www.aeaweb.org/webcasts/2020/aea-afa-joint-luncheon-nudges-are-not-enough>.
- Meager, Rachael. 2019. "Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments." *American Economic Journal: Applied Economics* 11 (1): 57-91.
- Miguel et al. 2014. "Promoting Transparency in Social Science Research." *Science* 10.1126/science.1245317.
- Milkman et al. 2020. "A mega-study approach to evaluating interventions." Working paper.
- Munscher, Robert, Max Vetter, and Thomas Scheuerle. 2016. "A Review and Taxonomy of Choice Architecture Techniques." *Journal of Behavioral Decision Making* 29: 511-524.
- Muralidharan, Karthik and Paul Niehaus. 2017. "Experimentation at Scale." *Journal of Economic Perspectives* 31 (4): 103-24.
- OECD. 2017. Behavioural insights and public policy: Lessons from around the world. OECD.
- Paule, Robert C. and John Mandel. 1989. "Consensus Values, Regressions, and Weighting Factors." *Journal of Research of the National Institute of Standards and Technology* 94 (3): 197-203.

- Shapiro, Bradley, Hitsch, Gunter J., and Tuchman, Anna. 2020. “Generalizable and robust TV advertising effects.” Working paper.
- Simonsohn, Uri, Leif D. Nelson, and Joseph P. Simmons. 2014. “P-curve: A key to the file-drawer.” *Journal of Experimental Psychology: General* 143 (2), 534–547.
- Sunstein, Cass. 2014. “Nudging: A Very Short Guide.” *Journal of Consumer Policy* 37: 583–588.
- Thaler, Richard and Cass Sunstein. 2008. *Nudge*. New Haven, CT: Yale University Press.
- van Houwelingen, Hans C., Arends, Lidia R., and Stijnen, Theo. 2002. “Advanced methods in meta-analysis: multivariate approach and meta-regression.” *Statistics in Medicine* 21: 589–624.
- Vivalt, Eva. 2020. “How Much Can We Generalize from Impact Evaluations?” *Journal of the European Economic Association* 18 (6), 3045–3089.

Figure 1: Example of nudges

(a) OES example: Control communication

GROUP A ROTH TSP: SMARTDOCS for January 2, 2015

Subject: Important! Your Action Needed in January to Continue Your Roth TSP Election

As a Roth TSP participant, your window to submit new contribution elections is here. You may submit your new Roth TSP elections based on percentages of basic pay, special pay, incentive pays and bonuses any time through Jan. 31, 2015, to avoid any interruption in your retirement investment plans.

Your elections may be submitted quickly and securely using myPay. You may also use the revised TSP-U-1 form available at www.tsp.gov. Forms must be submitted to your finance office to be applied to your military pay account.

We will send you reminders throughout January to make sure you have the information, worksheets and time to get your Roth TSP elections completed within the allotted time.

Election submissions received after Jan. 31, 2015, will result in a lapse in Roth TSP contributions.

For more information on the change to percentage-of-pay selections and how you can make sure your investment plans continue, visit www.dfas.mil/TSP_AC.html.

My POC for this effort is Matthew Ogles at matthew.ogles@dfas.mil

Matthew S. Ogles
Director, ESS Military Pay

(b) OES example: Treatment communication

GROUP B ROTH TSP: SMARTDOCS for January 2, 2015

Subject: Roth TSP - You Must Take Action Now to Avoid Interrupting Your 2015 Retirement Investment Contribution

Dear Servicemember,

It's a New Year! Re-enroll in your Roth TSP by submitting your new contribution percentages today! Because of changes to the way contributions are now being calculated, you must re-enroll this January or your contributions will be stopped February 1.

Avoid interrupting contributions by taking these three simple steps:

- 1) Log in at mypay.dfas.mil
- 2) Click on the "Traditional TSP and Roth TSP" link.
- 3) Enter your Roth TSP contribution percentages of basic, special, incentive, and bonus pay.

For more information on the change to percentage-of-pay selections, visit www.dfas.mil/TSP_AC.html. If you prefer to use a paper form, complete the TSP-U-1 form available at tsp.gov and submit it to your finance office.

Matthew Ogles (matthew.ogles@dfas.mil) is the POC for this Roth TSP update.

Sincerely,

Matthew S. Ogles
Director, ESS Military Pay

PS. Start 2015 off on the right foot - go to mypay.dfas.mil and take care of your future today. Make continuing your retirement investment plans an easy to do New Year's resolution.

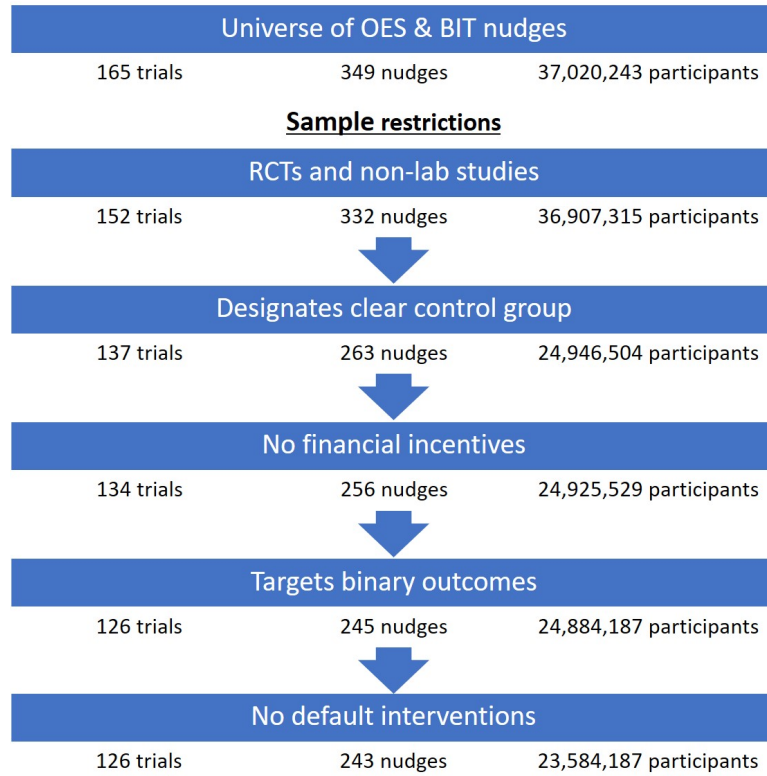
Annotations:

- Personalization² (points to Subject line)
- Fresh Start³ (points to "It's a New Year!" paragraph)
- Clear Action Steps⁴ (points to the three-step list)
- Loss Frame¹ (points to "You Must Take Action Now to Avoid Interrupting..." and "Avoid interrupting contributions...")
- Plain Language⁵ (points to the paragraph about information and paper forms)
- Postscript⁶ (points to the PS paragraph)

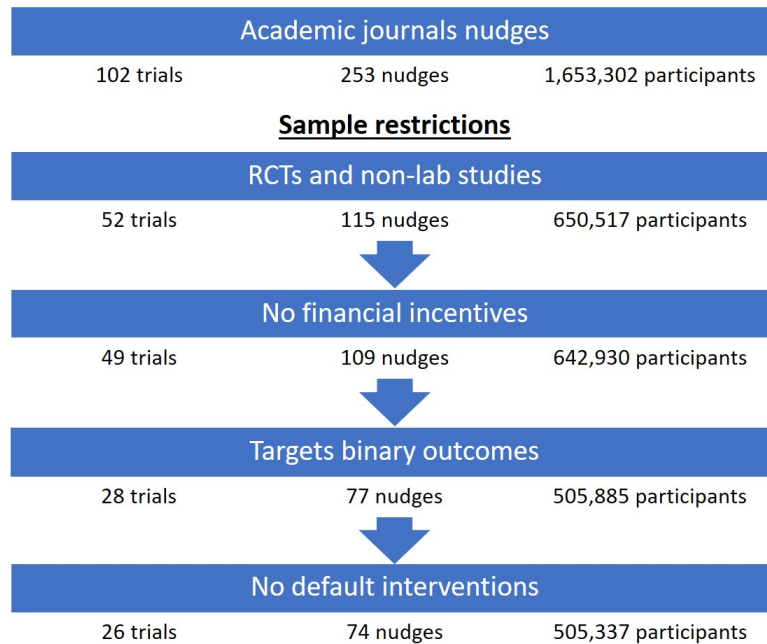
Figures 1a and 1b present an example of a nudge intervention from OES. This trial aims to increase service-member savings plan re-enrollment. The control group received the status-quo email (reproduced in Figure 1a), while the treatment group received a simplified, personalized reminder email with loss framing and clear action steps (reproduced in Figure 1b). The outcome in this trial is measured as savings plan re-enrollment rates.

Figure 2: Selection of nudge studies

(a) Selection among Nudge Units




(b) Selection among Academic Journals




This figure shows the number of trials, treatments, and participants remaining after each sample restriction.

Figure 3: Nudge treatment effects

(a) Academic Journals sample




(b) Nudge Units sample



This figure plots the treatment effect relative to control group take-up for each nudge with the quadratic fit. Nudges with extreme treatment effects are labeled for context. Error bars show 95% confidence intervals.

Figure 4: Minimum detectable effects and forecasts

(a) Minimum detectable effect sizes



(b) Forecasts by background





Figure 4a plots the CDF of the minimum detectable effects (MDE), or the size of the treatment effect that each treatment arm is powered to statistically detect 80% of the time given the control group take-up rate and the sample size. For 4 nudges (2 trials) in the Nudge Units sample that are missing control take-up data, the control group result is set to 50% to estimate a conservative measure of the MDE. Control take-up is bounded below at 1% when calculating MDE.


Figure 4b shows various distributions of forecasts made by Nudge Unit practitioners and academics (university faculty and post-docs) on the treatment effect of nudges.

Figure 5: Publication bias tests: Academic Journals


(a) Point estimate and minimum detectable effect



(b) t -stat distribution




(c) Most significant nudges by trial




This panel displays tests for publication bias in the Academic Journals sample. Figure 5a plots the relationship between the minimum detectable effect and the treatment effect size. The estimated equation is the linear fit with standard errors clustered at the trial level. Figure 5b shows the distribution of t -statistics (i.e., treatment effect divided by standard error) for all nudges, and Figure 5c shows the distribution for only the max t -stat within each trial. Figure 5c excludes 1 trial in which the most significant treatment arm uses financial incentives.

Figure 6: Publication bias tests: Nudge Units


(a) Point estimate and minimum detectable effect



(b) t -stat distribution




(c) Most significant nudges by trial




This panel displays tests for publication bias in the Nudge Units sample. Figure 6a plots the relationship between the minimum detectable effect and the treatment effect size. The estimated equation is the linear fit with standard errors clustered at the trial level. Figure 6b shows the distribution of t -statistics (i.e., treatment effect divided by standard error) for all nudges, and Figure 6c shows the distribution for only the max t -stat within each trial. Figure 6c excludes 2 trials in which the most significant treatment arm uses defaults/financial incentives.

Figure 7: Simulated densities from maximum likelihood and mixture of normals models

(a) Academic Journals



(b) Nudge Units



This figure plots the empirical histogram of observed nudge effects and compares the fit of a normal-based meta-analysis model (Panel A of Table 5) to the fit of a mixture of two normals model (Panel B of Table 5) for the Academic Journals sample with a correction for publication bias in Figure 7a for the Nudge Units sample in Figure 7b. 1 nudge in the Nudge Units sample with an effect less than -10 pp. and 3 nudges in the Academic Journals sample with effects greater than 35 pp. are not shown. The densities are kernel approximations from 500,000 simulated trials.

Table 1: Comparison of nudge categories

| | Nudge Units | | | Academic Journals | | |
|--------------------------------------|-------------|-----------------|-----------|-------------------|-----------------|-----------|
| | Freq. (%) | Nudges (Trials) | ATE (pp.) | Freq. (%) | Nudges (Trials) | ATE (pp.) |
| <i>Date</i> | | | | | | |
| Early* | 46.50 | 113 (49) | 1.84 | 48.65 | 36 (14) | 7.10 |
| Recent* | 53.50 | 130 (77) | 0.97 | 51.35 | 38 (12) | 10.18 |
| <i>Policy area</i> | | | | | | |
| Revenue & debt | 28.81 | 70 (30) | 2.43 | 17.57 | 13 (4) | 3.60 |
| Benefits & programs | 22.22 | 54 (26) | 0.89 | 10.81 | 8 (3) | 14.15 |
| Workforce & education | 18.52 | 45 (24) | 0.49 | 9.46 | 7 (2) | 2.56 |
| Health | 13.17 | 32 (18) | 0.65 | 28.38 | 21 (9) | 8.98 |
| Registration & regulation compliance | 8.64 | 21 (16) | 2.18 | 12.16 | 9 (2) | 3.16 |
| Community engagement | 7.82 | 19 (10) | 0.74 | 4.05 | 3 (2) | 2.80 |
| Environment | 0.82 | 2 (2) | 6.83 | 13.51 | 10 (3) | 22.95 |
| Consumer behavior | 0 | 0 (0) | – | 4.05 | 3 (1) | 3.19 |
| <i>Medium of communication</i> | | | | | | |
| Email | 39.51 | 96 (47) | 1.09 | 12.16 | 9 (6) | 3.75 |
| Physical letter | 29.63 | 72 (44) | 2.41 | 16.22 | 12 (4) | 1.67 |
| Postcard | 21.40 | 52 (21) | 0.82 | 6.76 | 5 (1) | 10.46 |
| Website | 2.88 | 7 (4) | -0.04 | 12.16 | 9 (3) | 6.24 |
| In person | 0.82 | 2 (2) | 3.05 | 28.38 | 21 (4) | 14.82 |
| Other | 11.11 | 27 (14) | 1.17 | 24.32 | 18 (8) | 9.38 |
| <i>Control group receives:</i> | | | | | | |
| No communication | 60.91 | 148 (66) | 1.42 | 43.24 | 32 (9) | 10.91 |
| Some communication | 39.09 | 95 (60) | 1.30 | 56.76 | 42 (17) | 6.99 |
| <i>Mechanism</i> | | | | | | |
| Simplification & information | 58.02 | 141 (70) | 1.19 | 5.41 | 4 (1) | 16.34 |
| Personal motivation | 57.20 | 139 (73) | 1.76 | 32.43 | 24 (7) | 9.59 |
| Reminders & planning prompts | 31.28 | 76 (47) | 2.54 | 35.14 | 26 (11) | 5.02 |
| Social cues | 37.04 | 90 (43) | 0.84 | 21.62 | 16 (5) | 13.81 |
| Framing & formatting | 31.69 | 77 (41) | 1.38 | 32.43 | 24 (8) | 13.53 |
| Choice design | 6.17 | 15 (8) | 7.01 | 20.27 | 15 (9) | 8.85 |
| Total | 100 | 243 (126) | 1.37 | 100 | 74 (26) | 8.68 |

This table shows the number of nudges and trials in each category, and the average treatment effect within each category. Frequencies for *Medium* and *Mechanism* are not mutually exclusive and frequencies may not sum to 1.

