

The Buy-In Effect: When Increasing Initial Effort Encourages Follow-Through

Holly Dykstra, Shibeal O’Flaherty, and Ashley Whillans*

July 21, 2025

Abstract

Behavioral interventions often focus on reducing friction to encourage behavior change. In contrast, we provide evidence that adding friction can promote behavior change when the target behavior requires follow-through. In collaboration with the Oregon Department of Transportation, we conducted a field experiment ($N = 27,227$) to test whether adding friction during an initial sign-up process for a new carpooling platform increases usage. Our results support this possibility: while a more effortful sign-up process led to a 25% decrease in sign-ups to the carpool platform, overall intensity of usage increased. Importantly, these results were only partly explained by selection effects: using an intention-to-treat analysis, participants who were randomly assigned to the more effortful sign-up process took 1.6 times more carpool trips per week over a four-month period as compared to those in the less effortful sign-up process. Of the 9,417 observed trips, the more effortful sign-up group took almost 800 more trips. To test generalizability and mechanisms, we conducted a second, pre-registered online experiment where participants completed transcription tasks and could sign up for a return opportunity the next day. Participants who were randomly assigned to a more effortful sign-up process were 37% more likely to return and complete more work overall. These results suggest that adding friction may be an overlooked strategy that could help promote behavior change, especially when follow-through, rather than initial uptake, is the primary goal.

Keywords: behavior change; friction; follow-through; subjective value

*Dykstra: holly.dykstra@uni-konstanz.edu, University of Konstanz; O’Flaherty: shibeal.oflaherty@kcl.ac.uk, King’s College London; Whillans: awhillans@hbs.edu, Harvard Business School. The authors thank the Oregon Department of Transportation for their collaboration on this project, as well as their partner Alta Planning and Design who facilitated implementation—in particular, Jessica Roberts and Hannah Mullin. Alessio Fikre and Matthew Gross provided excellent research assistance on the literature review, and Mihane Sermaxhaj provided excellent research assistance on the analysis. For helpful comments and feedback, the authors thank Syon Bhanot, Iris Bohnet, Blair Read, Todd Rogers, and Regina Strumpf, as well as seminar participants at Carnegie Mellon University, Harvard Business School, Harvard Kennedy School, the Office of Evaluation Sciences, and the University of Konstanz. We acknowledge funding from the University of Konstanz and Harvard Business School for Study 2.

1 Introduction

People regularly fail to follow through on decisions that they believe are in their best interest. This intention-action gap has encouraged behavioral scientists and policymakers to develop interventions that make goal-relevant behaviors easier to adopt (Rogers and Milkman, 2016; Duckworth, Milkman and Laibson, 2018). A popular approach is to reduce friction through simplification strategies such as using defaults or removing administrative barriers (Jachimowicz et al., 2019, DellaVigna and Linos, 2022). These approaches are successful in many contexts where a single decision produces lasting benefits, such as automatic enrollment in retirement savings plans (Choi et al., 2002; Carroll et al., 2009). However, many behaviors require follow-through rather than one-time engagement. In these settings, simplification may increase enrollment without facilitating the active buy-in necessary for follow-through. Indeed, emerging evidence suggests that easy enrollment processes can generate passive participation that may fail to translate into lasting engagement (Frey and Rogers, 2014; Braconnier, Dormagen and Pons, 2017).

Building on this research, we propose the “buy-in effect,” whereby modest, goal-relevant effort during an initial sign-up process increases follow-through. To provide an initial test of the buy-in effect, we conduct two large, pre-specified experiments in two very different settings: carpooling and online task work.

Study 1 is a large-scale, four-month field experiment ($N = 27,227$) with the Oregon Department of Transportation (ODOT), where we manipulate the difficulty of a carpooling platform sign-up process. Carpooling is a relevant context for testing the buy-in effect because it requires follow-through, rather than only an initial sign-up decision, and previous research suggests that carpooling is a difficult behavior to influence through standard behavioral approaches (Whillans et al., 2021). In Study 1, we recruited employees who had previously registered for ODOT’s carpool platform but had been inactive for six months. We randomly assigned these employees to receive one of two sign-up processes to join a new version of the carpooling platform. In the *High Effort* condition, the sign-up process introduced modest, goal-aligned friction by requiring employees to re-enter account and commute-relevant information (i.e., home address, work address, and employer information) as part of the initial sign-up process. In the *Low Effort* condition, employees were offered a streamlined, one-click sign-up process in which this information was automatically migrated from the previous

platform. Consistent with the existence of a buy-in effect, participants in the *High Effort* condition were more likely to follow-through: using an intent-to-treat analysis (ITT), employees logged about 1.6x more carpool trips per week over a four-month period, resulting in almost 800 more carpool trips in total. Using an ITT analysis indicates that this effect cannot be explained by selection effects alone. We also run a boundary exercise to further investigate the role of selection effects.

Study 2 is a pre-registered online experiment in which participants first complete a difficult transcription task, and are then offered the opportunity to return the following day to complete the task again. Similar to Study 1, participants were randomly assigned to either a *High Effort* condition, where the sign-up process involved a 15-question domain-relevant survey, or a *Low Effort* condition, where the sign-up process involved a single click. Consistent with the results of Study 1, participants in the *High Effort* condition were 37% more likely to return the following day and completed significantly more work overall as a result.

We develop a conceptual framework for the buy-in effect. First, we propose two categories of mechanisms through which the buy-in effect might increase follow-through: (1) enhancing the perceived value of the target behavior or (2) increasing attention and recall. In Study 2, to provide a direct test of this second category of mechanisms, we included a cross-randomized *Reminder* condition. Second, we describe the scope conditions under which the buy-in effect is likely to emerge: (1) the target behavior requires follow-through, (2) the initial effort and target behavior are voluntary in nature, and the initial effort is (3) modest and (4) connected to the target behavior. While the two studies presented in our paper are meant to provide an initial road test of the buy-in effect, and therefore do not formally test the proposed scope conditions, our conceptual framework helps to distinguish contexts where friction may be beneficial from those where it may be harmful.

Our research extends work on the intention-action gap, wherein people fail to follow through on decisions that they believe are in their best interest, due to factors such as present bias and self-control challenges (Laibson, 1997; O'Donoghue and Rabin, 1999; Augenblick and Rabin, 2019). First, our research adds to the set of interventions that have been developed to address this intention-action gap, which includes financial incentives (Charness and Gneezy, 2009; Acland and Levy, 2015), commitment devices (Bryan, Karlan and Nelson, 2010; Royer, Stehr and Sydnor, 2015), reminder systems (Dale and Strauss, 2009; Calzolari and Nardotto, 2017), planning prompts (Milkman et al., 2012), and temptation bundling (Milkman, Minson and Volpp, 2014). Our pa-

per introduces modest, goal-aligned friction as a new strategy to address this gap. Second, our research helps to reconcile seemingly contradictory findings in the literature by distinguishing between contexts where simplification succeeds (“set-and-forget” behaviors) versus contexts where simplification may generate passive enrollment without psychological commitment, as emerging evidence from environmental, tax, and civic domains suggest. For example, interventions designed to encourage pro-environmental behaviors have small effects while in place but no follow-through effects once they end (Nisa et al., 2019), while simplified tax filing processes can increase initial compliance without improving subsequent filing behavior or program understanding (Guyton et al., 2017; Fonseca and Grimshaw, 2017). Similarly, streamlined voter registration processes can boost early response rates without increasing total registrations (Sweeney et al., 2021), and door-to-door canvassing can increase immediate election turnout without always sustaining civic engagement (Nickerson, 2015; Braconnier, Dormagen and Pons, 2017). This pattern suggests that simplification can produce passive enrollment without encouraging follow-through, and we indeed find that, instead, adding modest, goal-relevant effort during a sign-up process can increase follow-through.

Our work has practical implications for designing interventions in domains where follow-through is critical. Rather than assuming that simplification is the optimal sign-up approach, our findings suggest that modest friction may be an important part of the practitioner toolkit in contexts where follow-through rather than initial uptake is the primary goal.

The rest of this paper is organized as follows. Section 2 describes the related literature in more detail, and distinguishes the buy-in effect from related interventions. Section 3 presents the conceptual framework, describing potential mechanisms as well as scope conditions. The research design and results of Study 1, the carpooling study, are described Section 4. The research design and results of Study 2, the online task study, are described in Section 5. This study also includes additional treatment arms that test the attentional mechanism as well conditions under which the buy-in effect generalizes. In Section 6, we interpret these results by discussing heterogeneity, mechanisms, alternative explanations, limitations, and future directions for this research. Section 7 concludes.

2 Literature Review

Why do people fail to follow through on decisions that they believe are in their best interest? Across domains such as saving for retirement, maintaining a healthy diet, exercising regularly, and engaging in environmentally sustainable behaviors, people frequently plan to act in ways that are aligned with their long-term goals, yet often fall short in accomplishing these aims. Economists and behavioral scientists attribute this intention-action gap to present bias and self-control challenges, wherein people prioritize immediate costs and benefits over future outcomes (Laibson, 1997; O’Donoghue and Rabin, 1999; Augenblick and Rabin, 2019).

To address this intention-action gap, researchers and policymakers have developed numerous intervention strategies ranging from structural changes that reduce participation barriers to psychological strategies that improve motivation and self-regulation (see Chater and Loewenstein, 2022 for an overview). These approaches include financial incentives (Charness and Gneezy, 2009; Acland and Levy, 2015; Loewenstein et al., 2015; DellaVigna and Pope, 2018), commitment devices (Bryan, Karlan and Nelson, 2010; Bhattacharya, Garber and Goldhaber-Fiebert, 2015; Royer, Stehr and Sydnor, 2015), interventions that harness the motivation for autonomy and competence (Deci and Ryan, 2000; Mellström and Johannesson, 2008; Aknin and Whillans, 2021), reminder systems (Coppock and Green, 2016; Calzolari and Nardotto, 2017), planning prompts that help people specify when and how they will act (Verplanken and Faes, 1999; Milkman et al., 2012), and temptation bundling, which pairs goal-directed behaviors with immediately gratifying activities, such as only allowing oneself to listen to engaging audiobooks while exercising at the gym (Milkman, Minson and Volpp, 2014).

One of the most widely adopted and discussed strategies among policymakers and researchers is to simplify the decision-making environment (Sunstein, 2013; Benartzi et al., 2017; Mertens et al., 2022). An influential simplification strategy is the use of defaults, which are pre-selected options that take effect if an individual does not actively make a different choice (Jachimowicz et al., 2019). For example, switching 401(k) retirement plans from opt-in to opt-out almost doubles employee participation (Madrian and Shea, 2001; Carroll et al., 2009). Beyond defaults, reducing administrative burden, such as using shorter forms or requiring fewer documents, can improve take-up (Herd and Moynihan, 2018); for example, simplifying the Earned Income Tax Credit (EITC) work-

sheet raised claim rates for eligible households (Bhargava and Manoli, 2015). In-person assistance, which reduces complexity in real time, is also highly effective (Bettinger et al., 2012; Daigneault et al., 2025). Pre-filled forms and other auto-completion tools similarly lower friction, for example in tax returns (Gillitzer and Skov, 2018). Choice-set simplification—presenting a ready-made bundle or narrower menu—can also be effective, like an enrollment packet that offered employees a pre-selected contribution rate and asset mix for their retirement savings plans (Beshears et al., 2013; Chernev, Boeckenholt and Goodman, 2015). Finally, information nudges that make key facts clearer or more salient at the moment of choice also change decision-making, such as nutrition labels on food (Cadario and Chandon, 2020; Daigneault et al., 2025; De-loyde et al., 2025).