**Early* refers to trials implemented between 2015-2016 for Nudge Units, and to papers published in 2014 or before for Academic Journals. *Recent* refers to trials and papers after these dates.

Table 2: Comparison of trial features

| | Academic Journals | Nudge Units | | | |
|--|-------------------|---|------------------------|-------------------------|-------------------------|
| | Mean [std. dev.] | Mean [std. dev.; <i>p</i> -value of difference from column 1] | | | |
| | | All | BIT | OES | Academic-affiliated OES |
| | (1) | (2) | (3) | (4) | (5) |
| <i>Academic faculty involvement</i> | 100% | 19% | 0% | 50% | 100% |
| <i>Outcome features</i> | | | | | |
| Control group take-up (%) | 26.0 | 17.2 | 15.4 | 19.5 | 26.4 |
| | [19.9] | [23.1; <i>p</i> =0.09] | [23.7; <i>p</i> =0.05] | [22.2; <i>p</i> =0.29] | [24.0; <i>p</i> =0.94] |
| Outcome time-frame (days) | 68.7 | 62.4 | 41.9 | 101.7 | 141.5 |
| | [91.7] | [76.4; <i>p</i> =0.67] | [45.0; <i>p</i> =0.15] | [104.9; <i>p</i> =0.25] | [110.9; <i>p</i> =0.04] |
| <i>Trial design</i> | | | | | |
| Mechanisms per treatment arm | 1.5 | 2.2 | 2.0 | 2.5 | 2.3 |
| | [0.7] | [1.0; <i>p</i> =0.00] | [1.0; <i>p</i> =0.00] | [0.9; <i>p</i> =0.00] | [0.9; <i>p</i> =0.00] |
| Treatment arms per trial | 2.8 | 1.9 | 1.7 | 2.3 | 1.9 |
| | [2.1] | [1.7; <i>p</i> =0.04] | [1.0; <i>p</i> =0.01] | [2.5; <i>p</i> =0.31] | [1.5; <i>p</i> =0.06] |
| Minimum detectable effect (pp.) | 8.2 | 1.7 | 2.2 | 1.2 | 1.7 |
| | [6.4] | [2.2; <i>p</i> =0.00] | [2.6; <i>p</i> =0.00] | [1.6; <i>p</i> =0.00] | [2.2; <i>p</i> =0.00] |
| Institutional constraints rating (1-5) | 4.0 | 3.0 | 3.0 | 3.0 | 2.8 |
| | [0.9] | [0.6; <i>p</i> =0.00] | [0.5; <i>p</i> =0.00] | [0.7; <i>p</i> =0.01] | [1.3; <i>p</i> =0.00] |
| <i>Planning and implementation</i> | | | | | |
| Total duration (months) | 21.3 | 11.1 | 8.6 | 15.0 | 17.0 |
| | [16.1] | [3.9; <i>p</i> =0.00] | [1.3; <i>p</i> =0.00] | [3.3; <i>p</i> =0.09] | [8.3; <i>p</i> =0.24] |
| Planning (including IRB) | 6.6 | 4.6 | 4.0 | 5.6 | 5.1 |
| | [6.1] | [2.3; <i>p</i> =0.17] | [1.1; <i>p</i> =0.06] | [3.4; <i>p</i> =0.61] | [2.5; <i>p</i> =0.28] |
| Intervention and data collection | 6.7 | 4.5 | 3.4 | 6.2 | 6.5 |
| | [7.1] | [2.0; <i>p</i> =0.16] | [1.2; <i>p</i> =0.03] | [1.8; <i>p</i> =0.77] | [2.3; <i>p</i> =0.91] |
| Data analysis and write-up | 7.8 | 2.0 | 1.3 | 3.2 | 3.9 |
| | [7.0] | [1.2; <i>p</i> =0.00] | [0.5; <i>p</i> =0.00] | [1.1; <i>p</i> =0.00] | [2.9; <i>p</i> =0.01] |
| Personnel full-time equivalent months | 14.9 | 5.8 | 4.3 | 8.3 | 6.2 |
| | [18.1] | [4.9; <i>p</i> =0.03] | [2.8; <i>p</i> =0.01] | [6.9; <i>p</i> =0.17] | [2.8; <i>p</i> =0.02] |
| Number of survey responses | 25 | 13* | 8* | 5* | 24 |
| Number of trials | 26 | 126 | 78 | 48 | 24 |

Data on the institutional constraints rating, duration, and personnel FTE months were collected from a survey of the researchers involved in the trials (see text and Appendix Section A.3 for details). Outcome duration is capped at 360 days, which only affects 1 trial in each of the Academic Journal and Nudge Unit samples. *In columns 2 to 4, the number of survey responses corresponds to the number of Nudge Unit staff members in leadership roles whom we surveyed.

Table 3: Unweighted treatment effects

| | Academic Journals | Nudge Units | | | |
|--|-------------------|------------------|------------------|------------------|--------------------------------|
| | (1) | All (2) | BIT (3) | OES (4) | Academic-affiliated OES (5) |
| Average treatment effect (pp.) | 8.682 (2.467) | 1.375 (0.302) | 1.666 (0.521) | 1.023 (0.206) | 0.978 (0.408) |
| Nudges | 74 | 243 | 133 | 110 | 45 |
| Trials | 26 | 126 | 78 | 48 | 24 |
| Observations | 505,337 | 23,577,537 | 2,029,731 | 21,547,806 | 8,923,186 |
| Average control group take-up (%) | 25.97 | 17.20 | 15.39 | 19.47 | 26.45 |
| <i>Distribution of treatment effects</i> | | | | | |
| 25th percentile | 1.05 | 0.04 | 0.00 | 0.15 | 0.10 |
| 50th percentile | 4.12 | 0.50 | 0.40 | 0.60 | 0.42 |
| 75th percentile | 12.00 | 1.40 | 1.63 | 1.22 | 1.20 |

This table shows the average treatment effect of nudges. Standard errors clustered by trial are shown in parentheses. pp. refers to percentage point.

Table 4: Predicting nudge effect sizes

| Dep. Var.: Treatment effect (pp.) | Full sample | | | | Academic-affiliated only | | |
|---|------------------|--------------------|-------------------|-------------------|--------------------------|-------------------|-------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Constant | 1.375 (0.302) | 4.235 (2.125) | 1.012 (0.340) | 2.781 (1.972) | 0.978 (0.405) | 4.117 (4.884) | 1.970 (4.404) |
| Omitted group: Nudge Units | | | | | | | |
| Academic Journals | 7.307 (2.449) | 2.353 (1.609) | -0.896 (1.930) | -0.026 (1.954) | 7.704 (2.487) | 6.122 (1.972) | -1.778 (2.693) |
| Publication bias controls (Egger's test) | | | | | | | |
| Minimum detectable effect (MDE) | | | 0.210 (0.246) | 0.227 (0.270) | | | -0.084 (0.168) |
| Academic Journals×MDE | | | 0.837 (0.386) | 0.351 (0.371) | | | 1.076 (0.372) |
| Nudge categories | | | | | | | |
| <i>Policy area</i> | | | | | | | |
| Benefits & programs | | -0.266 (1.003) | | -0.267 (0.924) | | | |
| Workforce & education | | -2.323 (1.006) | | -2.482 (0.943) | | | |
| Health | | -0.823 (1.533) | | -1.776 (1.449) | | | |
| Registrations & regulation compliance | | -1.012 (1.349) | | -0.991 (1.335) | | | |
| Community engagement | | -1.648 (1.594) | | -1.468 (1.289) | | | |
| Environment | | 9.390 (4.985) | | 5.561 (4.900) | | | |
| Consumer behavior | | -10.885 (3.667) | | -7.292 (3.562) | | | |
| <i>Medium of communication</i> | | | | | | | |
| Email | | -1.860 (1.423) | | -1.514 (1.382) | | | |
| Physical letter | | -0.836 (1.205) | | -0.302 (1.152) | | | |
| Postcard | | 0.137 (1.511) | | -0.015 (1.357) | | | |
| Website | | -2.127 (3.162) | | -1.382 (2.719) | | | |
| In person | | 7.237 (3.166) | | 5.386 (3.431) | | | |
| Other | | -0.341 (1.689) | | -0.100 (1.618) | | | |
| <i>Control group receives:</i> | | | | | | | |
| Some communication | | -1.194 (0.944) | | -1.200 (0.884) | | | |
| <i>Mechanism</i> | | | | | | | |
| Simplification & information | | 0.871 (1.118) | | 0.866 (1.208) | | | |
| Personal motivation | | -0.496 (0.849) | | -0.324 (0.909) | | | |
| Reminders & planning prompts | | 0.325 (0.831) | | 0.777 (0.774) | | | |
| Social cues | | 0.035 (0.961) | | 0.252 (0.923) | | | |
| Framing & formatting | | 1.224 (0.931) | | 0.978 (0.905) | | | |
| Choice design | | 6.199 (2.353) | | 5.508 (2.310) | | | |
| Trial features | | | | | | | |
| Control take-up (%) | | 0.106 (0.059) | | 0.045 (0.055) | | | |
| Control take-up ² | | -0.001 (0.001) | | -0.001 (0.001) | | | |
| Log(outcome time-frame days) | | -0.667 (0.401) | | -0.282 (0.356) | | | |
| Ideal nudge implemented rating (1-5) | | | | | 0.979 (1.291) | 0.467 (0.731) | |
| Log(personnel FTE months) | | | | | 0.671 (0.857) | 0.902 (0.711) | |
| Log(planning & implementation months) | | | | | -2.721 (1.562) | -1.419 (1.548) | |
| Nudges | 317 | 317 | 317 | 317 | 119 | 119 | 119 |
| Trials | 152 | 152 | 152 | 152 | 50 | 50 | 50 |
| R-squared | 0.18 | 0.46 | 0.39 | 0.49 | 0.14 | 0.22 | 0.45 |

This table shows OLS estimates with standard errors clustered by trial in parentheses. The MDE (minimum detectable effect) is calculated in pp. at power 0.8. Observations with missing data for outcome time-frame, control take-up result, trial duration, institutional constraints rating, or personnel FTE months are included with separate dummies.

Table 5: Generalized meta-analysis models

| | | Normal 1 | | | Normal 2 | | | | | |
|--|-----------------------------|----------------------------|-----------------|--------------------|--------------------|-----------------|--------------------|--------------------|----------------------|-----------------|
| | ATE (pp.) | $\hat{\gamma}$ (pub. bias) | $\hat{\beta}_1$ | $\hat{\tau}_{BT1}$ | $\hat{\tau}_{WI1}$ | $\hat{\beta}_2$ | $\hat{\tau}_{BT2}$ | $\hat{\tau}_{WI2}$ | \hat{P} (Normal 1) | -Log likelihood |
| Panel A. <i>Traditional parametric normal-based meta-analysis</i> | | | | | | | | | | |
| Academic Journals | 8.58 (2.00) | 1 (fixed) | 8.58 (2.00) | 7.89 (2.09) | 5.65 (2.71) | — | — | — | 1 (fixed) | 267.69 |
| Nudge Units | 1.49 (0.37; p=0.00) | 1 (fixed) | 1.49 (0.37) | 3.06 (1.22) | 2.36 (1.29) | — | — | — | 1 (fixed) | 651.36 |
| Panel B. <i>Generalized mixture model with selective publication correction</i> | | | | | | | | | | |
| Academic Journals | 3.16 (1.89) | 0.10 (0.10) | 0.33 (1.02) | 2.69 (1.14) | 0.04 (0.19) | 19.17 (5.38) | 5.91 (3.23) | 12.69 (2.78) | 0.85 (0.07) | 211.21 |
| Nudge Units | 1.38 (0.56; $p = 0.37$) | 1 (fixed) | 0.34 (0.13) | 0.41 (0.15) | 0.24 (0.10) | 5.10 (1.67) | 4.65 (3.35) | 6.40 (3.44) | 0.78 (0.11) | 397.95 |

This table shows the estimates from a normal-based meta-analysis method, and also from a model with a mixture of two normals. Under the normal-based meta-analysis assumptions, trial base effects β_i are drawn from a normal distribution centered at $\bar{\beta}$ with between-trial standard deviation τ_{BT} . Then, each treatment arm j within a trial i draws a base treatment effect $\beta_{ij} \sim N(\beta_i, \tau_{WI}^2)$, where τ_{WI} is the within-trial standard deviation. Each treatment arm also has some level of precision given by an independent standard error σ_{ij} . The observed treatment effect is $\hat{\beta}_{ij} \sim N(\beta_{ij}, \sigma_{ij}^2)$. The mixture of two normals model is a generalization of the normal-based meta-analysis, and allows trial base effects to be drawn from a second normal distribution. To capture the extent of selective publication, the probability of publication is allowed to differ depending on whether trial have at least one significant treatment arm. In particular, trials without any significant results at the 95% level are γ times as likely to be published as trials with significant results. Estimates are obtained using maximum likelihood. Standard errors from at least 200 bootstrap samples are shown in parentheses. The p -value of the difference in the estimated average treatment effect (ATE) between the Academic Journals and Nudge Units samples is shown in the parentheses below the Nudge Unit ATE.