Consistent with these benefits, a meta-analysis of 58 field experiments found that simplifying choices through the use of defaults increased engagement with decisions related to retirement savings, organ donation, and health-plan enrollment (Jachimowicz et al., 2019). These interventions were especially effective when the target behavior was “set-and-forget,” that is, when the behavior of interest required little to no personal effort after an initial decision was made, such as enrolling in a program that continued automatically once it was started.

2.1 The Limitations of Simplification

While defaults and other simplification strategies can increase initial participation, these strategies may fall short in settings that involve follow-through, where removing initial barriers to participation may generate passive enrollment without fostering the active buy-in needed. Studies on environmental, tax, and civic behaviors provide evidence for this proposition, showing that simplifying enrollment processes can increase initial uptake but may fail to produce lasting engagement (e.g., Guyton et al., 2017; Sunstein, 2017; Nisa et al., 2019; Sweeney et al., 2021).

In the environmental domain, a meta-analysis of 83 randomized trials of pro-environmental behavioral interventions, including ones designed to simplify the decision-making process, showed modest effects on behavior while the interventions were in place but no sustained effects once they ended (Nisa et al., 2019). In addition, defaults in environmental domains are less effective than in other domains where follow-through behavior is less important, like in retirement savings (Jachimowicz et al., 2019). Finally, pre-filled forms or initial planning tools can increase intentions to act but have limited success for improving actual engagement with environmental behaviors

that require follow-through, such as carpooling or taking public transit (Whillans et al., 2021); for example, in a study of five field experiments to reduce solo driving—including letters, emails, incentives, and personalized travel plans—none of the interventions succeed in shifting commuter behavior (Kristal and Whillans, 2020).

In the tax domain, simplification strategies such as pre-filled forms and reduced paperwork have been shown to increase initial compliance, filing rates, and benefit uptake (Guyton et al., 2017; Gillitzer and Skov, 2018). However, these simplification approaches often fail to increase long-term engagement or improve participants’ understanding of the tax system. Interventions to encourage likely-eligible people to use free tax preparation software or to claim the EITC increases returns in the current year but not in subsequent years (Guyton et al., 2017; Goldin et al., 2022). Moreover, a study of six nudge RCTs designed to increase EITC filing found no effect on filing even in the current year (Linos et al., 2022). Finally, pre-filled tax forms can also increase inaccuracies and tax evasion, suggesting that, with defaulted-in information, individuals may not engage deeply enough with the filing process (Fonseca and Grimshaw, 2017; van Dijk et al., 2020).

Similar results appear in studies of civic behavior, where reducing upfront effort can increase initial voting behaviors but may have little effect on sustained political engagement. Simplifying voter registration through redesigned envelopes and streamlined forms increases early response rates but does not affect the total number of voter registrations (Sweeney et al., 2021). Postage-paid ballots create no change in voter turnout (Michelson et al., 2012). People registered to vote using home visits are less likely to vote than those who registered on their own (Nickerson, 2015), and they turn out to vote at least once but are less likely to do so in subsequent elections (Braconnier, Dormagen and Pons, 2017). Finally, simplified election date and polling hours information increase turnout for the current election but show no effects in the next election five months later (Hill and Kousser, 2016).

2.2 Distinguishing the Buy-In Effect from Related Interventions

We propose that the buy-in effect is conceptually distinct from, but shares similarities with, several established behavior-change strategies. In the following section, we clarify how the proposed buy-in effect differs from commitment contracts, planning prompts, and financial incentives.

Commitment Devices. Commitment devices are mechanisms through which an individual can

impose personal costs after non-compliance, such as financial penalties for failing to meet savings or exercise goals (Bryan, Karlan and Nelson, 2010; Rogers, Milkman and Volpp, 2014). In contrast, the buy-in effect focuses on increasing effort at an initial decision point rather than creating future penalties. While both strategies involve upfront decision-making and therefore upfront effort, the buy-in effect does not rely on an additional strategy after this initial upfront effort.

In addition, commitment devices require people to anticipate self-control failures in advance and choose to actively bind themselves, meaning that engaged participants are typically highly motivated at baseline (Thaler and Benartzi, 2004; Bryan, Karlan and Nelson, 2010; Rogers, Milkman and Volpp, 2014; Beshears et al., 2020). While the buy-in effect does not require a person to be sophisticated about their future inability to follow through, upfront effort may similarly select for more motivated people who are most willing to engage in additional effort. However, this barrier to entry is lower, since both strategies require upfront effort, but the buy-in effect does not require the additional awareness and foresight that commitment devices do.

Planning Prompts. Planning prompts ask people to write down when and where they plan to complete a task, a strategy that can reliably increase follow-through (Milkman et al., 2012). Both the buy-in effect and planning prompts involve some form of upfront effort, and both may activate similar mechanisms: psychological ownership over the target behavior, identity-consistent signaling, and improved memory recall. However, the buy-in effect isolates the impact of modest effort alone, without requiring people to explicitly engage in scheduling.

Monetary rewards and costs. Previous work has studied the use of financial incentives on follow-through, including the use of monetary payments to reward behavior, or manipulating upfront financial costs to promote persistence. Receiving compensation to follow through with a target behavior has been successful across a variety of settings (Charness and Gneezy, 2009; Acland and Levy, 2015; Loewenstein, Price and Volpp, 2016; DellaVigna and Pope, 2018). However, the effects of financial incentives often decay once the incentives stop, which is usually attributed to the crowding out of personal motivation that happens due to extrinsic rewards, particularly in prosocial or identity-relevant contexts (Gneezy and Rustichini, 2000; Mellström and Johannesson, 2008).

Manipulating upfront cost-based frictions does not consistently promote follow-through. Research finds that requiring payment for items such as bed nets or water purifiers reduces initial

take-up and does not improve subsequent use (Ashraf, Berry and Shapiro, 2010; Cohen and Dupas, 2010). Similarly, eliminating financial barriers to access, such as reducing transit costs on election day, can increase intermediate behaviors like bus ridership without affecting ultimate outcomes like voting (Pereira et al., 2023; Cantoni, Pons and Schafer, 2025).

The buy-in effect differs from these literatures because we focus on introducing non-monetary, goal-aligned effort during an initial sign-up phase. Participants in our experiments do not pay financially or receive financial rewards. Instead, they are asked to expend effort that is meaningfully tied to the target behavior. This focus on effort may lead to important differences in the motivational framing of the target behavior, which could affect follow-through behavior (Charness and Gneezy, 2009; Acland and Levy, 2015).

3 Conceptual Framework

Building on previous research, which suggests but does not formally test the potential limitations of simplification, we propose the buy-in effect, a behavioral phenomenon wherein introducing modest effort during an initial sign-up process can increase follow-through.

Rather than assuming that all friction is detrimental for follow-through, we propose that initial effort has the potential to encourage follow-through when certain conditions are met. In the sections below, we (1) outline the potential mechanisms underlying this effect and (2) highlight the contexts in which the buy-effect is most likely to emerge.

3.1 Mechanisms

The buy-in effect may be explained through two broad channels of mechanisms: (1) increasing the perceived value of the target behavior and (2) improving attention and recall. Both proposed channels include specific underlying psychological processes.

Increasing Perceived Value. Modest effort at the point of an initial decision, such as during a sign-up process, may increase follow-through by increasing the perceived value of the target behavior. This proposed mechanism is consistent with research on psychological ownership and the “IKEA effect,” where individuals come to value outcomes more when they have invested effort into creating them (Shu and Peck, 2011; Norton, Mochon and Ariely, 2012). People facing modest effort may also experience a sunk-cost effect, which occurs when people value items or actions more when they have already spent time, money, or effort on them (Arkes and Blumer, 1985; Roth, Robbert

and Straus, 2015.

This mechanism is also consistent with research on self-signaling. When individuals expend effort during an initial sign-up process, they may infer that the target behavior is more important to them, reinforcing their self-concept as someone who values the behavior and thus increasing how much they believe the behavior is worth doing (Bénabou and Tirole, 2006; Gneezy et al., 2012; Bryan et al., 2016). Relatedly, the literature on cognitive dissonance—in which a person who holds two inconsistent beliefs seeks to resolve it—could lead to someone who exerts more upfront effort to move their perception of the behavior higher to match this previous effort (Hinojosa et al., 2017). Finally, people often infer value from contextual cues, such as how difficult something is to obtain (Kamenica, 2008; Inzlicht, Shenhav and Olivola, 2018), and research shows that time investments are an important signal of how valuable an activity is personally and to society more broadly (Kruger et al., 2004; Shaddy and Shah, 2018, 2022). Thus, adding modest friction to an initial sign-up process may increase follow-through by making the behavior seem more worthwhile.

Increasing Attention and Recall. Modest effort during an initial sign-up process might also improve follow-through by increasing attention and memory recall. Research suggests that when individuals engage in active processing during an initial decision-making step, they are more likely to encode intentions to complete the target behavior in their memory and retrieve them at relevant moments, thereby increasing follow-through (Gollwitzer and Sheeran, 2006; Schacter, Addis and Buckner, 2007; Kahneman, 2011). Studies on planning prompts and timely reminders show that deliberation at the time of an initial decision can encourage follow-through by making a target behavior more cognitively accessible. Specifically, increased deliberation can make people more likely to remember and follow-through with their intended behavior (Milkman et al., 2012; Rogers and Milkman, 2016; see also Gollwitzer and Sheeran, 2006; Duckworth, Gendler and Gross, 2016 for related research). Additionally, a systematic and meta-analytic review of 36 studies looked at the role of memory in implementation intentions, and found that they improve prospective memory, that is, remembering to take a planned action in the future (Chen et al., 2015). Consequently, a modest initial hurdle may encourage greater attention, improve recall, and increase the likelihood that an individual will follow through with the target behavior when they are presented with the opportunity to do so.

3.2 Scope Conditions

Despite the possibility that increasing friction during a sign-up process can increase follow-through, the buy-in effect is unlikely to generalize to all contexts or types of friction. Building on prior research, we propose that the buy-in effect is most likely to emerge (1) in domains that require follow-through, (2) when the sign-up process and engagement in the target behavior is voluntary, and when the initial friction is (3) modest and (4) goal-relevant.

3.2.1 Types of Contexts

Follow-Through. First, we propose that the buy-in effect is likely to emerge in domains that require follow-through, rather than for “set-and-forget” behaviors. This scope condition is supported by research showing that reducing friction can increase uptake, but often fails to produce behavioral change in contexts that require follow-through (Bhargava and Manoli, 2015; Braconnier, Dormagen and Pons, 2017; Nisa et al., 2019).

In addition, in contexts that require follow-through—unlike in context “set-and-forget” behaviors—people may be willing to exert more effort upfront because the target behavior is more effortful, and this cost is perceived as proportional to the demands of the target behavior. From a behavioral economics perspective, upfront effort may be viewed as a rational investment, particularly when the target behavior requires ongoing self-regulation (Laibson, 1997; O’Donoghue and Rabin, 1999). Follow-through contexts are more likely to involve behaviors that are effortful enough to make modest upfront costs seem reasonable and proportional.

Voluntary Participation. Second, we propose that the buy-in effect is more likely to emerge when participation is voluntary. This is because, in voluntary settings, the upfront investment of effort can activate self-signaling and personal commitment; when individuals cannot choose whether to sign up or engage in the target behavior, initial effort may not carry the same informational value about personal motivation (Bénabou and Tirole, 2003a, 2006). This proposed scope condition aligns with research on enhanced active choice, which shows that requiring individuals to actively choose to opt in or out (rather than defaulting them to a decision) can increase engagement by prompting self-reflection and encouraging individuals to connect the decision to their personal values (Keller et al., 2011; Beshears et al., 2021). This scope condition is also consistent with self-signaling theory (Bénabou and Tirole, 2003b) and findings from the social-psychological literature

which show that behaviors are more motivating and self-reinforcing when they are freely chosen (Deci and Ryan, 1985; Miller and Prentice, 2016). These studies suggest that freedom of choice may be important for effort-based interventions to enable individuals to interpret their effort as meaningful and self-relevant, and thus to increase follow-through.

3.2.2 Types of Friction

Modest Friction. Third, we propose that initial effort must be modest for the benefits of the buy-in effect to outweigh the potential costs. This is because excessive friction may reduce uptake without commensurately increasing commitment. In addition, excessive upfront requirements may backfire among those who complete the sign-up process, causing them to negatively perceive the target behavior. Indeed, previous work suggests that higher levels of friction can undermine participation when effort costs become too burdensome (Inzlicht, Shenhav and Olivola, 2018). In a large-scale field experiment, requiring in-person registration for a government cash-transfer program improved targeting when the registration sites were nearby; however, as the distance to these registration sites increased, take-up sharply declined and previously observed targeting improvements plateaued (Alatas et al., 2016). In consumer lending, increased application complexity reduced loan take-up (Bertrand, Mullainathan and Shafir, 2006; Bertrand et al., 2010), and in task-performance settings, the presentation of very effortful tasks lowered participation even when generous financial incentives were offered (DellaVigna and Pope, 2018). More broadly, this proposed scope condition aligns with conceptual and empirical research on “sludge,” which highlights how excessive administrative burdens can create costs that reduce access without improving program outcomes (Sunstein, 2021).

These findings suggest that there may be an optimal level of friction, which involves enough friction to increase the perception of value and subsequent commitment to the target behavior, but not so much friction as to create prohibitive barriers to participation or negative perceptions.

Goal-Relevant Friction. Finally, we propose that the buy-in effect is more likely to emerge when the friction is meaningfully tied to the behavior itself; that is, when the friction is goal-relevant. This proposition is supported by research showing that the effectiveness of friction depends not only on its presence, but on how it is perceived. When effort is seen as directly related to the desired outcome or as a signal of program quality, friction has the potential to increase follow-through

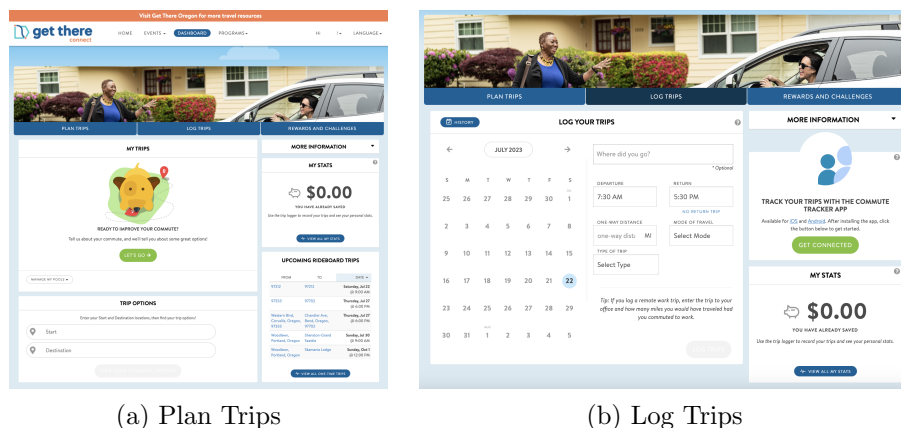
(Kruger et al., 2004; Buell and Norton, 2011). In contrast, when the friction feels arbitrary or disconnected from the target behavior, it may deter participation without improving outcomes, consistent with research on irrelevant administrative burdens (Herd and Moynihan, 2018; Emens, 2021; Sunstein, 2021).

4 Study 1: Field Experiment in Carpooling

4.1 Study Design

We conducted Study 1 with a field partner, the Oregon Department of Transportation (ODOT), which is responsible for systems of transportation in the U.S. State of Oregon. ODOT oversees a platform to promote carpooling and multimodal trips called Get There Connect. This platform, which is available to everyone who lives or works in the state of Oregon and is often promoted by major employers, helps match individuals with others driving in a similar direction. The underlying software is developed and run by RideAmigos.

As of our study time period, users could access the platform through the Get There Connect website or through RideAmigos’ CommuteTracker app. Users could also track trips using the physical exercise app Strava. Figure 1 shows images of what planning new trips and logging trips looked like to users.



Note: These screenshots show the Get There Connect website from the user perspective during our study time period. (a) Shows the “Plan Trips” function within the “Dashboard” tab where users can find carpool matches, and (b) shows the “Log Trips” function.

Figure 1: Get There Connect Platform

Finding carpool matches: In the “Plan Trips” tab, presented in Panel (a) of Figure 1, users could search for carpool matches by entering their start and end zip codes. On the next page, they

were presented with a list of matches, which they could narrow down by entering their schedule (i.e., days and times of week).

Logging trips: Users could log trips through the website or through the mobile apps CommuteTracker or Strava. If users manually logged trips, they must enter additional trip information, as shown in Panel (b) of Figure 1. If users enabled location tracking, the apps would automatically detect trips taken, which the user could review, edit, and confirm. Users must confirm automatically detected trips for these trips to be recorded.

In 2019, ODOT migrated its carpool platform to RideAmigos from a different service provider. This migration required users of the old platform to make an account on the new platform. This migration process provided us with an ideal setting to test the buy-in effect, since user information from the old platform was already available to our field partner, allowing us to randomize users into a treatment condition where that information was already migrated over (*Low Effort*) as compared to a treatment condition in which they had to enter it anew (*High Effort*), described in more detail in Section 4.1.1. In order to find carpool matches, all study participants had to re-enter their information as shown in Figure 1.

We pre-specified the study design, including directional hypotheses and treatment conditions, in pre-trial IRB documents submitted and approved by Harvard University’s Institutional Review Board (AEARCTR-0013133).¹

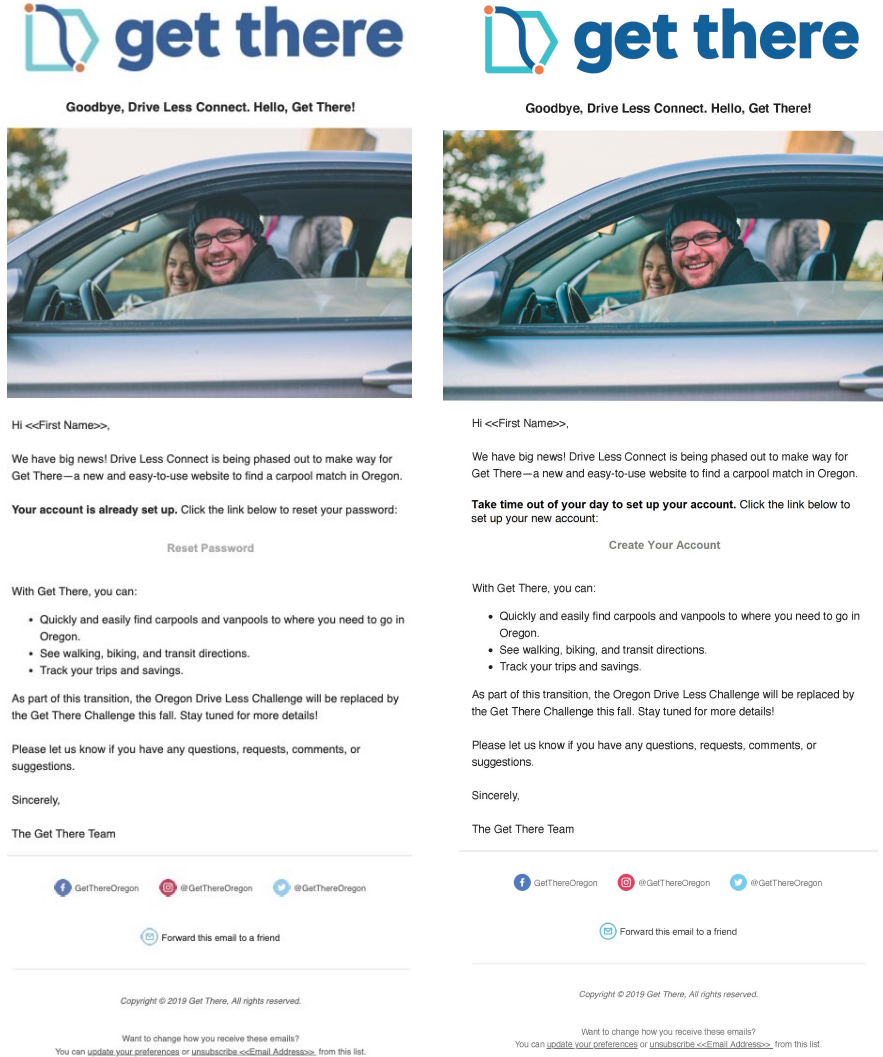
4.1.1 Treatment Conditions

Users of the old carpool platform received an email informing them of the new platform and asking them to sign up. Each participant was randomized into either the *Low Effort* condition or the *High Effort* condition, which involved a different email and sign-up process. Both versions of the email were designed to be similar and to differ only in their description of the effort required to migrate their account. These emails are shown in Figure 2.

Low Effort: In the *Low Effort* email and sign-up process, account information was automatically migrated and required only a password reset. In this version of the email, participants were informed that their account was already set up for them, and that they could click on the “Reset Password” link included in the email to access their account. Upon clicking the link, these participants

¹The AEA registration occurred after data collection, following prior work that uses pre-trial materials to document pre-analysis plans (e.g., Resnjanskij et al., 2024).

Figure 2: Treatment Emails



Note: These screenshots show the emails that study participants received. (a) Shows the *Low Effort* version of the email, and (b) shows the *High Effort* version.

completed only one step to access their accounts: creating a new password. After entering their new password, they saw a prompt indicating that they had received a verification email to their email address.

High Effort: In the *High Effort* email and sign-up process, participants had to remake their account. In this version of the email, participants were told that they would need to create their account, and that they could click on the “Create Your Account” link included in the email to do so. Upon clicking this link, these participants first had to enter their first name, last name, and

email address, and were prompted to create a new password. Then, they moved onto a second page, where they were asked to input additional information, including their home zip code, work zip code, and organization information. After submitting their information, participants saw the same prompt indicating that they had received a verification email to their email address.

Critically, all study participants received identical information about the new platform and, after creating their account and logging in, had to complete the same steps to find carpool matches and log trips. There were no automatic matches created for either group. The only change in participants' carpool platform experience was how much effort they had to exert to sign up. This design ensured that the treatment condition only manipulated the level of effort that was required to sign up for the new carpooling platform and did not change other factors that could impact the subsequent use of the carpooling platform, such as objective information about the program or how easy it was to find matches.

4.1.2 Sample and Data

Our sample consists of inactive users of ODOT's old carpool platform, who had previously signed up for the platform but had not logged in for at least six months. We used this sample for both conceptual and practical reasons. Conceptually, while these users were less likely to sign up for the new carpool platform themselves due to recent inactivity, their prior registration indicates their baseline interest in carpooling. This suggests that they might be amenable to interventions that are designed to increase their commitment to a personally relevant target behavior (i.e., carpooling). Practically, our implementation partner ODOT was interested in increasing the percentage of these inactive users—who accounted for 87% of their users—to successfully transfer over and engage with the new platform once it was launched.