Table 6: Model counterfactuals

| | Effect size distribution | Statistical power | Selective publication | Simulated ATE (pp.) |
|--|--------------------------|-------------------|------------------------|---------------------|
| (1) Acad. J. as observed | Acad. J. | Acad. J. | Yes (as in Acad. J.) | 7.85 (1.35) |
| <i>Counterfactuals – Academic Journal effect sizes with:</i> | | | | |
| (2) High power | Acad. J. | Nudge Units | Yes (as in Acad. J.) | 6.64 (1.25) |
| (3) No pub. bias | Acad. J. | Acad. J. | No (as in Nudge Units) | 3.16 (0.83) |
| (4) High power & no pub. bias | Acad. J. | Nudge Units | No (as in Nudge Units) | 3.12 (0.84) |
| <i>Counterfactuals – Nudge Unit effect sizes with:</i> | | | | |
| (5) Low power & pub. bias | Nudge Units | Acad. J. | Yes (as in Acad. J.) | 3.35 (0.69) |
| (6) Pub. bias | Nudge Units | Nudge Units | Yes (as in Acad. J.) | 2.39 (0.55) |
| (7) Low power | Nudge Units | Acad. J. | No (as in Nudge Units) | 1.38 (0.38) |
| (8) Nudge Units as observed | Nudge Units | Nudge Units | No (as in Nudge Units) | 1.38 (0.36) |

This table shows estimates for counterfactual simulated average treatment effects using the generalized model in Table 5. Each counterfactual exercise draws 1,000 rounds of effect sizes for 152 simulated trials from the estimated mixture distribution for the sample of nudges indicated under “Effect size distribution”. The number of experimental arms and their standard errors for these simulated trials are drawn with replacement from the sample listed under “Statistical power”. Under selective publication, simulated trials without any positively significant treatment arms at the 95% level are “published” with probability $\hat{\gamma} = 0.1$ (as estimated in Table 5). Simulated trials with at least one positively significant treatment arm are published with probability 1. When selective publication is suppressed, all simulated trials are published. The “Simulated ATE (pp.)” column reports the average treatment effect in percentage points for all “published” treatment arms from the $1,000 \times 152 = 152,000$ simulated trials. The standard deviation of the observed ATE in the 1,000 rounds is reported in parentheses.

A Online appendix

A.1 Additional Details on Sample

Nudge Units sample. For one nudge treatment, the trial report does not list a point estimate and simply indicates a result that is not statistically significant, and we were not able to track down the exact finding; in this case, we impute the outcome trial effect as zero. For two other nudge treatments, the result was also indicated as “not significant” without a point estimate, but we were able to infer the point estimate from the figure presented in the trial report. The information on take-up in the control group is missing for 4 nudges (2 trials); we still use these trials in our main analysis, but not in the additional log odds analysis. Finally, 7 nudges (3 trials) have control take-up of 0%, and 1 nudge has treatment take-up of 0%; these cases are also not used in the log odds analysis, but remain in the primary analysis.

In determining the sample, we exclude default changes, as discussed in the text. We define default interventions as changing “which outcome happens *automatically* if an individual remains passive” (Bronchetti et al., 2013), as in the retirement savings defaults. Sometimes a nudge that is labeled as a default intervention in an academic paper or in a Nudge Unit report did not meet this requirement. An example is a “default” appointment, in which participants are scheduled into an appointment slot, for instance to get a flu shot; still, participants would not be vaccinated if they remain passive. For a meta-analysis on nudges using defaults, see Jachimowicz et al. (2019). Adding in the default trials into our sample does not meaningfully change our main result.

Academic Journals sample. The number of nudges and participants are approximated from the data made available by Hummel and Maedche (2019). We focused on recording the main results of the paper for binary outcomes. After we applied our sample criteria to the sample of papers from these two sources, we re-coded the treatment effect sizes, standard errors, number of nudges and participants, and additional features of the interventions from the original papers. We took the treatment effect and standard error if they were readily available, for instance, in the main table. There were various cases in which we had to manually compute the treatment effect and standard errors; for example, sometimes we used the proportion of take-up in the treatment and control groups, and in other times, we translated logit coefficients. We transcribed all the significant digits. We calculate t -stats by dividing the treatment effect by the standard error. We also checked that the bunching to the right of the significant $t = 1.96$ threshold in Figure 5c is not due to rounding and lack of significant digits. In the Academic Journals sample, the 3 significant max t -stats closest to the $t = 1.96$ threshold are 1.9993, 2.0286, and 2.1189, and the three corresponding papers indicate that these results are indeed significant at the 95% level.

A.2 Published Nudge Units sample

To our knowledge, only 16 of the 126 Nudge Unit trials have been written or published as academic papers so far. (We note that all the OES trials have a public trial report shared online with the results.) These papers are listed in Table A1a. This section presents results for this subsample of trials.

Table A3b shows the impact of the 33 nudge interventions in these 16 Published Nudge Unit

trials. As mentioned in the text, they have an average treatment effect of 0.97 pp. (s.e.=0.23), similar to the one for the Nudge Units full sample (1.38 pp.). These studies also have similar statistical power: a median MDE of 0.81 pp. versus 0.78 pp. in the overall Nudge Units sample. Thus, the studies written up as academic papers do not appear to differ overall from the full sample of Nudge Unit trials.

Is there selective publication out of the Published Nudge Unit trials? In Figure A9a-e, we first show the Card and Krueger (1995) graph and the funnel plot for this subsample separately, and find suggestive patterns of publication bias with a missing mass of insignificant trials. In Panel B of Table A9a, we apply the estimation of the meta-analysis model with selective publication to this sample, as we do for the main sample. We estimate the degree of selective publication directly, and confirm a significant degree of publication bias with $\hat{\gamma} = 0.07$ (s.e.=0.14), which interestingly is very similar to the estimate for the Academic Journals sample. Yet the estimated average true treatment effect for this subsample (0.36 pp.) does not display a large bias relative to the observed effect size.

These estimates clarify the two factors behind the much smaller impact of publication bias. First, the Nudge Unit trials, being at scale, have much less noise in the treatment effects. Second, they also have less heterogeneity in treatment effects across trials, as visible in the estimates for τ^2 . Both factors limit the impact of selective publication on the observed effect size.

A.3 Sample Criteria for Meta-Analyses

To build their data set of papers on nudges, Hummel and Maedche (2019) conduct a systematic literature review. They begin by searching three databases of academic articles (ScienceDirect, EBSCOHost, and AISEL) for papers that include “nudge” or “nudging” in the title, abstract, or keywords since 2008. This initial search returned 2493 papers. From these papers, they exclude those that do not reference Thaler and Sunstein (2008), do not relate to nudges in the behavioral context (e.g., papers in the natural sciences where “nudge” has a different meaning), or do not report effect sizes. Their final sample consists of 100 papers.

Benartzi et al. (2017) determine their sample of nudge interventions as follows. They identify a list of policy areas from the 2015 summary reports of the Social and Behavioral Sciences Team and BIT-UK, identify the main outcome for each policy area, and search for papers that evaluated nudges, tax incentives, rewards, or education programs targeting those outcomes in the leading academic journals as ranked by Google Scholar. They found 18 relevant papers for four policy areas (Financial security in retirement, Education, Energy, and Health), and they compare the cost-effectiveness of the 5 nudge interventions against the other 13 traditional levers (such as financial incentives) within each policy area. Of the 5 nudge interventions, 2 are already included in the Hummel and Maedche (2019) sample, 1 does not target a binary outcome, and the remaining 2 are added to form our starting sample of 102 papers.

A.4 Categorizing nudges

While this paper does not focus on a taxonomy of nudges (see Johnson et al., 2012, Sunstein, 2014, and Munscher, Vetter, and Scheuerle, 2016), we categorized each nudge under six mechanisms from the descriptions in the trial reports: Simplification, Personal motivation,

Reminders & planning prompts, Social cues, Framing & formatting, and Choice design.

These six categories are broader than the nine groups used in Hummel and Maedche (2019), which are (1) default, (2) simplification, (3) social reference, (4) change effort, (5) disclosure, (6) warnings/graphics, (7) precommitment, (8) reminders, and (9) implementation intentions. Since we exclude defaults from our sample, there are eight remaining groups that can be linked to our categorization. (2) and (4) are both part of our “Simplification” category; (3) falls under “Social cues”; (5) and (6) share characteristics with “Personal motivation” though some aspects (6) can also be considered as “Framing & formatting”; lastly, (7), (8), and (9) are subcategories in “Reminders & planning prompts.” We illustrate the six categories below with examples.

Simplification and information This category includes interventions that simplify the language or design, or provide new information. In the Nudge Units sample, one nudge aimed to increase response rates to the American Housing Survey by rewriting the description of the survey in plain language for the advance letter. Another nudge simplified the payment instructions sent to businesses for fire inspections, false alarms, and permit fees. In the Academic Journals sample, Bettinger et al. (2012) pre-filled fields using tax returns to make signing up for FAFSA easier.

Personal motivation This category covers nudges that try to influence the recipient’s perception of how the targeted action will affect him/her. Specifically, these interventions may inform of the benefits (costs/losses/risks) from (not) taking-up, such as, in the Nudge Units sample, emphasizing the benefits of the flu shot or warning that parking violation fees will be sent to collections agencies if not paid on time. Personalizing communications (e.g., including the homeowner’s name on a letter for delinquent property taxes) or providing encouragement/inspiration (e.g., encouraging medical providers to use electronic flow sheet orders) also fall under this category. An example in the Academic Journals sample is Luoto et al. (2014), which marketed the health benefits of water treatment technologies in Kenya and Bangladesh.

Reminders & planning prompts This category consists of (i) communications that remind recipients to take up, for instance, veteran health benefits for transitioning service-members, and (ii) planning prompts, which remind recipients of deadlines or induce them to plan/set goals. Suggesting an appointment is an example; in one Nudge Unit trial, nurses called pre- and post-natal mothers to schedule a home visit. In the Academic Journals sample, Nickerson and Rogers (2010) study the effect of implementation intentions (i.e., forming a concrete plan) on voter turnout.

Social cues This category captures mechanisms that draw on social norms, comparisons, prosocial behavior, and messenger effects. Examples in the Nudge Units sample include: informing parking violators that most fines are paid on time, comparing quetiapine prescription rates among doctors to reduce over-prescriptions, encouraging double-sided printing, and addressing postcards from officers to promote applying for the police force. Rommel et al. (2015) in the Academic Journals sample provide households stickers to adhere on their mailboxes and reject unsolicited junk mail. In one treatment, households are told the average amount of paper waste from junk mail, and in another social pressure treatment, households are notified that researchers will return to check whether the sticker had been applied.

Framing & formatting This category encompasses mechanisms that target how the information is framed, or the format of the communication, which can include images or the visual layout. In the Nudge Units sample, one trial tests various wording of the subject line for an email encouraging borrowers to submit a form for loan forgiveness, while another trial added a red “Pay Now” logo with a handwritten signature to a letter sent to sewer bill delinquents.

From the Academic Journals sample, Wansink and Hanks (2013) investigate how the layout and order of menu items in a buffet line affect selection of healthy foods.

Choice design This category contains active choice interventions, which prompt recipients into making a decision. Nudge Units have used active choice nudges to enroll service-members into retirement savings plans, and to raise donations for a charity. In the Academic Journals sample, Chapman et al. (2010) apply active choice to flu vaccinations, Carroll et al. (2009) to 401(k) enrollment, and Stutzer et al. (2011) to blood donations.

A.5 Survey of nudge researchers

To gather information on trial features, we surveyed the authors of Academic Journals papers and the university faculty affiliated with Nudge Unit trials in our sample. We received responses from all the authors, except for one paper in the Academic Journals sample. We also asked staff members from OES and BIT to fill out the survey for a typical trial that they have conducted. For four OES trials, the affiliated university faculty stated that they could not accurately estimate these trial features. Thus, we supplement or substitute their responses with the medians reported by OES staff members as shown in Table 2. We distributed the survey and collected the responses by email. The exact wording is below.

Duration *Roughly how many months did you actively work on this project from the initial design steps until the first report/draft of the paper? (We understand these are just best guesses so please feel free to round.)*

--- months

If you remember, can you decompose the total months of active work into:

--- months of planning the intervention before implementation in the field (includes negotiating with partnering organizations and getting IRB approval),

--- months of implementation and data collection, and

--- months of analyzing the data and writing the report/draft?

Personnel *Including co-authors and RAs, approximately how many months of full-time work went into your project(s)? (For example, if you worked 1 day/week for 18 months and had a full-time research assistant who worked on 4 projects for 2 years, then that would be $0.2*18+0.25*24=9.6$ months total of full-time work.)*

--- months of full-time work

Institutional constraints *Working in the field often involves changing an intervention to fit institutional and legal constraints (such as the IRB or preferences of the partnering organization). For your project(s), how close was the intervention that you ultimately implemented compared to the one that you would have ideally wanted to run? Please answer on a scale from 1 (vastly different) to 5 (exactly the same).*

--- (Scale: 1-5)

A.6 Forecasting Survey

This section provides more detail on the 10-minute survey eliciting forecasts from behavioral scholars using a convenience sample through email lists and Twitter ($n=237$). As stated in the main text, the survey explained the methodology of our analysis, described the two samples, showed participants three nudge interventions randomly drawn out of 14 exemplars, and asked for predictions of: (a) the average effect size for the Nudge Units sample; (b) the average effect

size for the Academic Journals sample and (c) the effect size for the three nudge examples shown. Throughout, we asked predictions in percentage point units, just as reported in this paper. The survey also asked participants how many field experiments they have conducted.

Specifically, we asked *“Across all trials, what do you expect the average effect of a nudge to be? Please enter your answer as a percentage point (p.p.) difference. The average take-up in the control group across the trials is around 17%.”* We also added as a footnote, *“For our analysis, we will be taking the average effect across all the nudges (formally, a meta-analysis under a random effects model).”*

For the Academic Journals sample, we stated: *“Two recent meta-analyses (Benartzi et al., 2017; Hummel & Maedche, 2019) studied nudges and other behavioral interventions that have been published in academic journals. From their list of published trials that use nudges, we have extracted the trials that are comparable to those in our OES and BIT data set. These published trials also: are randomized controlled trials, target a binary outcome, do not feature defaults or monetary incentives. What do you expect the average effect of a nudge to be for nudges from these published trials?”*

As Figure A10a shows, the 237 participants belong to four main categories: academic faculty (27.9%), graduate students (24.1%), employees of non-profits or government agencies (16.9%), employees in the private sector (15.2%), and practitioners in nudge units (11.8%). Overall, the respondents expect a larger nudge impact in the Academic Journals sample than in the Nudge Units sample, as we indeed find. The respondents also make a rather accurate prediction for the average effect size among Academic Journals nudges, with the median (average) forecast of 6 pp. (8.02 pp.), close to the 8.7 pp. we estimate. They, however, broadly overestimate the impact in the Nudge Units sample, with a median (average) prediction of 4 pp. (5.84 pp.), compared to the 1.38 percentage point we estimate. This miscalibration on the effect of a nudge at scale could lead to sub-optimal policy decisions when policymakers choose between implementing a nudge and using traditional levers, such as taxes. Indeed, Hagmann, Ho, and Loewenstein (2019) survey policymakers and find that over-optimism on the effectiveness of nudges “crowds out” support for taxes.