We received the following data on each study participants' carpool activity on the platform: the total number of carpool trips, the total number of miles of carpooling, and the date of first sign-up to the platform.² We measured outcomes through four months after the initial emails were sent.³ For those participants who signed up, we also received data on their home and work zip codes.

²Note that we pre-specified that we would receive data on whether study participants looked for carpool matches. However, ODOT was not able to send us this data.

³Note that the number of carpool trips and miles had to be logged or confirmed by participants. Therefore, we may not see the full extent of participants' commute patterns.

4.1.3 Study Timeline

The study timeline ran as follows:

First wave randomization—June 2019: After accounting for unusable email addresses (e.g., no longer in use, syntax errors, or already unsubscribed from communications), we had a sample of $n = 17,036$ participants. $n = 8,589$ were randomly assigned to the *Low Effort* condition and $n = 8,447$ to the *High Effort* condition. The study emails were sent on June 29, 2019. After implementation, we discovered that one of the condition assignments—the *High Effort* condition—did not receive the emails until two days later, July 1, 2019. For that reason and others, our main outcome measure accounts for the number of days that participants were on the platform, as described in Section 4.2.

Second wave randomization—October 2019: Due to an administrative error, ODOT discovered that not all inactive users were originally pulled from the original database, so we conducted a second wave of randomization on these participants who had not yet been contacted by us. After again accounting for unusable email addresses, we had a total sample of $n = 10,191$ participants, with $n = 5,074$ then randomly assigned to the *Low Effort* condition and $n = 5,117$ to the *High Effort* condition in this wave. These emails were sent on October 7, 2019. ⁴

We collected study outcome data after four months, which included data from June 29, 2019 through October 31, 2019, to coincide with a time of consistent commuting before the U.S. holidays interrupted commuting patterns.⁵ Because of the administrative error described above, we analyze our results including an indicator variable for receiving the email in October as a robustness check below. We also received eight-month follow-up data through February 29, 2020. These results can be found in Section 4.2.1 and Section 4.2.2.

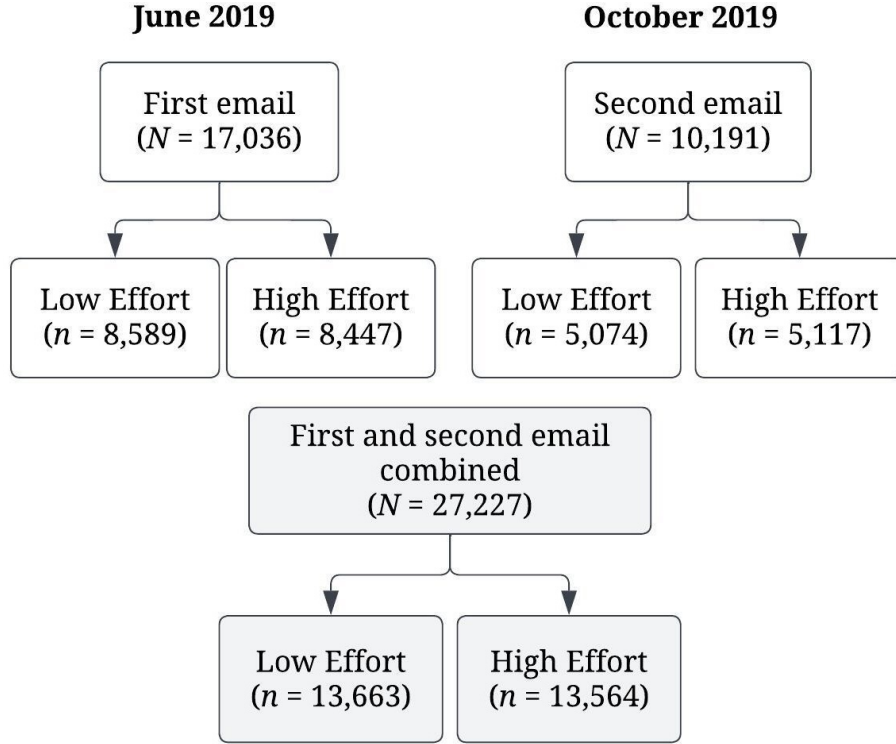
4.1.4 Descriptive Statistics

The descriptive statistics for the full ITT sample ($N = 27,227$) as well as for the conditional sample of participants who signed up ($n = 1,205$, or about 4.4%) are reported in Table 1. In the ITT sample, 0.35 trips were taken per person (9,417 in total), which resulted in 2.96 miles carpooled per person (80,581 in total) across an average of 3.29 days on the platform during our study period.

⁴Note that we pre-specified that we would have a sample size of 40,000, but the sample size sent to us by ODOT was lower after accounting for unusable email addresses and the administrative error described above.

⁵The U.S. holiday season is typically from mid-to-late November through to early January, incorporating multiple major U.S. holidays.

Figure 3: Study Timeline



Note: This figure shows the timeline and sample sizes for the two study waves and treatment conditions.

In the conditional sample, 7.81 trips were taken per person, which resulted in 66.9 miles per person across an average of 74 days on the platform.

We conduct a balance check to test whether conditions were equal with respect to the data we had at baseline, which was gender and race.⁶⁷ We also conduct a balance check with respect to being in our first and second wave, that is, whether participants received an email in June or October. See Table 2 for results. The balance of White study participants was slightly different, 76.5% vs. 77.8% ($p < 0.01$), which we control for in robustness checks in Section 4.2.1.

⁶This data was not held by our partner. Instead, we used the R package `predictrace` to predict gender using first names and race using last names based on U.S. Census and Social Security Administration Data (Kaplan, 2023). While R's `predictrace` was noted as the best performing algorithm in a recent review paper, it is still not very reliable (Lockhart, King and Munsch, 2023). However, we expected errors to be evenly distributed across conditions, which suffices for our purposes here.

⁷We only have geographic data for those who actually signed up to the platform, which we therefore expect to vary by condition, and cannot be used as a baseline check.

Table 1: Descriptive Statistics

	ITT		Conditional	
	<i>Mean</i>	<i>SE</i>	<i>Mean</i>	<i>SE</i>
Signed Up	0.04	0.21	1.00	0.00
Trips	0.35	3.69	7.81	15.79
Miles	2.96	49.06	66.87	223.95
Days on Platform	3.29	18.48	74.24	49.54
Trips per Week	0.06	1.01	1.40	4.58
Miles per Week	0.47	7.34	10.52	33.37
Female	0.58	0.49	0.58	0.49
<i>Race and Ethnicity</i>				
White	0.77	0.42	0.78	0.41
Asian	0.04	0.20	0.04	0.19
Hispanic	0.05	0.22	0.03	0.18
Black	0.01	0.11	0.01	0.10
American Indian	0.00	0.02	0.00	0.00
Observations	27,227		1,205	

Note: This table shows the mean and standard error of each of the displayed variables. The first two columns show them for the full ITT sample, and the last two columns show them for the sample conditional on having signed up.

4.2 Experimental Results

We analyze the effect of our experiment on initial sign-ups to the carpool platform and follow-through. We measure follow-through using two main outcomes of interest: the average number of carpool trips taken per week and the average number of carpool miles per week. We define “per week” as the total trips (or miles) that the participant took divided by the number of weeks that the participant was active on the platform (i.e., from first sign-up to the end of our study period). We use this usage intensity measure so that our results are not biased by differences in how long within our finite timeframe each participant was active on the platform. Since this is a standardized rate measure rather than an absolute total, it also makes it easier to compare our findings with other studies with different study period lengths (Mertens et al., 2022). Finally, as discussed in Section 4.1.3, administrative errors meant that certain groups received the treatment email at different times, making this adjustment important.

Throughout our main results, we estimate the following ordinary least squares (OLS) equation:

Table 2: Balance Table

Variable	(1) 0 Mean/(SE)	(2) 1 Mean/(SE)	(1)-(2) Pairwise t-test Mean difference
October	0.371 (0.004)	0.377 (0.004)	-0.006
Female	0.563 (0.004)	0.556 (0.004)	0.007
<i>Race and Ethnicity</i>			
American Indian	0.000 (0.000)	0.000 (0.000)	0.000
Asian	0.042 (0.002)	0.039 (0.002)	0.004
Black	0.013 (0.001)	0.013 (0.001)	0.001
Hispanic	0.054 (0.002)	0.049 (0.002)	0.004
Two or More Races	0.000 (0.000)	0.000 (0.000)	0.000
White	0.765 (0.004)	0.778 (0.004)	-0.014***
Number of observations	13663	13564	27227

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table shows the mean and standard error of each of the pre-sign-up variables. These include whether the study participant received the email in June or October, as well as their gender, race, and ethnicity. The last column shows the p-values of a pairwise t-test of the mean difference.

$$y_i = \alpha + \beta \text{Effort}_i + X_i' \gamma + \epsilon_i \quad (1)$$

where y is one of our three outcome measures: sign-ups, trips per week, or miles per week for a given participant i . *Effort* is an indicator for being in the *High Effort* condition, and its coefficient, β , measures the average difference between the *High Effort* and *Low Effort* conditions. X_i is a vector of additional covariates that we include in certain specifications (but omit from the main specifications), and γ contains their associated coefficients.⁸

To measure treatment effects, we use the entire sample of participants that we randomized into

⁸We also re-run our results using a Poisson regression model in Section 4.2.5.

Table 3: Estimated Treatment Effects of Effort

	(1) Sign-Ups	(2) Trips per Week	(3) Miles per Week	(4) Sign-Ups	(5) Trips per Week	(6) Miles per Week
High Effort	-0.013*** (0.002)	0.057*** (0.012)	0.414*** (0.089)	-0.015*** (0.003)	0.050*** (0.013)	0.403*** (0.097)
October Email				-0.021*** (0.003)	0.017 (0.017)	0.061 (0.111)
Female				-0.000 (0.003)	0.017 (0.012)	0.137 (0.090)
<i>Race and Ethnicity</i>						
American Indian				-0.042*** (0.005)	-0.062*** (0.013)	-0.465*** (0.092)
Asian				-0.005 (0.007)	-0.003 (0.026)	-0.067 (0.193)
Black				-0.009 (0.010)	0.135 (0.183)	0.641 (1.024)
Hispanic				-0.016*** (0.005)	-0.036*** (0.009)	-0.226* (0.118)
Two or More				-0.060*** (0.003)	-0.036*** (0.006)	-0.297*** (0.057)
Constant	0.051*** (0.002)	0.034*** (0.003)	0.259*** (0.034)	0.061*** (0.003)	0.019 (0.014)	0.159* (0.088)
Observations	27,227	27,227	27,227	23,115	23,115	23,115

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a bivariate OLS model with an indicator for being in the *High Effort* condition as the independent variable. Columns 1 and 4 estimate the effect on an indicator for having signed up to the carpool platform. Columns 2, 3, 5, and 6 estimate the treatment effects on the number of trips taken per week on the platform and the miles driven per week. *October Email* is an indicator variable for receiving the treatment email in October, *Female* is an indicator for being female, and *Race and Ethnicity* show indicator variables for each described race or ethnicity, with White as the base group.

treatment conditions, giving us the ITT effect. We also conduct a boundary exercise using the sample conditional on sign-up in Section 4.2.4.

4.2.1 Effects of Effort on Initial Sign-Ups

To understand the impact of modest effort during the sign-up process on follow-through, we first analyze the effect of effort on signing up to the carpool platform. Results can be found in the first column of Table 3. We find that a more effortful sign-up process reduces the likelihood of signing up to the platform: participants in the *High Effort* condition were 1.3 percentage points less likely to sign up, off of a baseline 5.1% sign-up rate for the *Low Effort* condition ($p < 0.01$). This corresponds to a 25% decrease in sign-up rates when the sign-up process required more effort.

4.2.2 Effects of Effort on Follow-Through

Next, we analyze the effect of effort in the sign-up process on follow-through using the number of trips taken per week and the number of miles driven per week on the carpool platform during our study period. Results are shown in columns two and three of Table 3. Participants who were randomly assigned to the *High Effort* condition take an average of 0.057 more trips per week and carpool about 0.41 miles more per week than those assigned to the *Low Effort* condition, off of baseline usage rates of 0.034 and 0.26, respectively ($p < 0.01$). We therefore find evidence for the buy-in effect.

Looked at differently, the *High Effort* condition took more trips overall, despite there being fewer participants who signed up to the platform: while 694 signed up from the *Low Effort* condition, only 511 signed up from the *High Effort* condition. During our 124-day study period, we observed 9,147 total trips; of these, the *High Effort* condition took 5,106 trips, while the *Low Effort* condition took 4,311, meaning that the *High Effort* condition took 795 more trips overall than the *Low Effort* condition. The drop off in initial sign-up rates was overcome by the increase in usage rates.

Finally, columns four through six show that these results are robust to controlling for whether the participant received the treatment email in June or October, as well as for gender and race.

4.2.3 Intent-to-Treat Estimate as Evidence of a Buy-in Effect

How can we know that these results are not simply evidence of a selection effect, where only the most motivated people signed up in the *High Effort* condition? If the more intensive usage in the *High Effort* condition was only due to a selection effect, we would expect the ITT result—the impact of being assigned to the *High Effort* condition on the number of trips taken—to result in fewer trips per week than being assigned to the *Low Effort* condition, and thus analogously, to result in fewer overall trips taken. This is because the *High Effort* condition would lose some participants who would have signed up under easier conditions, and these lost participants could only have taken zero or a positive numbers of trips, resulting in an overall reduction in total trip-taking for the *High Effort* condition.

To see this, consider that our sample of participants can be divided into three groups: the always-takers, who sign up to the platform regardless of their condition assignment; the never-

takers, who never sign up to the platform; and the complacents,⁹ who sign up when it is easy to do so in the *Low Effort* condition, but do not sign up when the process is more difficult. By having the same number of always-takers in both conditions, but losing the complacents in the *High Effort* condition—which we see through the reduced number of sign-ups—we would expect to also see fewer trips taken in the *High Effort* condition than in the *Low Effort* condition in the ITT analysis. This is because the addition of the complacents can only increase the number of trips taken, since they must take either zero or positive numbers of trips. However, what we find is that the number of trips increases, suggesting that the always-takers in the *High Effort* condition are using the platform more intensively, in a way that compensates for the reduction in the extensive margin through the loss of the complacents that drop out during the sign-up process.

4.2.4 Boundary Exercise on the Buy-in Effect

To investigate how much our results might be driven by selection (i.e., the idea that the *High Effort* condition simply filters out less motivated participants) we conduct a boundary exercise. In particular, we ask, how much do the most motivated participants from the *Low Effort* condition use the platform compared to participants in the *High Effort* condition?

To investigate this question, we remove the 25% of people in the *Low Effort* condition who used the platform the least. Since this is the difference in the sign-up rate between the two conditions, we thus assume a “worst-case” scenario where the people who signed up when it was easy but did not sign up when it was hard were the least motivated participants. Therefore, we are now comparing the follow-through of the 3.8% of participants with the highest usage in the *Low Effort* condition with the full 3.8% of participants in the *High Effort* condition. If selection accounts for the treatment effect, we would expect that limiting the *Low Effort* sample to the most motivated participants would eliminate the observed difference in usage.

These results are shown in Table 4. Columns (1) and (2) present the effect of effort on follow-through, conditional on having signed up, for the full sample. We find that participants in the *High Effort* condition took 1.7 more trips per week and 12.8 more miles per week on average ($p < 0.01$). Looking at columns (3) and (4), the modified sample where the least motivated participants of the *Low Effort* condition are removed, we find that the treatment effect is only somewhat reduced:

⁹We follow (Fowlie et al., 2021) in using the term complacents rather than compliers, as is traditional in the LATE framework literature, because compliers implies active adherence to a treatment condition, but in our study, these individuals drop out in the *High Effort* condition when it requires more effort.

Table 4: Boundary Exercise: Estimated Treatment Effects for the Most Motivated

	Full Sample		Modified Sample	
	(1)	(2)	(3)	(4)
	Trips per Week	Miles per Week	Trips per Week	Miles per Week
High Effort	1.73*** (0.300)	12.76*** (2.146)	1.51*** (0.305)	11.06*** (2.216)
Constant	0.67*** (0.063)	5.11*** (0.649)	0.89*** (0.082)	6.81*** (0.854)
Observations	1,205	1,205	1,031	1,031

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a bivariate OLS model with an indicator for being in the *High Effort* condition as the independent variable, conditional on having signed up. Columns 1 and 2 estimate the effect using the main sample, whereas Columns 3 and 4 estimate the effect on a modified sample that removes the 25% of participants in the *Low Effort* condition with the lowest number of trips per week and miles per week, respectively.

participants in the *High Effort* condition took 1.5 more trips per week and 11.1 more miles per week on average ($p < 0.01$). In other words, using this boundary exercise, we find that selection constitutes a part but not most of the buy-in effect that we observe in this setting.

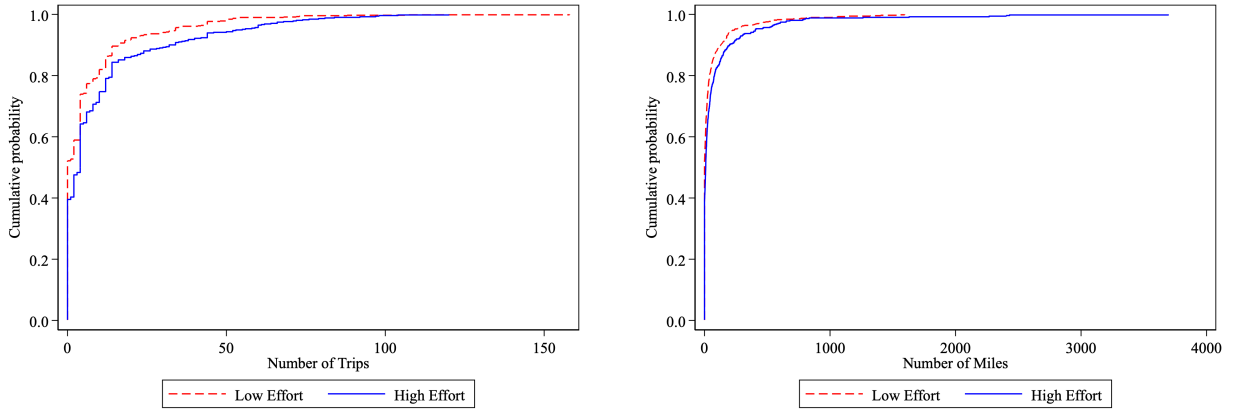
4.2.5 Robustness Checks and Further Analysis

We run a series of robustness checks and additional analyses to view our results using a different regression model, investigate the presence of outliers, analyze the impact of study wave, and look at follow-up data.

Results using a Poisson regression model. First, we re-run our main results using a Poisson regression model to account for the discrete, non-negative nature of the count outcome measures (see results in Appendix Table A1). The direction and significance of the results remain unchanged.

Winsorized Results. We also investigate the sensitivity of our results to the presence of outliers. We re-run our regression models under increasingly restrictive one-sided winsorization of usage at the 99th, 95th, and 90th percentiles of usage for participants who signed up. These percentiles are 13.3, 5.7, and 3.8 trips per week, and 23.6, 48.0, and 174.1 miles per week, respectively. These results are shown in Appendix Table A2. Although the coefficient estimates reported in these columns are smaller in magnitude than in our un-winsorized sample, as expected, they remain significant at each level of winsorization.

Figure 4: Cumulative Distribution Functions by Effort



Note: $n = 1,205$. These figures plot the CDFs of number of trips and number of miles for participants who signed up to the carpool platform. Formal Kolmogorov-Smirnov tests of the equality of treatment and control distributions reject the null for each ($p < 0.01$).

CDFs of Results. We further investigate the presence of outliers in Figure 4, by plotting the cumulative distribution functions (CDFs) of the number of trips and miles on the carpool platform for those who signed up. For both outcome measures, we see a consistent difference in intensity of usage across the distribution, instead of only in the lower or upper tails. This provides further evidence that outliers were not producing the differences we observed across conditions. We confirm this formally with non-parametric Kolmogorov-Smirnov tests of the null hypothesis that the treatment distribution was drawn from the control distribution, which we reject ($p < 0.01$).

Study Wave. To test whether the buy-in effect is consistent across both waves of data collection, we adjust our main regression models for the month of email receipt. These results are presented in Table 5.

The first four columns use the full sample and include an indicator variable for receiving the study email in October (Columns 1 and 2) and interacting this variable with condition assignment (Columns 3 and 4). In all specifications, the treatment effect is significant and similar in size to our main treatment results. Receiving the email in October is associated with directionally more usage. However, the interaction term between *High Effort* and *October* is not significant, indicating that the effect of the treatment condition did not vary by wave.

Table 5: Estimated Treatment Effects of Effort By Study Wave

	Full Sample				June Only		October Only	
	(1) Trips Per Week	(2) Miles Per Week	(3) Trips Per Week	(4) Miles Per Week	(5) Trips Per Week	(6) Miles Per Week	(7) Trips Per Week	(8) Miles Per Week
High Effort	0.056*** (0.012)	0.413*** (0.089)	0.062*** (0.009)	0.459*** (0.089)	0.062*** (0.009)	0.459*** (0.089)	0.047 (0.029)	0.337* (0.186)
October Email	0.017 (0.015)	0.067 (0.103)	0.025*** (0.008)	0.128 (0.082)				
High Effort \times October Email			-0.016 (0.030)	-0.122 (0.206)				
Constant	0.028*** (0.006)	0.235*** (0.046)	0.025*** (0.003)	0.212*** (0.032)	0.025*** (0.003)	0.212*** (0.032)	0.050*** (0.008)	0.340*** (0.075)
Observations	27227	27227	27227	27227	17036	17036	10191	10191

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a multivariate OLS model with the number of trips per week or miles per week as the outcome variable. *High Effort* is an indicator for being assigned the more effortful sign-up process. *October Email* is an indicator for receiving the treatment email in October. Columns 1-4 estimate the effect for the full sample, while Columns 5-6 and 7-8 estimate it using only the sample who received the treatment email in June or in October, respectively.

Columns 5-6 and 7-8 run the main regression using only the June sample and the October sample, respectively. The June results ($n = 17,036$) remain robust. The October results ($n=10,191$) estimate a similar but smaller treatment effect, and this estimate is not significant. This indicates that the effect was weaker for the October sample, potentially because we did not have the statistical power to detect the effect due to the shorter time frame, smaller sample size, and smaller percentage who signed up for the platform (3.1%, or 314 people).

8-Month Follow-Up. Our main study ended in October, as we expected the U.S. holiday season to disrupt commuting patterns. However, ODOT provided us with follow-up data through February 29, 2020, right before the COVID-19 pandemic kept people at home, allowing us to observe whether the treatment effects persisted in the longer run. Results are shown in Table 6.