Interestingly, there is significant heterogeneity in these forecasts. In Figure A11b, we plot the predictions for the Nudge Unit results separately for researchers with no (reported) experience in running field experiments ($n=86$), for researchers with a sizable experience (having run at least 5 field experiments, $n=42$), and for practitioners working in Nudge Units ($n=28$). The median researcher with no experience expects an average impact of a Nudge Unit treatment of 5.00 pp., the median experienced researcher expects an impact of 3.50 pp., and the median nudge practitioner expects an average impact of 1.95 pp. Thus, experience with the setting at hand—running field experiments and especially nudge treatments—significantly increases the accuracy in predictions. The fact that expertise improves prediction, while intuitive, is not obvious: for example, DellaVigna and Pope (2018) found that experience with MTurk experiments did not improve the accuracy of prediction of the results of an MTurk experiment. Further, this result was not obvious, as, to the best of our knowledge, the Nudge Unit practitioners did not have an in-house systematic estimate prior to our study.

This result raises a next question: are nudge practitioners more knowledgeable about all estimated nudge impacts? As Figure A11a shows, nudge practitioners actually make a biased forecast for the sample of Academic Journal nudges, with a median prediction of 3.3 pp., compared to the finding of 8.7 pp. impact. One interpretation of these findings is that each group (over-)extrapolates based on the setting they most observe: researchers are quite aware of

the Academic Journal nudge papers, but over-extrapolate for the Nudge Unit results, possibly because they assume that there is less publication bias in academic journals than there actually is. Conversely, the nudge practitioners are focused on the trials they run, for which they have an approximately correct estimate, and they may not pay as much attention to the results in the Academic Journal papers.

We consider one last issue. Are the respondents able to predict *which* treatments will have a larger impact? This is a relevant question, as researchers are implicitly using predictions to decide which treatments to run. The respondents make predictions for three (randomly drawn) interventions, after seeing some detail of the nudge (including visual images of the letter/email/nudge when possible). In Figure A12a, we plot the median forecasted effect size against the estimated treatment effect for each of the 14 treatments used as examples. The median prediction is correlated with the actual effect size, but the correlation is not statistically significant at traditional significance levels ($t=1.39$). This correlation is approximately the same both for experienced and inexperienced predictors (Figure A12b). Predictions on a larger sample of trials will be necessary to conclusively address this issue.

A.7 Mixture of Normals Meta-Analysis with Publication Bias

Consider a population of trials i with base trial effects β_i drawn from Normal 1 $\sim N(\bar{\beta}_1, \tau_{BT1}^2)$ with probability $q \equiv Pr(\text{Normal 1})$, and from Normal 2 $\sim N(\bar{\beta}_2, \tau_{BT2}^2)$ with probability $1 - q$. The between-trial variance in base effects is τ_{BT}^2 , which can differ for Normal 1 and for Normal 2, and the grand average treatment effect is $q\bar{\beta}_1 + (1 - q)\bar{\beta}_2$.

Trials can have multiple arms indexed by j , and each treatment has a true effect β_{ij} centered around the base trial effect β_i . In particular, β_{ij} is drawn from $N(\beta_i, \tau_{WI}^2)$, where τ_{WI}^2 is the within-trial variance in true effects. Furthermore, τ_{WI}^2 can differ depending on whether the base trial effect β_i is drawn from Normal 1 or Normal 2 (i.e., there are separate τ_{WI1} and τ_{WI2}). Lastly, each treatment arm has some level of precision given by an independent standard error σ_{ij} . The final treatment effect observed by the researcher is $\hat{\beta}_{ij} \sim N(\beta_{ij}, \sigma_{ij}^2)$.

To correct for selective publication, we follow Andrews and Kasy (2019)¹¹ that identifies the extent of publication bias in a sample of published studies, and produces bias-corrected parameters for the underlying distribution of true effect sizes. In our case, we assume that the publication decision depends on the highest t -stat among the treatments. That is,

$$Pr(\text{Publish}_i) = \begin{cases} 1 & \text{if } \max_j(\hat{\beta}_{ij}/\sigma_{ij}) \geq 1.96 \\ \gamma & \text{otherwise} \end{cases}$$

The probability of publishing insignificant trials is identified up to scale, i.e., relative to the probability of publishing significant trials.

This model is estimated via maximum likelihood, where the likelihood of trial i is:

$$\mathcal{L}_i(\hat{\beta}_{i1}, \dots, \hat{\beta}_{iK}, \sigma_{i1}, \dots, \sigma_{iK}, |\bar{\beta}, \tau_{BT}, \tau_{WI}, q, \gamma) = \frac{1 - (1 - \gamma)\mathbf{1}\{\max_j(\hat{\beta}_{ij}/\sigma_{ij}) < 1.96\}}{E[1 - (1 - \gamma)\mathbf{1}\{\max_j(\hat{\beta}_{ij}/\sigma_{ij}) < 1.96\}]} \mathbf{f}_{N(\bar{\beta}, \Sigma, q)}$$

where K is the number of treatment arms j in trial i , and $\mathbf{f}_{N(\bar{\beta}, \Sigma, q)}(\hat{\beta}_{i1}, \dots, \hat{\beta}_{iK})$ is the den-

¹¹We thank Andrews and Kasy for their comments in helping us adapt their model to our setting.

sity of the mixture of two normals under the parameters $\bar{\beta} = (\bar{\beta}_1, \bar{\beta}_2)$, $\tau_{BT} = (\tau_{BT1}, \tau_{BT2})$, $\tau_{WI} = (\tau_{WI1}, \tau_{WI2})$ and q . The estimates of $\bar{\beta}_1, \bar{\beta}_2, \tau_{BT1}, \tau_{BT2}, \tau_{WI1}, \tau_{WI2}, q, \gamma$ from this procedure back out the latent distribution of effects before any selective publication.

Extension. As an alternative approach, we present here the results (in Tables A9a-c) under the assumption that the Academic Journals trials and the Nudge Unit trials are drawn from the same underlying distribution of results, modeled with a mixture of 3 normals, but the two sets of trials are drawn with a different probability from the higher normals. The third normal distribution, Normal 3 $\sim N(\bar{\beta}_3, \tau_{BT3}^2)$, also has its own within-trial variance. Now the grand average treatment effect is $q_1\bar{\beta}_1 + q_2\bar{\beta}_2 + q_3\bar{\beta}_3$, where $q_1 + q_2 + q_3 = 1$ and q_m is the probability of drawing a trial base effect from the m -th normal. The likelihood function is the analog of the mixture of two normals version.

The results in Tables A9b-c differ from those in Panel C of Table A9a for two reasons. First, in Table A9a, the mixture of three normals model is estimated on the Academic Journals and Nudge Units samples separately; in Tables A9b-c, it is instead estimated on the stacked data set combining both samples. The latter assumes that the parameters of the three normals ($\bar{\beta}_1, \bar{\beta}_2, \bar{\beta}_3, \tau_{BT1}, \tau_{BT2}, \tau_{BT3}, \tau_{WI1}, \tau_{WI2}, \tau_{WI3}$) are the same for both samples.

Second, in Tables 9b-c, the probabilities of drawing from each of the normals (q_1, q_2, q_3) are estimated under an ordinal probit framework. Specifically, the probability that a trial i draws its effect size from the first (lowest) normal is $P(X_i'\beta + \varepsilon < \theta_1)$, where X_i is a $k \times 1$ vector of trial characteristics, such as being in the Academic Journals sample. β is a $k \times 1$ vector of coefficients, and the error ε follows a standard normal distribution. The probability that a trial i draws its effect size from the second (middle) normal is $P(\theta_1 \leq X_i'\beta + \varepsilon < \theta_2)$, and the probability of drawing from the third (highest) normal is $P(\theta_3 \leq X_i'\beta + \varepsilon)$. The thresholds θ_1, θ_2 and the coefficient vector β are jointly estimated.

Similar to our benchmark estimates of Panel B in Table 5, we estimate in Table A9b a high degree of selective publication $\gamma_{AJ} = 0.07$ (s.e. 0.06) and an ATE for the Academic Journals sample at 2.72 pp. (s.e.=1.06), again suggesting a somewhat larger impact than the Nudge Unit trials. In Table A9c we reproduce this result in Column 2 and then further generalize the set of predictors X to include some of the most common categories of nudges for both samples. We largely replicate qualitatively the findings of Table 4, such as the fact that in-person nudges are more likely to be associated with higher normals, and thus to yield higher effect sizes.

A.8 Additional Meta-analysis models (without selective publication)

In Table A8 we consider additional meta-analyses models, all with the feature that they do *not* model selective publication: (1) DerSimonian and Laird (1986) (DL), (2) empirical Bayes (Paule and Mandel, 1989), (3) (restricted) maximum likelihood; (4) the method from Card, Kluve, and Weber (2018).

The DL method uses the statistic $Q = \sum_i \frac{1}{\sigma_i^2} (\beta_i - \tilde{\beta})^2$, where β_i is the effect size for study i , σ_i is the standard error, and $\tilde{\beta} = \frac{\sum_i (\beta_i / \sigma_i^2)}{\sum_i (1 / \sigma_i^2)}$ is the weighted average using inverse-sampling variance weights. Under random-effects assumptions, the expectation of Q is:

$$E[Q] = (n - 1) + \left(\sum_i (1 / \sigma_i^2) - \frac{\sum_i (1 / \sigma_i^2)^2}{\sum_i (1 / \sigma_i^2)} \right) \tau^2$$

where n is the number of studies in the sample. Solving this equation for the between-study variance results in $\tau_{DL}^2 = \max \left\{ 0, \frac{E[Q] - (n-1)}{\sum_i w_i - \frac{\sum_i w_i^2}{\sum_i w_i}} \right\}$, from which the sample estimates for σ_i and β_i can be plugged in for estimation.

The empirical Bayes and (restricted) maximum likelihood methods assume that each study draws its true effect from some normal distribution $N(\tilde{\beta}, \tau^2)$. The empirical Bayes procedure can be derived using the generalized Q -statistic, which takes the form:

$$Q = \sum_i W_i (\beta_i - \tilde{\beta})^2,$$

$$W_i = \frac{1}{\tau^2 + \sigma_i^2}, \tilde{\beta} = \frac{\sum_i W_i \beta_i}{\sum_i W_i}$$

Under the normal distributional assumption, the expected value of Q equals $n-1$. The empirical Bayes procedure iteratively estimates τ_{EB}^2 using a derivation of the equation

$$\sum_i W_i (\beta_i - \tilde{\beta})^2 = n - 1$$

The (restricted) ML method maximizes the likelihood function

$$L(\hat{\beta}, \hat{\sigma} | \bar{\beta}, \tau^2) = \prod_i \phi \left(\frac{\hat{\beta}_i - \bar{\beta}}{\sqrt{\tau^2 + \hat{\sigma}_i^2}} \right)$$

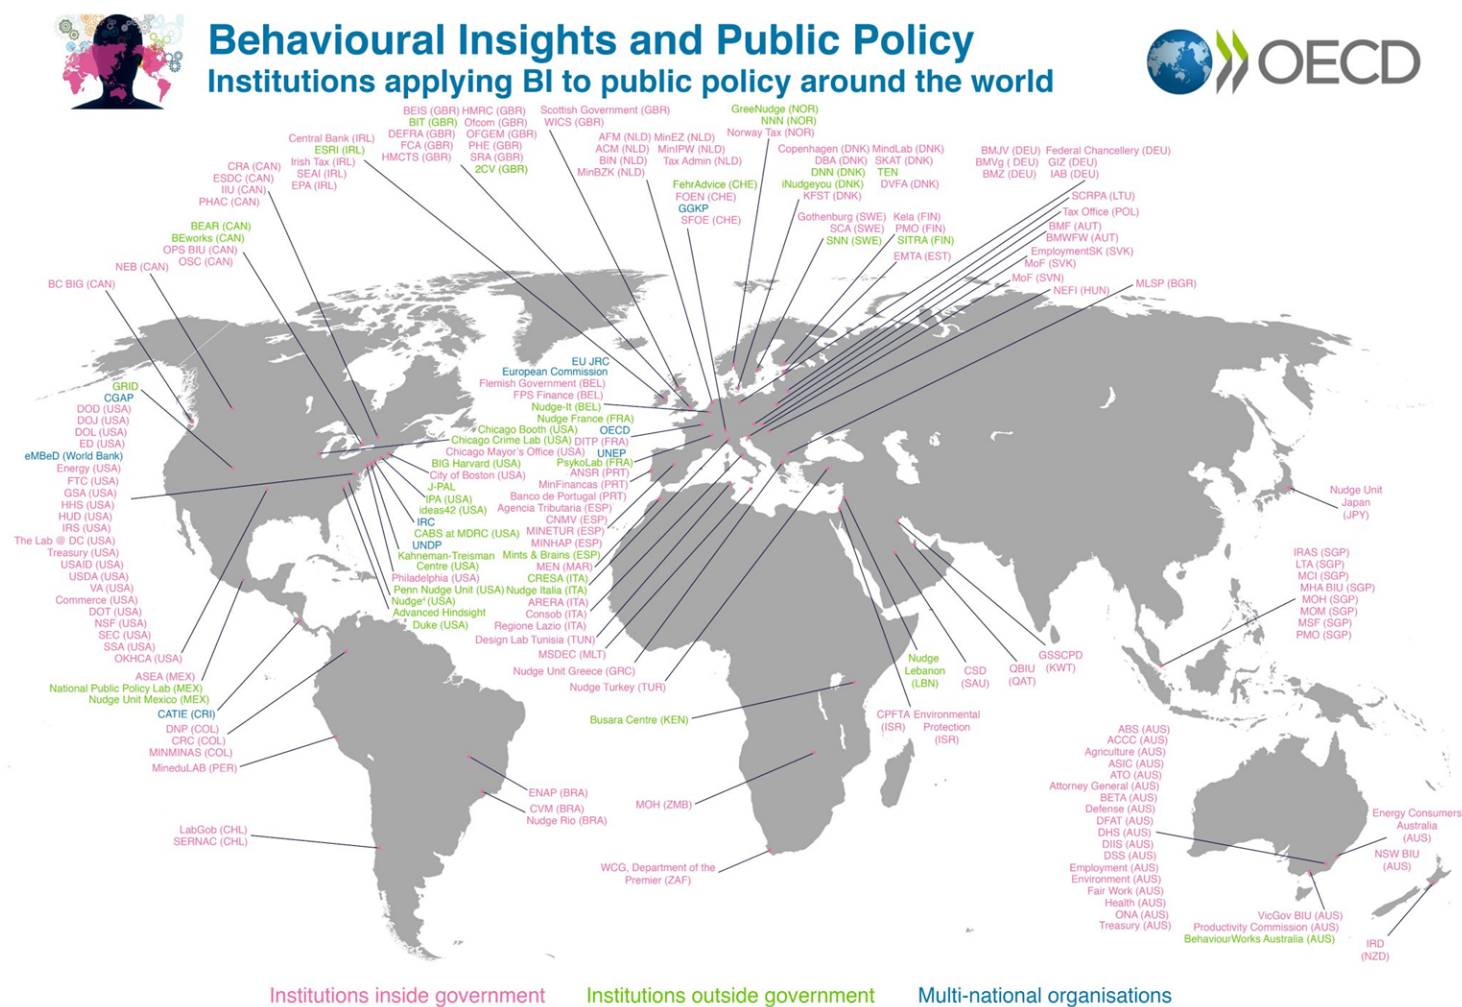
where ϕ is the standard normal density.

The Card, Kluve, and Weber (2018) method decomposes the two random-effects components of variance via linear regression. Regressing the squares of the effect sizes around the (weighted) mean on a constant and the inverse of the effective sample size N_i separates the between-study variance (coefficient on the constant) and the variation attributable to sampling error (coefficient on $1/N_i$). The procedure is conducted in the following steps: Take demeaned effect sizes and square them to obtain $(\beta_i - \bar{\beta})^2$

1. Regress the squared residuals on a constant and the inverse of effective sample size $1/N_i$
2. Re-estimate $\bar{\beta}$ by weighting each effect by $1 / (\hat{\tau}^2 + \hat{k}/N_i)$, where $\hat{\tau}^2$ is the coefficient on the constant and \hat{k} the coefficient on $1/N_i$
3. Iterate steps 1-3 until convergence

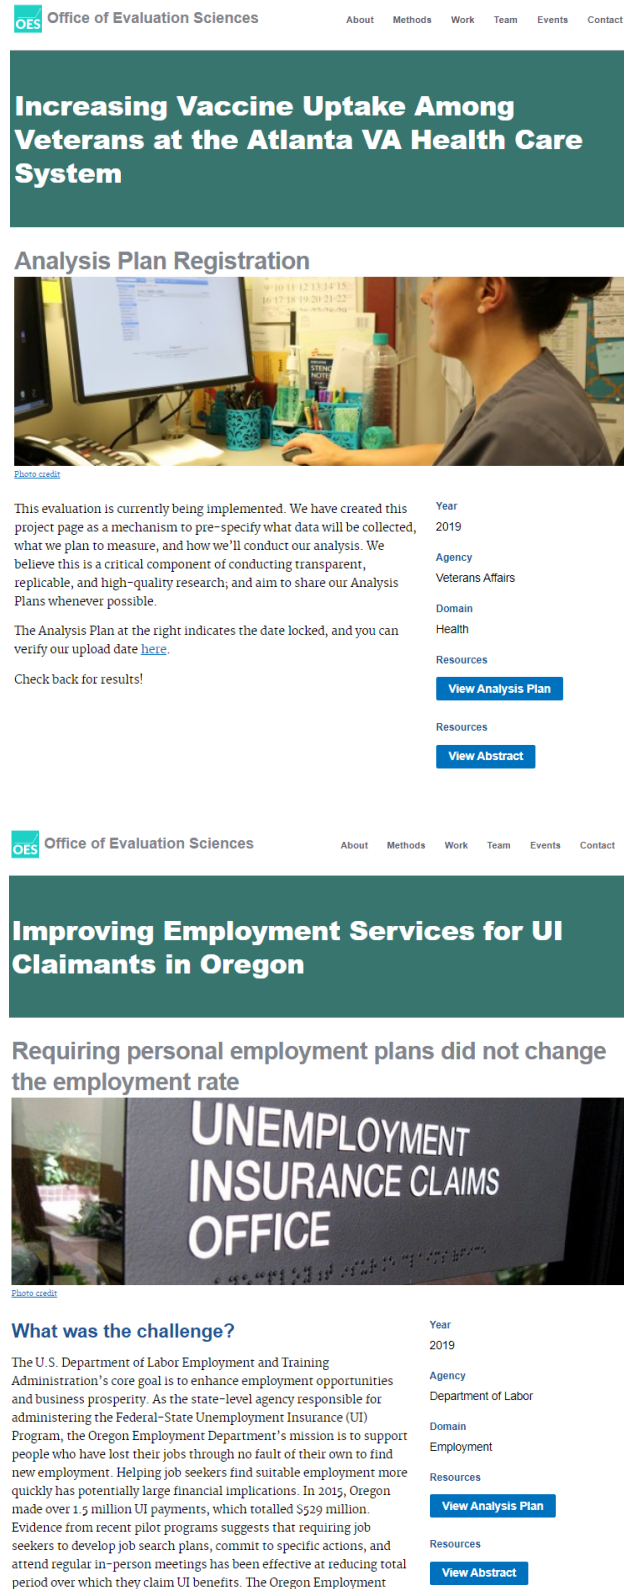
From this iterative variance decomposition, the coefficient on $1/N$ for the Academic Journals sample is 27162.0 (s.e.=12053.1), and the constant is estimated at -3.38 (s.e.=47.13). For the Nudge Units, the estimates are 6362.6 (s.e.=3446.6) and 11.00 (s.e.=6.46) respectively, and for the Published Nudge Units, 576.7 (s.e.=198.5) and 0.647 (s.e.=0.325). The coefficient on the inverse sample size $1/N_i$ is significantly positive as expected.

Figure A1: Nudge Units around the world



This figure shows the various Nudge Units across the world.

Figure A2a: Additional examples of nudges: OES website




This figure shows screen captures directly from the Office of Evaluation Sciences website. The top page documents the analysis plan registration for an ongoing trial, whereas the bottom page presents the trial report from a concluded trial.

Figure A2b: Additional examples of nudges: BIT-NA example




This figure presents an example of a nudge intervention run by BIT-NA. This trial encourages utilities customers to enroll in AutoPay and e-bill using bill inserts. The control group received the status quo utility bill that advertises e-bill and AutoPay on the back, while the treatment group received an additional insert with simplified graphics. The outcome in this trial is measured as AutoPay/e-bill enrollment rates.

Figure A3: Comparison of nudge categories



This figure shows the frequencies of nudges in category of characteristics. Categories for Medium and Mechanism are not mutually exclusive and frequencies may not sum to 1.


Figure A4: Comparison of trial features between Nudge Units and Academic Journals



This figure compares the distribution of nudge-by-nudge treatment arm sample sizes (i.e. excluding the control group sample size) between the Nudge Units and the Academic Journals samples.

Figure A4: Comparison of trial features between Nudge Units and Academic Journals

(b) Treatment arms per trial



(c) Mechanisms per treatment arm

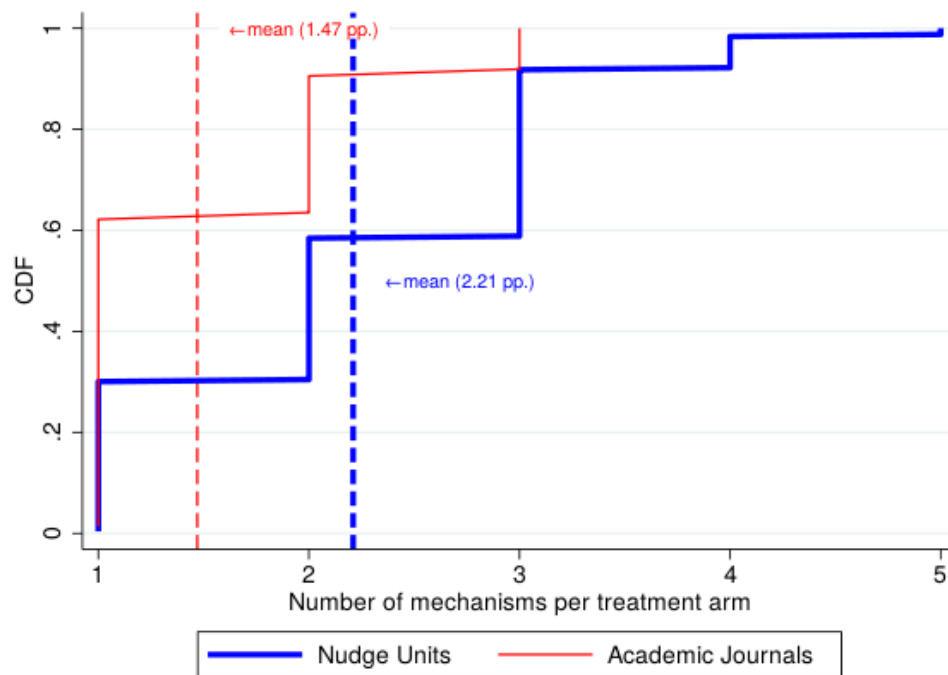




Figure A5: Publication bias tests: Point estimate and standard error




This figure compares the nudge-by-nudge relationship between the standard error and the treatment effect for the Academic Journals sample (A5a) versus the Nudge Units sample (A5b). The estimated equation is the linear fit with standard errors clustered at the trial level.

Figure A5: Publication bias tests: Andrews-Kasy funnel plot


(c) Academic Journals: All nudges




(e) Nudge Units: All nudges



(d) Academic Journals: Most significant nudges by trial




(f) Nudge Units: Most significant nudges by trial




This figure plots the nudge-by-nudge treatment effect (horizontal axis) against the standard error (vertical axis). Nudges within the two gray lines are insignificant at the 5% level (i.e., $|t| < 1.96$). Figures A5c and A5e show all the nudges in the samples, while A5d and A5f show only the nudges with the highest t -stat within each trial. 1 trial in the Academic Journals sample and 2 trials from the Nudge Units sample in which the most significant treatment uses defaults/financial incentives are excluded from A5d and A5f respectively.

Figure A5: Publication bias tests: t -stat distribution (bin-width ≈ 0.15)


(g) Academic Journals: All nudges




(i) Nudge Units: All nudges



(h) Academic Journals: Most significant nudges by trial

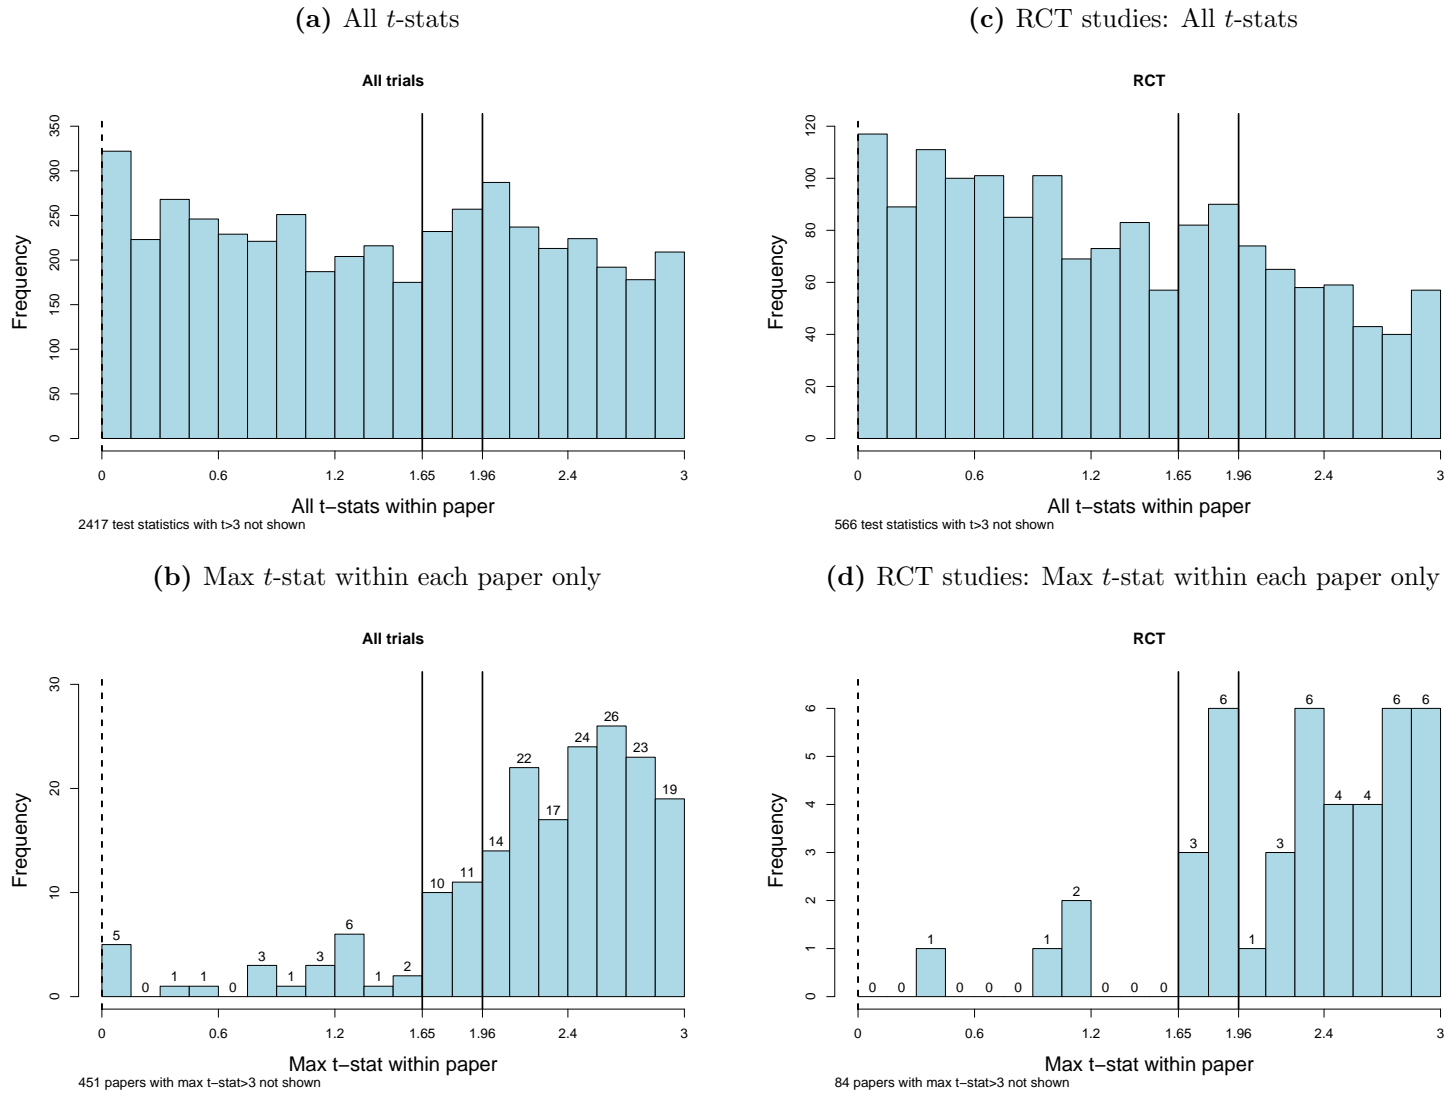


(j) Nudge Units: Most significant nudges by trial



1 trial in the Academic Journals sample and 2 trials from the Nudge Units sample in which the most significant treatment uses defaults/financial incentives are excluded from A5h and A5j respectively.