Table 6: Estimated Treatment Effects of Effort at 8 Months

	(1) Sign-Ups	(2) Trips per Week	(3) Miles per Week
High Effort	-0.013*** (0.003)	0.010*** (0.004)	0.068 (0.049)
Constant	0.052*** (0.002)	0.031*** (0.002)	0.263*** (0.035)
Observations	27,227	27,227	27,227

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a bivariate OLS model with an indicator for being in the *High Effort* condition as the independent variable, with all outcome variables calculated through February 29, 2020. Column 1 estimates the effect on an indicator for having signed up to the carpool platform. Columns 2 and 3 estimate the treatment effects on the number of trips taken and miles driven per week on the platform.

At 8 months, we see the same treatment effects on sign-ups, and we also see persistent and significant treatment effects on the number of trips taken per week. These effect sizes are five times lower than in our main study period but show those in the *High Effort* condition continuing to log trips at a higher rate. However, the miles per week measure, while still positive, is no longer significant, making it difficult to interpret the persistence of these effects.

5 Study 2: Online Experiment with Work Tasks

While Study 1 provides initial evidence for the buy-in effect in a real-world setting, several questions remain about generalizability and mechanisms. First, carpooling is a prosocial and interdependent behavior—it has environmental benefits and requires coordination with other people. These characteristics raise questions about whether the buy-in effect can extend to other contexts. Second, the mechanisms underlying the result from Study 1 remain unclear: does initial effort increase follow-through by increasing the perceived value of the behavior or by improving recall?

To address these questions, we conducted a second pre-registered experiment in a controlled online setting to test the replicability of the main effect and to formally test one plausible psychological mechanism.

5.1 Study Design

We run Study 2 on Amazon Mechanical Turk (MTurk) where we offer online study participants a Human Intelligence Task (or HIT, the basic unit of jobs on MTurk) consisting of a set of transcription tasks. Upon completion of the tasks, participants are offered the chance to sign up for a return opportunity, a new HIT that will be offered the next day with a second set of transcription tasks. We measure follow-through using the return rate, effort, and accuracy in this second HIT.¹⁰ The study is described below. For full study instructions, see Appendix B.

Transcription Tasks. In each HIT, participants complete a set of three transcription tasks. These are adapted from [Augenblick, Niederle and Sprenger \(2015\)](#) and [Augenblick and Rabin \(2019\)](#)’s meaningless Greek letters task, in which participants are presented with a row of 35 blurry Greek letters and must choose the correct corresponding Greek letter. An example is presented in Figure 5.¹¹ These tasks are designed to be effortful and unpleasant. Participants are not required to transcribe each letter, but to receive a bonus payment, they must score 80% correct on a transcription task, which corresponds to 28 of the 35 letters.

Sign-Up Process. Directly after completing the transcription tasks in the first HIT, participants are presented with the chance to sign up for the return opportunity, which consists of the same

¹⁰In this way, this study is a “field-in-the-lab” study, in which we incorporate real-world conditions into laboratory experiments ([Dykstra, Forthcoming](#)): the study participants are online task workers who—by choosing whether to participate—are making a labor force participation decision similar to the kind they regularly make.

¹¹The full transcription task materials are generously provided by the authors at <https://faculty.haas.berkeley.edu/ned/>. We use the exact task images by the authors but change how participants select the correct letter. An image of this is provided in Appendix B Figure 7.

Figure 5: Transcription Task Example

Βαβα . χηηΒββλχηηδ . λθ . ηδχΒββββ . δλσηχχαη

number of tasks for the same amount of earnings. Like in Study 1, participants are randomized to two versions of the sign-up process, *Low Effort* or *High Effort*.

In the *Low Effort* condition, participants simply click “yes” to sign up. They are then brought to the end of the study.

In the *High Effort* condition, they must first complete an additional survey to sign up. This additional survey is on the next page and consists of 15 questions that ask about past experiences with transcription tasks and computer work. Consistent with our proposed scope conditions, this friction is goal-relevant (i.e., meaningfully tied to the target behavior) because the survey items focus on the same domain participants engage with in the return opportunity. This survey can be found in Appendix C.

Pre-Registration. We follow a pre-registration that was submitted to the AEA RCT Registry on March 11, 2025 (AEARCTR-0015492). This registration was submitted prior to data collection and specified the hypotheses, treatments, sample size, and outcome measures.

Implementation. Study participants were recruited on MTurk on March 12, 2025. To participate in the study, individuals had to be located in the United States, have completed at least 500 tasks on MTurk, and received a 99% or higher approval rating on average. They were also required to pass a CAPTCHA screen at the beginning of the survey. We included an attention check question that asked participants to select a particular response; all of our analyses are robust to controlling for this attention check. Two MTurkers completed the first HIT twice, so these observations were dropped from the sample.

5.2 Results

We measure the effect of the *High Effort* condition using the same empirical strategy as in our carpooling experiment, using the entire sample of participants from the first HIT to examine the ITT effect on return, effort, and accuracy.

Table 7: Estimated Treatment Effects of Effort

	(1) Sign-Ups	(2) Returns	(3) Num. Entered	(4) Num. Correct
High Effort	-0.040 (0.037)	0.136*** (0.044)	12.927*** (4.278)	11.909*** (3.882)
Constant	0.808*** (0.025)	0.364*** (0.030)	32.256*** (2.910)	22.392*** (2.532)
Observations	496	496	496	496

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a bivariate OLS model with an indicator for being in the *High Effort* condition as the independent variable. Column 1 estimates the effect on an indicator for having signed up to the return opportunity, Column 2 on an indicator for returning, and Columns 3 and 4 on the number of letters transcribed and the number of letters correctly transcribed, respectively.

5.2.1 Effects of Effort on Sign-Ups and Follow-Through

Results can be found in Table 7. First, we estimate the effect of effort during the sign-up process on the number of sign-ups, which is shown in Column 1. Those who completed the more effortful sign-up process were directionally less likely to sign-up — 76.8% in the *High Effort* condition versus 80.8% in the *Low Effort* condition — but this difference was not distinguishable from zero.

Column 2 shows the effect of effort on returning the next day to complete the task. While 36.4% of participants assigned to the *Low Effort* condition complete the return opportunity, 50.0% of participants in the *High Effort* condition return, which is a 37% increase in the return rate ($p < 0.01$). This increase in the return rate also results in a higher number of tasks completed: participants who were assigned to the *High Effort* condition transcribe 13 more letters on average and have 12 more correct on average ($p < 0.01$). Overall, the *Low Effort* condition correctly transcribe 5,598 letters, while the *High Effort* condition correctly transcribe 8,438 letters, despite marginally fewer people signing up in the *High Effort* condition. We therefore show evidence of the buy-in effect in a different setting.

Next, we conduct the same boundary exercise as in Study 1 to test whether selection into the *High Effort* condition can explain the treatment effect. How much do the most motivated participants from the *Low Effort* condition follow through compared to participants in the *High*

Effort condition? To do this, we remove the 5% of people in the *Low Effort* condition who entered the lowest number of letters. Since this is the difference in the sign-up rate between the two conditions, we thus assume a “worst-case” scenario where the people who signed up when it was easy but did not sign up when it was hard were the least motivated participants. The full table of these results can be found in Appendix Table A5. Conditional on signing up, participants in the *High Effort* condition transcribe 19 more letters and transcribe 17 more letters correctly on average ($p < 0.01$). When removing the least motivated participants from the *Low Effort* condition, we find that the treatment effect is somewhat reduced: participants in the *High Effort* condition transcribe 17 more letters and transcribe 15.5 more letters correctly on average ($p < 0.01$). Thus, like in Study 1, selection accounts for a portion, but not all of the buy-in effect that we observe.

Finally, we also conduct the same robustness checks and further analyses as in Study 1. First, we re-run our main results using a Poisson regression model to take into account the discrete, non-negative nature of the count outcome variables (see Appendix Table A6). The direction and significance of the results remain unchanged. Second, we investigate the sensitivity of our results to the presence of outliers using one-sided winsorization of follow-through at the 99th, 95th, and 90th percentiles. Our results remain significant (see Appendix Table A7). Third, we further investigate the presence of outliers in Appendix A Figure 1 by plotting the CDFs of the number of letters transcribed and the number of letters correctly transcribed. For both outcomes, we see a consistent difference in the intensity of usage across the distribution, instead of only in the lower or upper tails. We confirm this formally with non-parametric Kolmogorov-Smirnov tests of the null hypothesis that the treatment distribution was drawn from the control distribution, which we reject ($p < 0.05$).

5.2.2 Additional Treatment Conditions

In addition to the main *High Effort* and *Low Effort* condition assignments, we also cross-randomize participants into a set of additional conditions to test the attentional mechanism as well as the generalizability conditions—does the buy-in effect require the target behavior to be prosocial and dependent on other people, as was the case in Study 1?

The first generalizability condition that we study is whether the buy-in effect requires the target behavior to be prosocial, since carpooling benefits the environment and helps community members

commute more affordably. Prior research suggests that people are more willing to persist in effortful tasks when outcomes support a prosocial cause rather than a personal reward (Imas, 2014). Similarly, individuals are willing to forgo financial incentives to signal other-oriented motivations in prosocial contexts (Kirgios et al., 2020). We include a *Volunteer* condition, where participants are told that their earnings in the return opportunity will be donated to a charity of their choice.

The second generalizability condition that we study is whether the buy-in effect requires the target behavior to be dependent on other people, since carpooling depends on the commitment and participation of other people to show up reliably and on time. Conceptual research suggests that individuals make inferences about the likely quality of an experience based on features of the choice environment (Kamenica, 2008). When sign-up requires effort, individuals may infer that other people who opted in are also committed, reliable, and aligned with the shared goal, increasing confidence in the feasibility of the target behavior and follow-through. To test this possibility, we incorporate the *Social: Quantity* and *Social: Quality* conditions into Study 2. In *Social: Quantity*, participants are informed that their earnings in the return opportunity will depend on the number of other participants who complete the return opportunity. In *Social: Quality*, participants are informed that their earnings in the return opportunity will depend on the accuracy of the other people who return. In this way, *Social: Quantity* reflects the fact that higher numbers of people on a carpooling platform are better, and *Social: Quality* reflects the fact that the quality of people you carpool with matter.

Finally, we explicitly test the role of recall, one of the mechanisms we describe in our conceptual framework in Section 3. To do this, we include the *Reminder* condition, wherein participants receive a reminder email on the following day to remind them that the return opportunity is now available.

5.2.3 Generalizability Conditions

The results of the generalizability conditions are presented in Table 8. The *Volunteer* condition has a strongly negative impact on the return rate and the number of correctly transcribed letters. Those in the *Volunteer* condition are 23.3 p.p. less likely to return, which is almost double the magnitude of the *High Effort* result ($p < 0.01$). There is no interaction effect between the *High Effort* and the *Volunteer* conditions, indicating that the *High Effort* condition has an impact regardless of the prosociality of the behavior.

Table 8: Estimated Treatment Effects by Generalizability Condition

	Volunteer				Social			
	(1) Returns	(2) Returns	(3) Num. Correct	(4) Num. Correct	(5) Returns	(6) Returns	(7) Num. Correct	(8) Num. Correct
High Effort	0.122*** (0.043)	0.173*** (0.062)	11.293*** (3.864)	15.520*** (5.571)	0.136*** (0.044)	0.117 (0.076)	12.032*** (3.885)	10.870 (6.959)
Volunteer	-0.233*** (0.043)	-0.181*** (0.060)	-10.154*** (3.855)	-5.881 (5.080)				
High Effort \times Volunteer		-0.105 (0.086)		-8.638 (7.718)				
Quantity					-0.042 (0.054)	-0.042 (0.074)	-5.015 (4.770)	-6.017 (6.182)
Quality					0.049 (0.054)	0.019 (0.075)	-1.649 (4.792)	-2.360 (6.298)
High Effort \times Quantity						-0.000 (0.108)		2.043 (9.577)
High Effort \times Quality						0.059 (0.108)		1.483 (9.635)
Constant	0.485*** (0.039)	0.458*** (0.046)	27.672*** (3.303)	25.450*** (3.787)	0.361*** (0.043)	0.371*** (0.052)	24.518*** (3.801)	25.067*** (4.511)
Observations	496	496	496	496	496	496	496	496

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a multivariate OLS model. Columns 1, 2 and 5, 6 use an indicator for returning as the dependent variable. Columns 3, 4 and 7, 8 use the number of variables correctly transcribed as the dependent variable. *Volunteer* is an indicator for the earnings in the return opportunity being donated to charity. *Quantity* and *Quality* are indicators for the return opportunity payment being affected by the number of other people who returned and for the accuracy of other people who returned, respectively.