Figure A6: Distribution of t -stats from Brodeur, Cook, and Heyes (2020)




We thank Abel Brodeur for promptly sharing the data for this analysis. Brodeur et al. (2020) gather this data from the universe of papers published in the top 25 economics journals in 2015 and 2018. They categorize papers by empirical method (DID, IV, RCT, and RDD) and record the point estimate and standard error from the results in the main table of each article. Figure A6a shows the distribution of all the t -stats from the main table of each paper for the entire sample of articles, while Figure A6b shows the distribution of only the maximum t -stat within each paper. Figures A6c and A6d show the analog for the subsample of RCT papers.

Figure A7: Academic Journals: Comparison of meta-analysis models

(a) Normal-based meta-analysis vs. mixture of two normals




(b) With and without publication bias correction



This figure plots the empirical and estimated distribution of observed nudge effects and compares various meta-analysis specifications from Tables 5 and A9a. Figure A7a compares the fit of a normal-based meta-analysis model and that of a mixture of two normals model. These two models also include a correction for publication bias in Figure A7b. 3 nudges with effects greater than 35 pp. are not shown. The densities are kernel approximations from 500,000 simulated trials.

Figure A8: Within-collaboration Nudge Unit effects

(a) Nudge Unit treatment effects in early vs. later collaborations with the same agency/city



(b) Success of first collaboration and number of total collaborations with the same agency/city






Figure A8a compares the CDF of the treatment effects in percentage points between the first half of trials (“early”) in a series of collaborations with the same government agency or city and the latter half of trials in the same series of collaborations (“latter”). Trials that were one-time collaborations with an agency or city are not included. When there is an odd number of trials in a collaboration, the median trial is not included. Figure A8b categorizes the first trials in each series of collaborations with a partnering government agency or city (which may be one-time) into deciles based on the treatment effect of their most effective arm. This figure shows the average total number of collaborations for each decile. The labels for each point reports the range of treatment effect sizes in each decile.

Figure A9: Publication bias tests: Published Nudge Unit trials


(a) Point estimate and minimum detectable effect



(b) t -stat distribution




(c) Most significant nudges by trial




This panel displays tests for publication bias in the Nudge Units sample. Figure A9a plots the relationship between the minimum detectable effect and the treatment effect size. The estimated equation is the linear fit with standard errors clustered at the trial level. Figure A9b shows the distribution of t -statistics (i.e., treatment effect divided by standard error) for all nudges, and Figure A9c shows the distribution for only the max t -stat within each trial.

Figure A9: Publication bias tests: Published Nudge Unit trials

(d) Andrews-Kasy funnel plot




(e) Andrews-Kasy funnel plot: Most significant treatments



This figure plots the nudge-by-nudge treatment effect (horizontal axis) against the standard error (vertical axis). Nudges within the two gray lines are insignificant at the 5% level (i.e., $t < 1.96$). Figure A9d shows all the nudges in the Published Nudge Units sample, while A9e shows only the nudges with the highest t -stat within their trial.


Figure A10: Characteristics of forecasters




This figure shows the characteristics of the forecasters along several dimensions. Figure A10a categorizes forecasters by their professional affiliation, A10b by their academic background (if they are university faculty/(under)graduate students), and A10c by their experience in conducting field experiments.

Figure A11: Findings vs. expert forecasts

(a) Forecasts for Academic Journals by forecaster experience




(b) Forecasts for Nudge Units by forecaster experience



Figures A11a and A11b show the distributions of forecasts for treatment effects in the Academic Journals and Nudge Units samples respectively, separated by the forecasters' experience in running field experiments.


Figure A12: Example-by-example forecasts

(a) All respondents



14 examples. Numeric labels are the number of forecasts for each example.
45 degree dashed line shown.

(b) Forecasts by forecaster experience



Experienced respondents: >5 field experiments experience/nudge practitioners.
14 examples. Numeric labels are the number of forecasts for each example.
45 degree dashed line shown.

This figure plots the median forecasted treatment effect for each of the 14 examples shown on the forecast survey against the true treatment effect. Figure A12a presents forecasts from all the respondents, and A12b splits the forecasts by experience.

Table A1a: List of published papers in the Nudge Units sample

Published papers featuring OES trials

1. Anteneh et al. 2020. "Appraising praise: experimental evidence on positive framing and demand for health services." *Applied Economics Letters*. Cited by 0 (Insignificant)
2. Benartzi et al. 2017. "Should Governments Invest More in Nudging?" *Psychological Science*, 28(8): 1041-1055. Cited by 281
3. Bowers et al. 2017. "Challenges to Replication and Iteration in Field Experiments: Evidence from Two Direct Mail Shots." *American Economic Review, Papers and Proceedings*, 107(5): 462-65. Cited by 0
4. Castleman and Page. 2017. "Parental influences on postsecondary decision-making: Evidence from a text messaging experiment." *Educational Evaluation and Policy Analysis*, 39(2): 361-77. Cited by 26
5. Chen et al. forthcoming. "The Effect of Postcard Reminders on Vaccinations Among the Elderly: A Block-Randomized Experiment." *Behavioural Public Policy*. Cited by 0
6. Guyton et al. 2017. "Reminders and Recidivism: Using Administrative Data to Characterize Nonfilers and Conduct EITC Outreach." *American Economic Review, Papers & Proceedings*, 107(5): 471-75. Cited by 8
7. Leight and Safran. 2019. "Increasing immunization compliance among schools and day care centers: Evidence from a randomized controlled trial." *Journal of Behavioral Public Administration*, 2(2). Cited by 2 (Insignificant)
8. Leight and Wilson. 2019. "Framing Flexible Spending Accounts: A Large-Scale Field Experiment on Communicating the Return on Medical Savings Accounts." *Health Economics*, 29(2): 195-208. Cited by 0 (Insignificant)
9. Kramer and Cooper. 2020. Paper based on trial "Using Proactive Communication to Increase College Enrollment for Post-9/11 GI Bill Beneficiaries", R&R at *Education Finance and Policy*.
10. Sacarny, Barnett, and Le. 2018. "Effect of Peer Comparison Letters for High-Volume Primary Care Prescribers of Quetiapine in Older and Disabled Adults." *JAMA Psychiatry*, 75(10): 1003-1011. Cited by 21
11. Yokum et al. 2018. "Letters designed with behavioural science increase influenza vaccination in Medicare beneficiaries." *Nature Human Behaviour*, 2: 743-749. Cited by 5

Published papers featuring BIT NA trials

1. Linos. 2017. "More Than Public Service: A Field Experiment on Job Advertisements and Diversity in the Police." *Journal of Public Administration Research and Theory*, 28(1): 67-85. Cited by 25
2. Linos, Ruffini, and Wilcoxon. 2019. "Belonging Affirmation Reduces Employee Burnout and Resignations in Front Line Workers." Working paper. Cited by 0
3. Linos, Quan, and Kirkman. 2020. "Nudging Early Reduces Administrative Burden: Three Field Experiments to Improve Code Enforcement." *Journal of Policy Analysis and Management*, 39(1): 243-265. (covers 3 trials) Cited by 0 (2/3 trials are insignificant)

Table A1b: List of papers in the Academic Journals sample

1. Altmann and Traxler. 2014. "Nudges at the Dentist." *European Economic Review*, 11(3): 634-660. Cited by 69
2. Apesteguia, Funk, and Iriberry. 2013. "Promoting Rule Compliance in Daily-Life: Evidence from a Randomized Field Experiment in the Public Libraries of Barcelona." *European Economic Review*, 63(1): 66-72. Cited by 36
3. Bartke, Friedl, Gelhaar, and Reh. 2016. "Social Comparison Nudges—Guessing the Norm Increases Charitable Giving." *Economics Letters*, 67: 8-13. Cited by 16
4. Bettinger and Baker. 2011. "The Effects of Student Coaching in College: An Evaluation of a Randomized Experiment in Student Mentoring." *Educ. Eval. & Policy Analysis*, 33: 433-461. Cited by 31
5. Bettinger, Long, Oreopoulos, and Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H & R Block FAFSA Experiment." *Quarterly Journal of Economics*, 8(10): e77055. Cited by 780

6. Carroll, Choi, Laibson, Madrian, and Metrick. 2009. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics*, 53(5): 829-846. [Cited by 581](#)
7. Castleman and Page. 2015. "Summer Nudging: Can Personalized Text Messages and Peer Mentor." *Journal of Economic Behavior and Organization*, 16(1): 15-22. [Cited by 273](#)
8. Chapman et al.. 2010. "Opting in Vs. Opting out of Influenza Vaccination." *Journal of the American Medical Association*, 76: 89-97. [Cited by 135](#)
9. Cohen et al.. 2015. "Effects of Choice Architecture and Chef-Enhanced Meals on the Selection and Consumption of Healthier School Foods: A Randomized Clinical Trial." *JAMA Pediatrics*, 124(4): 1639-1674. [Cited by 77](#)
10. Damgaard and Gravert. 2016. "The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising." *Journal of Public Economics*, 121(556): F476-F493. [Cited by 66 \(Insignificant\)](#)
11. Fellner, Sausgruber, and Traxler. 2013. "Testing Enforcement Strategies in the Field: Appeal, Moral Information, Social Information." *Journal of the European Economic Association*, 108(26): 10415-10420. [Cited by 285](#)
12. Gallus. 2016. "Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia." *Management Science*, 115: 144-160. [Cited by 68](#)
13. Goswami and Urminsky. 2016. "When Should the Ask Be a Nudge? The Effect of Default Amounts on Charitable Donations." *Journal of Marketing Research*, 60(573): e137-43. [Cited by 57](#)
14. Holt, Thorogood, Griffiths, Munday, Friede, and Stables. 2010. "Automated electronic reminders to facilitate primary cardiovascular disease prevention: randomised controlled trial." *British Journal of General Practice*, 152: 73-75. [Cited by 35](#)
15. Kristensson, Wästlund, and Söderlund. 2017. "Influencing Consumers to Choose Environment Friendly Offerings: Evidence from Field Experiments." *Journal of Business Research*, 304(1): 43-44. [Cited by 22](#)
16. Lehmann, Chapman, Franssen, Kok, and Ruiter. 2016. "Changing the default to promote influenza vaccination among health care workers." *Vaccine*, 36(1): 3-19. [Cited by 22](#)
17. Löfgren, Martinsson, Hennlock, and Sterner. 2012. "Are Experienced People Affected by a Pre-Set Default Option—Results from a Field Experiment." *Journal of Env. Econ. & Mgmt.*, 64: 266-284. [Cited by 69 \(Insignificant\)](#)
18. Luoto, Levine, Albert, and Luby. 2014. "Nudging to Use: Achieving Safe Water Behaviors in Kenya and Bangladesh." *Journal of Development Economics*, 63(12): 3999-4446. [Cited by 30](#)
19. Malone, and Lusk. 2017. "The Excessive Choice Effect Meets the Market: A Field Experiment on Craft Beer Choice." *Journal of Behav. & Exp. Econ.*, 129: 42-44. [Cited by 13](#)
20. Miesler, Scherrer, Seiler, and Bearth. 2017. "Informational Nudges As An Effective Approach in Raising Awareness among Young Adults about the Risk of Future Disability." *Journal of Consumer Behavior*, 169(5): 431-437. [Cited by 7](#)
21. Milkman, Beshears, Choi, Laibson, and Madrian. 2011. "Using Implementation Intentions Prompts to Enhance Influenza Vaccination Rates." *PNAS*, 34(11): 1389-92. [Cited by 297](#)
22. Nickerson, and Rogers. 2010. "Do You Have a Voting Plan? Implementation Intentions, Voter Turnout, and Organic Plan Making." *Psychological Science*, 127(3): 1205-1242. [Cited by 243](#)
23. Rodriguez-Priego, Van Bavel, and Monteleone. 2016. "The Disconnection Between Privacy Notices and Information Disclosure: An Online Experiment." *Economia Politica*, 21(2): 194-199. [Cited by 4](#)
24. Rommela, Vera Buttmannb, Georg Liebig, Stephanie Schönwetter, and Valeria Svart-Gröger. 2015. "Motivation Crowding Theory and Pro-Environmental Behavior: Experimental Evidence." *Economics Letters*, 157: 15-26. [Cited by 14](#)
25. Stutzer, Goette, and Zehnder. 2011. "Active Decisions and Prosocial Behaviour: A Field Experiment on Blood Donation." *Economic Journal*, 72: 19-38. [Cited by 65 \(Insignificant\)](#)
26. Wansink and Hanks. 2013. "Slim by Design: Serving Healthy Foods First in Buffet Lines Improves Overall Meal Selection." *PLoS ONE*, 110: 13-21. [Cited by 93](#)

[Citations](#) are updated as of March 5, 2020. The "[\(Insignificant\)](#)" label applies to papers that have no nudge treatment arms with a *t*-stat above 1.96.

Table A2: Comparison of nudge categories

| | Nudge Units | | | Academic Journals | | |
|--------------------------------------|-------------|---------------------|------------------|-------------------|---------------------|-----------------|
| | Freq. (%) | Control take-up (%) | Trial-level N | Freq. (%) | Control take-up (%) | Trial-level N |
| <i>Date</i> | | | | | | |
| Early* | 46.50 | 13.78 | 191,571 | 48.65 | 25.34 | 24,208 |
| Recent* | 53.50 | 20.06 | 142,634 | 51.35 | 26.58 | 5,518 |
| <i>Policy area</i> | | | | | | |
| Revenue & debt | 28.81 | 11.90 | 151,075 | 17.57 | 10.98 | 23,380 |
| Benefits & programs | 22.22 | 17.37 | 381,021 | 10.81 | 27.66 | 4,312 |
| Workforce & education | 18.52 | 14.39 | 134,726 | 9.46 | 66.16 | 3,950 |
| Health | 13.17 | 18.31 | 81,810 | 28.38 | 24.57 | 4,854 |
| Registration & regulation compliance | 8.64 | 45.41 | 7,981 | 12.16 | 14.42 | 8,917 |
| Community engagement | 7.82 | 8.77 | 196,286 | 4.05 | 40.27 | 135,912 |
| Environment | 0.82 | 23.37 | 9,478 | 13.51 | 28.20 | 419 |
| Consumer behavior | 0 | – | 0 | 4.05 | 15.43 | 7,253 |
| <i>Medium of communication</i> | | | | | | |
| Email | 39.51 | 13.03 | 205,076 | 12.16 | 21.06 | 17,962 |
| Physical letter | 29.63 | 26.05 | 184,759 | 16.22 | 13.17 | 14,911 |
| Postcard | 21.40 | 15.39 | 122,838 | 6.76 | 8.90 | 1,227 |
| Website | 2.88 | 9.85 | 22,822 | 12.16 | 10.83 | 2,492 |
| In person | 0.82 | 27.50 | 4,242 | 28.38 | 35.40 | 2,299 |
| Other | 11.11 | 20.65 | 114,979 | 24.32 | 38.28 | 26,304 |
| <i>Control group receives:</i> | | | | | | |
| No communication | 60.91 | 15.14 | 230,798 | 43.24 | 29.51 | 25,709 |
| Some communication | 39.09 | 20.37 | 83,508 | 56.76 | 23.28 | 8,149 |
| <i>Mechanism</i> | | | | | | |
| Simplification & information | 58.02 | 17.23 | 217,529 | 5.41 | 24.08 | 4,057 |
| Personal motivation | 57.20 | 15.81 | 205,815 | 32.43 | 30.97 | 4,347 |
| Reminders & planning prompts | 31.28 | 27.13 | 160,849 | 35.14 | 25.17 | 26,246 |
| Social cues | 37.04 | 17.19 | 95,540 | 21.62 | 31.11 | 8,230 |
| Framing & formatting | 31.69 | 12.74 | 205,766 | 32.43 | 23.78 | 1,614 |
| Choice design | 6.17 | 14.05 | 334,554 | 20.27 | 23.60 | 2,723 |
| Total | 100 | 17.20 | 23,577,537 (sum) | 100 | 25.97 | 505,337 (sum) |

This table shows the frequency of nudges in each category, and the average control group take-up and trial-level N within each category. Frequencies for *Medium* and *Mechanism* are not mutually exclusive and frequencies may not sum to 1.

**Early* refers to trials implemented between 2015-2016 for Nudge Units, and to papers published in 2014 or before for Academic Journals. *Recent* refers to trials and papers after these dates.

Table A3a: Unweighted treatment effects in log odds ratio

| | Academic Journals | Nudge Units | | | |
|---|-------------------|-------------------|-------------------|------------------|--------------------------------|
| | (1) | All (2) | BIT (3) | OES (4) | Academic-affiliated OES (5) |
| Average treatment effect (log odds ratio) | 0.499 (0.110) | 0.267 (0.0667) | 0.245 (0.0713) | 0.292 (0.120) | 0.339 (0.265) |
| Nudges | 74 | 231 | 125 | 106 | 44 |
| Trials | 26 | 121 | 75 | 46 | 23 |
| Observations | 505,337 | 23,391,985 | 1,935,014 | 21,456,971 | 8,919,795 |
| Average control group take-up (%) | 25.97 | 17.79 | 16.37 | 19.47 | 26.45 |
| <i>Distribution of treatment effects</i> | | | | | |
| 25th percentile | 0.12 | 0.01 | -0.00 | 0.02 | 0.01 |
| 50th percentile | 0.32 | 0.10 | 0.11 | 0.08 | 0.04 |
| 75th percentile | 0.69 | 0.34 | 0.47 | 0.23 | 0.17 |

This table shows the average treatment effect of nudges. Standard errors clustered by trial are shown in parentheses.

Table A3b: Unweighted treatment effects for Published Nudge Unit trials

| | Percentage points (1) | Log odds ratio (2) |
|--|--------------------------|-----------------------|
| Average treatment effect | 0.970 (0.234) | 0.202 (0.0981) |
| Nudges | 33 | 33 |
| Trials | 16 | 16 |
| Observations | 2,136,014 | 2,136,014 |
| Average control group take-up (%) | 31.93 | 31.93 |
| <i>Distribution of treatment effects</i> | | |
| 25th percentile | 0.20 | 0.02 |
| 50th percentile | 0.50 | 0.05 |
| 75th percentile | 1.20 | 0.14 |

This table shows the average treatment effect of nudges. Standard errors clustered by trial are shown in parentheses.

Table A4a: Categorization of treatment effects

| | Academic Journals | | Nudge Units | |
|--------------------------|-------------------|-----------|-------------|-----------|
| | Nudges | Freq. (%) | Nudges | Freq. (%) |
| Significant & positive | 40 | 54.05 | 116 | 47.74 |
| Insignificant & positive | 28 | 37.84 | 79 | 32.51 |
| Insignificant & negative | 6 | 8.11 | 34 | 13.99 |
| Significant & negative | 0 | 0 | 14 | 5.76 |
| Total | 74 | 100 | 243 | 100 |

Significance is determined at the 95% level.

Table A4b: Robustness checks

| | Academic Journals | Nudge Units | Published Nudge Units |
|---|-------------------|----------------|-----------------------|
| | (1) | (2) | (3) |
| Average treatment effect (pp.) | 8.68 (2.47) | 1.37 (0.30) | 0.97 (0.23) |
| <i>Panel A. ATE including:</i> | | | |
| Defaults | 9.57 (2.60) | 1.45 (0.31) | 1.09 (0.26) |
| Most policy relevant | 6.47 (1.73) | 1.55 (0.47) | 1.00 (0.24) |
| Low cost interventions | — | 1.35 (0.36) | 1.87 (0.67) |
| <i>Panel B. ATE weighted by:</i> | | | |
| Citations | 7.89 (2.01) | — | 0.76 (0.14) |
| asinh(citations) | 8.25 (2.19) | — | 0.92 (0.19) |
| Nudges | 74 | 243 | 33 |
| Trials | 26 | 126 | 16 |
| Observations | 505,337 | 23,852,753 | 2,335,924 |

This table shows the average treatment effects including default nudges, only the outcomes in the top half of policy relevance, or only nudges with low cost interventions, and weighting treatment effects by citations. Standard errors clustered by trial are shown in parentheses. The Nudge Units sample has 2 nudges (from 1 trial) that use defaults on 1.3 million participants and have treatment effects in p.p. (standard errors) of 9.4 (0.15) and 11.2 (0.15). The Academic Journals sample has 3 nudges (from 3 trials) that use defaults on 548 participants and have treatment effects in pp. (standard errors) of -0.1 (3.6), 3.9 (7.78), and 91 (2.87). Policy relevance is determined by priority scores in response to the question: *How much of a priority is this outcome to its policy area?* Seven undergraduate research assistants reported their scores for each trial outcome on a 3-point scale (1-Low, 2-Medium, 3-High). The most policy relevant nudges are defined as those in the top half of average priority scores. For the Academic Journal outcomes, the Cronbach's alpha for the scoring is 0.83, and for the Nudge Unit outcomes, 0.62. 65 percent of Nudge Unit trials are considered low cost interventions, which are either email communications or cases in which the control group was receiving a status quo communication. Citations are updated as of March 5, 2020. Trials with zero citations are assigned a citation count of 1 in the weighting analysis. See Tables A1a and A1b for the list of published trials and their citation counts.

Table A5: Targeted power in MDE calculations from AEA registry trials

| | Number of trials |
|---|------------------|
| (1) All trials in AEA registry as of March 2020 | 3379 |
| (2) Trials registered prior to intervention start date | 1315 |
| (2a) Trials with non-empty MDE field | 555 |
| (2b) Trials specifying targeted power level for MDE calculation | 267 |
| (2c) Trials using a target power level of 0.8 for MDE calculation | 240 |

The trials included in this table were scrapped from the AEA RCT Registry in March 2020. The registry contains an optional field titled “Minimum detectable effect size for main outcomes (accounting for sample design and clustering)”. We use the responses in this field to compile data on targeted power levels in minimum detectable effect size (MDE) calculations for trials that were registered prior to the start of their intervention. Row (2a) includes trials that (i) stated a MDE without specifying the target power level, (ii) referred to a separate document without stating the MDE and its target power level in the MDE field, or (iii) calculated the power based on an expected effect size (instead of calculating the minimum detectable effect size based on a target power level); these trials are excluded in rows (2b) and (2c).

Table A6a: Heterogeneity in effects by nudge categories: Academic Journals

| Dep. Var.: Treatment effect (pp.) | OLS (1) | OLS (2) | OLS (3) | OLS (4) | OLS (5) | OLS (6) | OLS (7) | OLS (8) | Lasso (9) |
|---------------------------------------|------------------|-------------------|------------------|-------------------|-------------------|-------------------|-------------------|---------------------|--------------|
| Min. detectable effect (MDE) | 1.047 (0.303) | | | | | | | -0.820 (0.457) | 0.572 |
| Control take-up % | | 0.706 (0.289) | | | | | | 1.078 (0.332) | |
| Control take-up % ² | | -0.009 (0.004) | | | | | | -0.011 (0.006) | |
| Log(outcome time-frame days) | | -1.676 (0.945) | | | | | | -3.542 (1.432) | |
| <i>Date</i> | | | | | | | | | |
| Recent (published after 2014) | | | 3.086 (4.760) | | | | | 0.295 (3.302) | |
| <i>Policy area</i> | | | | | | | | | |
| Benefits & programs | | | | 10.548 (5.170) | | | | 6.892 (6.455) | |
| Workforce & education | | | | -1.045 (3.483) | | | | -11.559 (11.008) | |
| Health | | | | 5.379 (3.885) | | | | -1.754 (6.904) | |
| Registrations & regulation compliance | | | | -0.447 (3.482) | | | | -22.885 (8.069) | |
| Community engagement | | | | -0.802 (4.039) | | | | -20.176 (9.863) | |
| Environment | | | | 19.352 (7.723) | | | | 1.318 (8.461) | 2.937 |
| Consumer behavior | | | | -0.409 (3.436) | | | | -23.614 (10.004) | |
| <i>Medium of communication</i> | | | | | | | | | |
| Email | | | | | -5.629 (3.683) | | | 9.886 (5.623) | |
| Physical letter | | | | | -7.710 (3.253) | | | -1.022 (4.866) | |
| Postcard | | | | | 1.078 (3.124) | | | 19.467 (7.729) | |
| Website | | | | | -3.144 (4.307) | | | 10.777 (11.767) | |
| In person | | | | | 5.442 (5.331) | | | 3.703 (6.083) | 0.152 |
| <i>Control group receives:</i> | | | | | | | | | |
| Some communication | | | | | | -3.920 (5.319) | | -5.335 (4.553) | |
| <i>Mechanism</i> | | | | | | | | | |
| Simplification & information | | | | | | | 14.333 (4.649) | 13.567 (5.847) | |
| Personal motivation | | | | | | | 0.288 (3.984) | 1.571 (4.114) | |
| Reminders & planning prompts | | | | | | | 0.286 (3.183) | 2.870 (4.388) | |
| Social cues | | | | | | | 9.382 (6.724) | 9.953 (4.639) | |
| Framing & formatting | | | | | | | 8.999 (4.496) | 8.429 (4.363) | |
| Choice design | | | | | | | 3.766 (4.183) | 10.424 (6.037) | |
| Constant | 0.116 (1.935) | 3.720 (4.566) | 7.098 (1.638) | 3.602 (3.436) | 9.382 (3.124) | 10.907 (5.047) | 2.003 (3.679) | 1.106 (7.969) | 3.568 |
| Nudges | 74 | 74 | 74 | 74 | 74 | 74 | 74 | 74 | 74 |
| Trials | 26 | 26 | 26 | 26 | 26 | 26 | 26 | 26 | 26 |
| Observations | 505,337 | 505,337 | 505,337 | 505,337 | 505,337 | 505,337 | 505,337 | 505,337 | 505,337 |
| R-squared | 0.34 | 0.24 | 0.02 | 0.35 | 0.17 | 0.03 | 0.23 | 0.72 | |
| Avg. control take-up | 25.97 | 25.97 | 25.97 | 25.97 | 25.97 | 25.97 | 25.97 | 25.97 | 25.97 |

Standard errors clustered by trial are shown in parentheses. The minimum detectable effect (MDE) is calculated in pp. at power 0.8. The penalty parameter in the linear lasso model is selected with cross-validation.

Table A6b: Heterogeneity in effects by nudge categories: Nudge Units

| Dep. Var.: Treatment effect (pp.) | OLS (1) | OLS (2) | OLS (3) | OLS (4) | OLS (5) | OLS (6) | OLS (7) | OLS (8) | Lasso (9) |
|---------------------------------------|------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|--------------|
| Min. detectable effect (MDE) | 0.210 (0.246) | | | | | | | 0.225 (0.252) | 0.099 |
| Control take-up % | | 0.092 (0.059) | | | | | | -0.002 (0.048) | |
| Control take-up % ² | | -0.001 (0.001) | | | | | | -0.000 (0.001) | |
| Log(outcome time-frame days) | | 0.253 (0.261) | | | | | | 0.264 (0.314) | 0.088 |
| <i>Date</i> | | | | | | | | | |
| Recent (2017-) | | | -0.863 (0.632) | | | | | -0.123 (0.634) | |
| <i>Policy area</i> | | | | | | | | | |
| Benefits & programs | | | | -1.541 (1.003) | | | | -1.230 (0.726) | -0.049 |
| Workforce & education | | | | -1.935 (0.935) | | | | -1.217 (0.822) | -0.134 |
| Health | | | | -1.774 (0.964) | | | | -2.771 (1.197) | -0.502 |
| Registrations & regulation compliance | | | | -0.251 (1.233) | | | | -0.717 (1.445) | |
| Community engagement | | | | -1.685 (1.537) | | | | -1.538 (1.176) | |
| Environment | | | | 4.404 (1.180) | | | | 4.880 (1.852) | 3.195 |
| <i>Medium of communication</i> | | | | | | | | | |
| Email | | | | | -0.211 (0.644) | | | -1.026 (0.807) | |
| Physical letter | | | | | 1.229 (0.806) | | | 1.043 (0.700) | 0.879 |
| Postcard | | | | | -0.682 (0.647) | | | -0.357 (0.698) | |
| Website | | | | | -1.309 (3.372) | | | -0.169 (2.848) | |
| In person | | | | | 1.274 (1.612) | | | 1.275 (2.809) | |
| <i>Control group receives:</i> | | | | | | | | | |
| Some communication | | | | | | -0.118 (0.623) | | -0.290 (0.581) | |
| <i>Mechanism</i> | | | | | | | | | |
| Simplification & information | | | | | | | -0.202 (0.479) | -0.770 (0.684) | |
| Personal motivation | | | | | | | 0.848 (0.511) | 0.942 (0.544) | 0.496 |
| Reminders & planning prompts | | | | | | | 1.356 (0.630) | 1.100 (0.587) | 0.738 |
| Social cues | | | | | | | -0.354 (0.452) | -0.435 (0.610) | -0.074 |
| Framing & formatting | | | | | | | 0.016 (0.584) | -0.423 (0.676) | |
| Choice design | | | | | | | 5.625 (3.030) | 4.861 (2.610) | 4.904 |
| Constant | 1.012 (0.339) | 0.010 (0.809) | 1.837 (0.521) | 2.426 (0.919) | 1.267 (0.547) | 1.421 (0.378) | 0.361 (0.555) | 1.213 (1.715) | -0.012 |
| Nudges | 243 | 243 | 243 | 243 | 243 | 243 | 243 | 243 | 243 |
| Trials | 126 | 126 | 126 | 126 | 126 | 126 | 126 | 126 | 126 |
| Observations | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 | 23,577,537 |
| R-squared | 0.01 | 0.04 | 0.01 | 0.06 | 0.03 | 0.00 | 0.17 | 0.27 | |
| Avg. control take-up | 17.20 | 17.20 | 17.20 | 17.20 | 17.20 | 17.20 | 17.20 | 17.20 | 17.20 |

Standard errors clustered by trial are shown in parentheses. The minimum detectable effect (MDE) is calculated in pp. at power 0.8. The penalty parameter in the linear lasso model is selected with cross-validation. The 4 nudges (2 trials) missing control take-up data are dummied out when including control take-up in the regression.