The *Social: Quantity* and *Social: Quality* conditions did not have an impact on the return rate or the number of correctly transcribed letters. In addition, there was no interaction between these conditions and the *High Effort* condition, indicating that the social component of the return opportunity did not impact the effect of the *High Effort* condition. These results indicate that the buy-in effect can also exist in settings that are not prosocial, and furthermore, that the effect does not depend on the behaviors or perceived quality of other participants.

5.2.4 Reminder Condition

The results of the *Reminder* condition testing the recall mechanism are presented in Table 9. This condition did not have an impact on returns or the number of correctly transcribed letters.¹²¹³ An interpretation of this result is that the buy-in effect in this setting works primarily through the perceived value channel, which we discuss more in Section 6.2.

Table 9: Estimated Treatment Effects in the Reminder Condition

	(1) Returns	(2) Returns	(3) Num. Correct	(4) Num. Correct
High Effort	0.200*** (0.049)	0.214*** (0.070)	16.869*** (4.606)	16.233** (6.775)
Reminder	-0.032 (0.049)	-0.019 (0.070)	-8.311* (4.587)	-8.926 (6.023)
High Effort \times Reminder		-0.027 (0.099)		1.271 (9.226)
Constant	0.467*** (0.043)	0.460*** (0.050)	31.910*** (3.930)	32.220*** (4.511)
Observations	391	391	391	391

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Each column presents the results of a multivariate OLS model. Columns 1 and 2 use an indicator for returning as the dependent variable. Columns 3 and 4 use the number of variables correctly transcribed as the dependent variable. *Reminder* is an indicator for receiving a reminder email about the return opportunity.

¹²Note that the sample is smaller because participants were randomized into the reminder condition after signing up for the return opportunity.

¹³Because we were puzzled by this non-result, we ran a follow-up survey the next day for study participants in the *Reminder* condition to find out whether they saw the email ($n = 50$). All but one reported seeing it.

6 Discussion

We proposed and tested the buy-in effect, whereby introducing modest, goal-relevant effort during an initial sign-up process can increase follow-through. We tested this possibility across two pre-specified studies using diverse contexts and participants. Study 1 was a large-scale field experiment in carpooling in which participants who were randomly assigned to a more effortful sign-up process took more trips overall, despite fewer people signing up. Study 2 was a pre-registered online experiment in which participants who were randomly assigned to a more effortful sign-up process were more likely to return the next day and complete more work overall. These results challenge the assumption that reducing friction universally encourages behavior change, and instead suggest that friction can enhance follow-through under specific conditions. As we highlight below, future research is needed to directly test when and where these benefits are most likely to emerge.

6.1 What is the Role of Heterogeneity?

While our studies were designed to estimate average effects on follow-through using ITT analyses, understanding how heterogeneity in the target population might affect the buy-in effect is important for policymakers, and aligns with recent work documenting the importance of understanding heterogeneity in behavioral science interventions (Bryan, Tipton and Yeager, 2021). We therefore present evidence on heterogeneity of the buy-in effect by observable characteristics and baseline motivation. We also discuss the importance of the composition of the target population by the number of complacents.

By Baseline Motivation. Using data from Study 2, we analyze the relationship between baseline motivation and follow-through. To do this, we use effort in the first HIT as a measure of baseline motivation before participants signed up for the return opportunity. These results can be found in Appendix Table A8.

Panel A presents results using the number of letters entered in the first HIT as a measure of baseline motivation. The measure is not associated with a higher return rate but is associated with a higher number of correctly transcribed letters in the return task; this can be interpreted as a mechanical relationship between the number of letters entered and the number of letters correct, since baseline number entered is highly predictive of the number entered in the return task and also of the number correct. There is no interaction between *High Effort* and the number of letters

entered at baseline.

Panel B presents results using the number of letters correctly transcribed in the first HIT as a measure of baseline motivation. Here, we do see a significant relationship between this measure of motivation and the return rate, although this effect is not moderated by the *High Effort* treatment. We also see a relationship between motivation and the number of correct letters in the second task, and this time, this relationship is moderated by the *High Effort* treatment condition: one more baseline correct letter is associated with 0.45 more letters correct in the second task on average, and those in the *High Effort* treatment have an additional 0.16 letters correct in the second task ($p < 0.05$).

However, while this analysis provides initial, exploratory evidence that the buy-in effect may sometimes be more effective for people with higher baseline motivation, more work is needed to test whether this finding is robust and generalizes to other settings and populations.

By Composition of the Target Population. Finally, the composition of the target population could be a crucial determinant of whether the buy-in effect is effective in a certain setting. Specifically, the proportion of complacent individuals in the population could be an important factor that shapes the impact of the buy-in effect. Across our studies, the proportion of complacent individuals was relatively low (25% in Study 1 and 5% in Study 2). As a result, the increase in follow-through behavior that was encouraged by the buy-in effect was large enough to compensate for the loss of these individuals. However, this will not be true in every context, where the level of drop-outs might be high enough as to no longer be compensated by the increase in follow-through. Thus, the composition of the target population is important for program designers to consider when deciding whether to use initial effort as a policy tool.

6.2 What is the Role of Perceived Value vs. Recall?

Understanding the psychological processes underlying the buy-in effect has important theoretical and practical implications. We described the two broad channels of mechanisms that could explain the buy-in effect: increasing the perceived value of the behavior through psychological ownership, sunk-cost effects, or effort justification (Arkes and Blumer, 1985; Shu and Peck, 2011; Norton, Mochon and Ariely, 2012), and increasing attention and recall (Chen et al., 2015; Rogers and Milkman, 2016).

Study 2 included a direct test of the attention-based mechanism using the *Reminder* condition. If the effect was driven primarily by enhanced recall, then reminders should have substituted for initial effort, reducing differences between the *High Effort* and *Low Effort* conditions. Instead, reminders had no effect on return behavior, and did not interact with the effort manipulation. These results suggest that the buy-in effect is more consistent with a value-based rather than an attention-based mechanism in this setting. There are many reasons why this might be the case. One interpretation is that MTurk workers already decided whether or not to work that day, so the reminder email did not change their behavior. Instead, when MTurk workers logged in, and were deciding between different HITs, the *High Effort* condition affected whether or not they thought the return opportunity was worth completing.

While these findings allow us to rule out a recall-based account in the online context we studied, these results do not isolate which value-based process, such as psychological ownership, identity alignment, or sunk-cost perceptions, underpins our findings. Ideally, we would directly measure such constructs after participants faced a more effortful sign-up process. However, doing so in the same study as measuring behavioral outcomes poses a methodological challenge, because any additional measures included during the sign-up process would increase the amount of effort in the *Low Effort* condition, contaminating the manipulation and potentially reducing any between-condition differences. While the observed pattern of increased return rates in Study 2 supports a value-based interpretation of our results, we leave open the question of which specific cognitive and motivational processes underpin the buy-in effect. Additional research using alternative methodological approaches is needed to directly measure and differentiate between these mechanisms.

6.3 Could Higher Observability and Identifiability Explain our Results?

An alternative explanation is differential observability or identifiability across conditions. Identifiability and reputational concerns can influence behavior in contexts where actions are visible to relevant peers (Yoeli et al., 2013). Thus, participants who provide more information during a sign-up process could feel more observable and accountable, leading to increased behavioral compliance to maintain their reputation or avoid social sanctions.

Several features of our research design make this explanation unlikely. In Study 1, participants interacted with a third-party organization, ODOT, that did not share personal information with

their employers, as outlined in their Privacy Policy as part of their Terms & Conditions. In addition, all participants had previously provided identical personal information (i.e., zip codes, organization details) to the old platform. While it is possible that this information might have been more salient in the *High Effort* condition due to the re-entry requirement, the actual level of identifiability was identical across conditions.

In Study 2, this alternative explanation is also unlikely. The online task was completed anonymously through MTurk, where workers interact using anonymous worker IDs rather than personal identifiers. Participants' decisions to return and their task performance were entirely private, with no mechanism for individual accountability or reputational feedback. We did include two treatment conditions in which participants' bonus payments were dependent on other people: the *Social: Quantity* condition and the *Social: Quality* condition. However, in these cases, participants' performance remained anonymous and individually unobservable, and they were told that their bonus payments depended on the overall number of people who returned and the overall accuracy of the other people, respectively. Furthermore, in both conditions, participants were explicitly told that their bonus would depend on others' behavior, and that their own completion or performance would have limited impact. Thus, even when social framing was made salient, reputational concerns were minimized, making it unlikely that peer influence or reputational concerns were responsible for the follow-through behavior that we observed in Study 2.

6.4 Limitations & Future Directions

6.4.1 Scope Conditions

We theorize that the buy-in effect is most likely to emerge under certain scope conditions: when the target behavior requires follow-through, when the initial effort and target behavior are voluntary, and when the friction is modest and goal-relevant. We describe these scope conditions and the reasons behind them in Section 3. However, a limitation of these studies is that we do not directly test these scope conditions.

First, while we propose that the buy-in effect will most likely emerge in settings that require follow-through, it remains an open question as to whether the buy-in effect operates similarly for lower-effort or one-time follow-through behaviors, such as casting a single vote in a referendum or receiving a seasonal flu shot. It is theoretically plausible for the buy-in effect to emerge with an

effortful one-time behavior, and Study 2 provides initial evidence for this possibility. However, the motivational benefits of upfront effort are likely to be strongest when there is a calibrated match between the initial effort required and the demands of the target behavior (Inzlicht, Shenhav and Olivola, 2018). When both the buy-in task and the target behavior require modest effort, this initial effort could serve as a credible signal of commitment and better prepare people psychologically for the effortful follow-through that is required. In contrast, requiring significant upfront effort for a simple, or one-time behavior, may feel disproportionate and therefore be unlikely to increase follow-through.

From a policy perspective, this research suggests different intervention strategies for different behavioral contexts. For more effortful behaviors like medication adherence or exercise routines, buy-in mechanisms may be particularly valuable because they could help people overcome motivational barriers. For simpler, one-time behaviors like flu vaccinations, policy makers might achieve better population-level outcomes by removing barriers rather than adding them, since the primary challenge might be initial participation rather than sustained motivation. Clarifying the precise conditions under which initial friction serves as a motivational tool will be critical for refining empirical predictions and guiding the design of effective interventions.

Second, we propose that the buy-in effect will most likely emerge in voluntary settings like our two studies, where participants can decide whether to continue at any point. We hypothesize that the voluntary nature of the task creates conditions for meaningful self-signaling to occur, where initial effort may provide relevant information about personal commitment and motivation. This proposition aligns with research on self-determination theory, which emphasizes the importance of autonomy for intrinsic motivation and follow-through (Deci and Ryan, 2000). This proposition is also consistent with research showing that costly initial prosocial acts promote follow-through by signaling a prosocial identity that people then seek to uphold (Gneezy et al., 2012). However, an avenue for future research is to test whether the buy-in effect emerges in contexts where participation is mandatory instead of voluntary. Understanding this scope condition has practical implications for designing interventions in organizational, educational, and policy contexts where participation may not always feel freely chosen.