Table A7: Weighted decomposition between Nudge Units and Academic Journals

| Dep. Var.: Treatment effect (pp.) | (1) Egger's test | (2) | (3) | (4) |
|---|-------------------|------------------|------------------|------------------|
| Academic Journals | -0.284 (0.100) | 1.688 (1.313) | 3.990 (2.160) | 0.235 (0.727) |
| Standard error (SE) | 4.172 (1.108) | | | |
| Academic Journals×SE | -0.750 (1.285) | | | |
| Constant | 0.044 (0.041) | 1.096 (0.391) | 1.567 (0.478) | 1.095 (0.368) |
| Nudges | 317 | 317 | 317 | 317 |
| Trials | 152 | 152 | 152 | 152 |
| R-squared | 0.110 | 0.021 | 0.077 | 0.001 |
| Weighted by 1/SE ² | ✓ | | | |
| Weighted by 1/MDE | | ✓ | | ✓ |
| Weighted by P-score from nudge categories | | | ✓ | ✓ |

Standard errors clustered by trial are shown in parentheses. The coefficient on Academic Journals sample is the estimated average difference in percentage point (pp.) treatment effects between the Academic Journals and Nudge Units samples. MDE (minimum detectable effect) is calculated in pp. at power 0.8. P-score is the propensity score of being in the Academic Journals sample using predicted probabilities from a logit regression that includes the same nudge category controls as in Column 3 of Table 4.

Table A8: Traditional meta-analysis models (without correction for publication bias)

| | | Academic Journals | | Nudge Units | | Published Nudge Units | |
|---|------------|-------------------|---------------------|------------------|---------------------|-----------------------|---------------------|
| True study-level effects distributional assumption | | (1) ATE (pp.) | (2) $\hat{\tau}$ | (3) ATE (pp.) | (4) $\hat{\tau}$ | (5) ATE (pp.) | (6) $\hat{\tau}$ |
| Unweighted | None | 8.68 (2.47) | – | 1.37 (0.30) | – | 0.97 (0.23) | – |
| Maximum Likelihood | Normal | 7.86 (2.11) | 9.68 | 1.31 (0.27) | 3.49 | 0.55 (0.14) | 0.34 |
| Empirical Bayes | Normal | 7.95 (2.15) | 10.40 | 1.31 (0.27) | 3.70 | 0.62 (0.14) | 0.49 |
| DerSimonian-Laird | None | 5.41 (1.42) | 2.53 | 0.93 (0.17) | 0.63 | 0.57 (0.14) | 0.38 |
| Card, Kluve, and Weber (2018) | None | 2.54 (1.26) | – | 1.25 (0.25) | – | 0.82 (0.18) | – |
| Fixed effect | Degenerate | 2.40 (1.09) | 0.00 | 1.22 (0.38) | 0.00 | 0.71 (0.16) | 0.00 |

This table shows the average treatment effects using various meta-analysis methods. Standard errors clustered by trial are shown in parentheses. $\hat{\tau}$ is the estimated standard deviation in between-study true effect sizes. Following Card, Kluve, and Weber (2018), we winsorize weights from their method at the 10th and 90th percentiles. Mantel-Haenszel weights are used for the fixed-effect model.

Table A9a: Generalized meta-analysis models: Additional specifications

| | | | Normal 1 | | | | Normal 2 | | | | Normal 3 | | | | |
|---|----------------|----------------------------|-----------------------|--------------------|--------------------|----------------|-----------------------|--------------------|--------------------|----------------|-----------------------|--------------------|--------------------|----------------|-----------------|
| | ATE (pp.) | $\hat{\gamma}$ (pub. bias) | $\hat{\hat{\beta}}_1$ | $\hat{\tau}_{BT1}$ | $\hat{\tau}_{WI1}$ | $\hat{P}(N1)$ | $\hat{\hat{\beta}}_2$ | $\hat{\tau}_{BT2}$ | $\hat{\tau}_{WI2}$ | $\hat{P}(N2)$ | $\hat{\hat{\beta}}_3$ | $\hat{\tau}_{BT3}$ | $\hat{\tau}_{WI3}$ | $\hat{P}(N3)$ | -Log likelihood |
| <i>Panel A. Traditional parametric normal-based meta-analysis</i> | | | | | | | | | | | | | | | |
| Academic Journals | 6.32 (3.32) | 0.35 (0.75) | 6.32 (3.32) | 9.00 (2.53) | 5.53 (2.76) | 1 | — | — | — | — | — | — | — | — | 259.93 |
| Published Nudge Units | 0.68 (0.33) | 1 (fixed) | 0.68 (0.33) | 0.45 (0.28) | 0.14 (0.06) | 1 | — | — | — | — | — | — | — | — | 31.66 |
| Published Nudge Units | 0.35 (0.23) | 0.07 (0.08) | 0.35 (0.23) | 0.42 (0.20) | 0.13 (0.06) | 1 | — | — | — | — | — | — | — | — | 26.15 |
| <i>Panel B. Mixture of two normals meta-analysis</i> | | | | | | | | | | | | | | | |
| Academic Journals | 8.47 (2.16) | 1 (fixed) | 3.09 (1.00) | 2.48 (0.81) | 0.04 (0.20) | 0.69 (0.11) | 20.43 (4.65) | 5.44 (2.95) | 12.41 (3.07) | 0.31 (0.11) | — | — | — | — | 216.56 |
| Published Nudge Units | 1.07 (0.35) | 1 (fixed) | 0.29 (0.16) | 0.13 (0.11) | 0.26 (0.06) | 0.53 (0.16) | 2.74 (0.65) | 0.00 (0.04) | 0.00 (0.04) | 0.47 (0.16) | — | — | — | — | 28.69 |
| Published Nudge Units | 0.36 (0.23) | 0.07 (0.14) | 0.08 (0.06) | 0.04 (0.03) | 0.41 (0.17) | 0.91 (0.15) | 0.75 (0.53) | 0.11 (0.28) | 0.17 (0.08) | 0.09 (0.15) | — | — | — | — | 23.96 |
| <i>Panel C. Mixture of three normals meta-analysis</i> | | | | | | | | | | | | | | | |
| Academic Journals | 0.26 (1.58) | 0.17 (0.10) | 0.03 (0.49) | 0.60 (0.36) | 3.17 (0.16) | 0.07 (0.14) | 3.12 (1.64) | 2.88 (1.34) | 0.00 (0.28) | 0.29 (0.11) | 19.18 (5.14) | 5.91 (2.92) | 12.82 (3.16) | 0.11 (0.06) | 205.65 |
| Nudge Units | 1.47 (0.37) | 1 (fixed) | 0.20 (0.19) | 0.28 (2.16) | 0.11 (0.52) | 0.60 (0.14) | 2.29 (0.84) | 1.83 (1.23) | 0.66 (0.28) | 0.33 (0.10) | 8.54 (5.39) | 0.00 (5.37) | 13.19 (5.23) | 0.07 (0.08) | 359.55 |

This table shows additional results from generalized normal-based meta-analysis model in Table 5. Under the normal-based meta-analysis assumptions in Panel A, trial base effects β_i are drawn from a normal distribution centered at $\bar{\beta}$ with between-trial standard deviation τ_{BT} . Then, each treatment arm j within a trial i draws a base treatment effect $\beta_{ij} \sim N(\beta_i, \tau_{WI}^2)$, where τ_{WI} is the within-trial standard deviation. Each treatment arm also has some level of precision given by an independent standard error σ_{ij} . The observed treatment effect is $\hat{\beta}_{ij} \sim N(\beta_{ij}, \sigma_{ij}^2)$. In Panel B, the mixture of normals model is a generalization of the normal-based meta-analysis, and allows trial base effects to be drawn from a second normal distribution (and a third, in Panel C). $\hat{P}(N)$ is the estimated proportion of effects drawn from each normal distribution. To capture the extent of selective publication, the probability of publication is allowed to differ depending on whether trial have at least one significant treatment arm. In particular, trials without any significant results at the 95% level are γ times as likely to be published as trials with significant results. Estimates are obtained using maximum likelihood, and bootstrap standard errors are shown in parentheses.

Table A9b: Mixture of three normals with stacked data

| | Sample | | | Parameters of normals | | |
|----------------------|-------------------|-------------|----------|-----------------------|------------------|-----------------|
| | Academic Journals | Nudge Units | | Mean | Between-trial SD | Within-trial SD |
| $P(\text{Normal 1})$ | 0.49 (0.14) | 0.62 (0.08) | Normal 1 | 0.41 (0.08) | 0.52 (0.08) | 0.24 (0.06) |
| $P(\text{Normal 2})$ | 0.39 (0.10) | 0.31 (0.07) | Normal 2 | 2.54 (0.63) | 2.11 (0.48) | 0.65 (0.18) |
| $P(\text{Normal 3})$ | 0.12 (0.07) | 0.07 (0.03) | Normal 3 | 13.32 (3.43) | 7.77 (4.55) | 12.92 (1.52) |
| ATE (pp.) | 2.72 (1.06) | 1.76 (0.29) | | | | |
| Pub. bias | 0.07 (0.06) | 1 (fixed) | | | | |
| -Log likelihood | 208.15 | 360.78 | | | | |

This table shows the joint estimation of the mixture of three normals meta-analysis combining both the Academic Journals and Nudge Units samples of nudges. (Panel C of Table A9a presents the results when the model is estimated separately for the two samples.) The mean, between-trial variance, and within-trial variance of each of the three normal distributions are assumed to be the same for both samples of nudges, and the two samples only differ in the probability of drawing a trial from each of the normals. The probabilities of drawing from the three normals are modeled using ordinal probit assumptions (see notes in Table A9c for details). The results in this table correspond to Column 2 in Table A9c. Standard errors from 50 bootstrapped samples are shown in parentheses.

Table A9c: Ordinal probit estimates from a mixture of three normals meta-analysis with stacked data

| | (1) | (2) | (3) | (4) |
|--------------------------------|----------------|----------------|-----------------|-----------------|
| Academic Journals | 1.22 (0.28) | 0.34 (0.35) | 1.08 (0.34) | 0.17 (0.43) |
| Email | | | -0.07 (0.30) | -0.09 (0.32) |
| In-person | | | 0.99 (0.73) | 1.19 (0.53) |
| Control receives communication | | | 0.19 (0.34) | 0.16 (0.33) |
| Revenue & debt | | | 0.05 (0.41) | 0.09 (0.36) |
| Health | | | -0.16 (0.36) | -0.20 (0.33) |
| θ_1 | 0.31 (0.24) | 0.31 (0.20) | 0.39 (0.27) | 0.37 (0.22) |
| θ_2 | 1.58 (0.28) | 1.50 (0.22) | 1.69 (0.37) | 1.60 (0.33) |
| γ | 1 (fixed) – | 0.07 (0.06) | 1 (fixed) – | 0.06 (0.08) |
| <i>Academic Journals</i> | | | | |
| ATE at means (pp.) | 6.66 (1.75) | 2.72 (1.05) | 7.23 (1.81) | 3.17 (1.57) |
| $P(\text{Normal } 3)$ | 0.36 (0.11) | 0.12 (0.07) | 0.39 (0.11) | 0.15 (0.07) |
| <i>Nudge Units</i> | | | | |
| ATE at means (pp.) | 1.88 (0.32) | 1.81 (0.29) | 1.84 (0.35) | 1.73 (0.31) |
| $P(\text{Normal } 3)$ | 0.06 (0.08) | 0.07 (0.03) | 0.05 (0.02) | 0.06 (0.02) |
| -Log likelihood | 577.85 | 568.93 | 575.86 | 566.44 |
| Nudges | 317 | 317 | 317 | 317 |
| Trials | 152 | 152 | 152 | 152 |

This table shows results from the mixture of three normals meta-analysis on a stacked data set combining both Academic Journal and Nudge Unit samples of nudges. The parameters of each of the three normals (mean, between-trial variance, and within-trial variance) are held constant between both samples. The two samples of nudges differ in the probability of drawing a trial from each of the three normals. These probabilities are estimated under an ordinal probit model. Specifically, the probability that a trial i draws its effect size from the first (lowest) normal is $P(X_i'\beta + \varepsilon < \theta_1)$, where X_i is a $k \times 1$ vector of trial characteristics, such as being in the Academic Journal sample. β is a $k \times 1$ vector of coefficients, and the error ε follows a standard normal distribution. The probability that a trial i draws its effect size from the second (middle) normal is $P(\theta_1 \leq X_i'\beta + \varepsilon < \theta_2)$, and the probability of drawing from the third (highest) normal is $P(\theta_2 \leq X_i'\beta + \varepsilon)$. The thresholds θ_1, θ_2 and the coefficient vector β are jointly estimated. This table shows the estimated coefficients for trial-level features (e.g., being in the Academic Journal sample or the Revenue & debt policy area), the thresholds in the ordinal probit, and γ , the probability that a trial with no significant results is published relative to a trial with at least one significant result. Below these estimates, the table shows the average treatment effect (ATE) and the probability of drawing a trial from the third (highest) normal for the two samples separately. The ATE is calculated holding X_i at the averages for each of the two samples. Standard errors from at least 50 bootstrap samples are reported in parentheses.