Third, we propose that the buy-in effect will most likely emerge when the friction is modest and commensurate to the target behavior, because excessive friction may significantly reduce initial

sign up or create a backlash effect. However, we refrain from defining what modest means, because this will vary by setting and target behavior. Future work could experimentally vary the amount of friction across contexts to directly test this question.

Finally, we propose that the buy-in effect will most likely emerge when the effort during sign-up is relevant to the target behavior. While past research suggests overly burdensome or non-goal-relevant friction may deter follow-through (Kruger et al., 2004; Bertrand et al., 2010; Herd and Moynihan, 2018; Sunstein, 2021), it is an open question as to whether similar effects would emerge if the friction was seen as arbitrary, bureaucratic, or unrelated to the target behavior. Future research should experimentally vary the content and perceived meaningfulness of the friction to examine whether the buy-in effect is specific to goal-aligned effort or whether it generalizes to more incidental or externally imposed effort costs.

6.4.2 Unobserved Behavior

In both studies, we could not observe behavior outside our platform contexts, creating three key limitations. First, without data on users' total transportation behavior in Study 1, we cannot assess whether observed differences in carpooling represent changes in carpooling as a fraction of total trips and miles taken by the person, meaning that our results might be platform-specific rather than reflecting behavior change. For example, participants might have increased the amount of driving they did overall or sought out alternative means of carpooling. We believe that it is unlikely that participants were carpooling through other means because ODOT's Get There Connect was one of the region's primary carpool-matching services. However, we cannot rule out the possibility of substitution outside of the platform. Similarly, in Study 2, since we did not observe what other tasks the study participants were completing, we cannot say that they increased their overall rate of work. This limitation is relevant given research showing that interventions can redistribute behavior across channels rather than increasing overall engagement (Tiefenbeck et al., 2013; Blanken, Van De Ven and Zeelenberg, 2015; Dolan and Galizzi, 2015).

Second, in Study 1, since users must self-log or confirm trips on the platform, it is possible for people to take carpooling trips that were matched through the platform that we did not observe. Although we believe this is unlikely because ODOT and Oregon employers offer various incentives for logging trips through the platform, we cannot rule out the possibility that people are taking

trips that they did not log. In this case, our results would be conservative estimates of the amount of carpooling that happened because of the platform.

Third, and relatedly, we cannot rule out the possibility that the *High Effort* condition in Study 1 influenced participants’ diligence in logging carpooling trips without necessarily changing their actual carpooling behavior. Although we believe that differential logging is unlikely because both treatment groups received identical instructions and incentives for trip logging, we cannot definitively separate changes in actual carpooling from changes in logging behavior. However, the buy-in effect remains robust in Study 2, where we automatically observe follow-through behavior.

Future work could explore these possibilities by incorporating cross-platform or behavioral-trace data that captures mobility patterns more broadly, including total driving usage, and by testing whether friction increases follow-through within the intended channel or redistributes this behavior across alternative channels, while also using objective measures of carpooling behavior. Understanding these dynamics is important for policy applications where the goal is to increase behavior overall, not just within a specific program or platform.

6.4.3 Long-Run Behavior Change

While the buy-in effect could plausibly cause long-run changes in the target behavior, our current studies find limited evidence for this possibility. As described in Section 2, there are theoretical reasons why the buy-in effect might produce more durable results than other interventions that address the intention-action gap. Behavioral interventions often fail to produce persistent effects once the intervention is removed (Nisa et al., 2019), and financial incentives often decay once the incentives are removed (Charness and Gneezy, 2009; Acland and Levy, 2015). In contrast, the buy-in effect could theoretically create more lasting commitment due to using upfront “work” rather than external rewards.

In Study 1, we received follow-up data through 8 months, past the U.S. holiday season. We observed a persistent but small effect of the *High Effort* condition on trips taken per week, but we no longer saw an effect on miles per week. One interpretation of these results is that a small group of people successfully changed their long-run carpooling behavior, while a larger group stopped carpooling altogether. In Study 2, we only examined follow-through behavior the next day, providing no additional evidence for long-run effects. Studying longer-run impacts of the buy-in effect

and comparing these results to other behavior change strategies would be an interesting avenue for future research.

7 Conclusion

The intention-action gap represents a fundamental challenge across domains ranging from personal finance to healthy eating to recycling and carpooling. Our research conceptualizes and presents evidence for the buy-in effect, a behavioral phenomenon in which the intentional design of initial effort increases follow-through. Across two experiments involving over 27,000 participants, we find that individuals who complete more effortful sign-up processes show greater follow-through than those who face minimal friction to sign up.

Our findings suggest that policymakers and practitioners should reconsider the universal application of simplification strategies, particularly in domains that require follow-through. While reducing friction is appropriate for “set-and-forget” behaviors, and in contexts where initial enrollment is the critical behavior of interest, modest increases in effort may be beneficial when follow-through is essential. The key insight that emerges from this work is not that friction is inherently good or bad, but rather that the effect of friction depends on the setting, the nature of the friction, and the psychological mechanisms it activates. Goal-aligned friction that increases the perceived value of a behavior may support follow-through. Future work that tests the limits of these scope conditions directly would help to better understand when and how to implement this friction.

This more nuanced view has implications for designing interventions in sustainability, health, civic engagement, and other domains where follow-through is critical. Rather than defaulting to simplification, program designers might consider incorporating modest, meaningful effort requirements that enhance buy-in while preserving accessibility.

References

- Acland, Dan, and Matthew Levy.** 2015. “Naiveté, Projection Bias, and Habit Formation in Gym Attendance.” *American Economic Journal: Applied Economics*, 7(4): 1–35.
- Aknin, Lara B., and Ashley V. Whillans.** 2021. “Helping and Happiness: A Review and Guide for Public Policy.” *Social Issues and Policy Review*, 15(1): 3–50.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi.** 2016. “Self-Targeting: Evidence from a Field Experiment in Indonesia.” *Journal of Political Economy*, 124(2): 371–427.
- Arkes, Hal R., and Catherine Blumer.** 1985. “The Psychology of Sunk Cost.” *Organizational Behavior and Human Decision Processes*, 35(1): 124–140.
- Ashraf, Nava, James Berry, and Jesse M. Shapiro.** 2010. “Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia.” *American Economic Review*, 100(5): 2383–2413.
- Augenblick, Ned, and Matthew Rabin.** 2019. “An Experiment on Time Preference and Misprediction in Unpleasant Tasks.” *The Review of Economic Studies*, 86(3): 941–975.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger.** 2015. “Working Over Time: Dynamic Inconsistency in Real Effort Tasks.” *The Quarterly Journal of Economics*, 130(3): 1067–1115.
- Bénabou, Roland, and Jean Tirole.** 2003*a*. “Intrinsic and Extrinsic Motivation.” *Review of Economic Studies*, 70(3): 489–520.
- Bénabou, Roland, and Jean Tirole.** 2003*b*. “Self-Knowledge and Self-Regulation: An Economic Approach.” *The Psychology of Economic Decisions*, 1: 137–167.
- Bénabou, Roland, and Jean Tirole.** 2006. “Incentives and Prosocial Behavior.” *American Economic Review*, 96(5): 1652–1678.
- 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.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman.** 2010. “What’s Advertising Content Worth? Evidence From a Consumer Credit Marketing Field Experiment.” *The Quarterly Journal of Economics*, 125(1): 263–306.
- Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir.** 2006. “Behavioral Economics and Marketing in Aid of Decision Making Among the Poor.” *Journal of Public Policy & Marketing*, 25(1): 8–23.
- Beshears, John, James J. Choi, Christopher Harris, David Laibson, Brigitte C. Madrian, and Jung Sakong.** 2020. “Which Early Withdrawal Penalty Attracts the Most Deposits to a Commitment Savings Account?” *Journal of Public Economics*, 183: 104144.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian.** 2013. “Simplification and Saving.” *Journal of Economic Behavior & Organization*, 95: 130–145.
- Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Katherine L. Milkman.** 2021. “Building a Commitment Device That Works.” *Management Science*, 67(7): 4386–4400.
- Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu.** 2012. “The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment*.” *Quarterly Journal of Economics*, 127(3): 1205–1242.
- Bhargava, Saurabh, and Dayanand Manoli.** 2015. “Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment.” *American Economic Review*, 105(11): 3489–3529.
- Bhattacharya, Jay, Alan M. Garber, and Jeremy D. Goldhaber-Fiebert.** 2015. “Nudges in Exercise Commitment Contracts: A Case Study in Singapore.” *Health Affairs*, 34(11): 1980–1985.
- Blanken, Irene, Niels Van De Ven, and Marcel Zeelenberg.** 2015. “A Meta-Analytic Review of Moral Licensing.” *Personality and Social Psychology Bulletin*, 41(4): 540–558.

- Braconnier, Céline, Jean-Yves Dormagen, and Vincent Pons.** 2017. “Voter Registration Costs and Disenfranchisement: Experimental Evidence from France.” *American Political Science Review*, 111(3): 584–604.
- Bryan, Christopher J., David S. Yeager, Cintia P. Hinojosa, Aimee Chabot, Holly Bergen, Mari Kawamura, and Fred Steubing.** 2016. “Harnessing Adolescent Values to Motivate Healthier Eating.” *Proceedings of the National Academy of Sciences*, 113(39): 10830–10835.
- Bryan, Christopher J., Elizabeth Tipton, and David S. Yeager.** 2021. “Behavioural Science is Unlikely to Change the World Without a Heterogeneity Revolution.” *Nature Human Behaviour*, 5(8): 980–989.
- Bryan, Gharad, Dean Karlan, and Scott Nelson.** 2010. “Commitment Devices.” *Annual Review of Economics*, 2(1): 671–698.
- Buell, Ryan W., and Michael I. Norton.** 2011. “The Labor Illusion: How Operational Transparency Increases Perceived Value.” *Management Science*, 57(9): 1564–1579.
- Cadario, Romain, and Pierre Chandon.** 2020. “Which Healthy Eating Nudges Work Best? A Meta-Analysis of Field Experiments.” *Marketing Science*, 39(3): 465–486.
- Calzolari, Giacomo, and Mattia Nardotto.** 2017. “Effective Reminders.” *Management Science*, 63(9): 2915–2932.
- Cantoni, Enrico, Vincent Pons, and Jerome Schafer.** 2025. “Voting Rules, Turnout, and Economic Policies.” *Annual Review of Economics*, 17.
- Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick.** 2009. “Optimal Defaults and Active Decisions.” *Quarterly Journal of Economics*, 124(4): 1639–1674.
- Charness, Gary, and Uri Gneezy.** 2009. “Incentives to Exercise.” *Econometrica*, 77(3): 909–931.

- Chater, Nick, and George Loewenstein.** 2022. “The i-Frame and the s-Frame: How Focusing on Individual-Level Solutions Has Led Behavioral Public Policy Astray.” *Behavioral and Brain Sciences*, 46: e147.
- Chen, Xing-jie, Ya Wang, Lu-lu Liu, Ji-fang Cui, Ming-yuan Gan, David HK Shum, and Raymond CK Chan.** 2015. “The Effect of Implementation Intention on Prospective Memory: A Systematic and Meta-Analytic Review.” *Psychiatry Research*, 226(1): 14–22.
- Chernev, Alexander, Ulf Boeckenholt, and Joseph Goodman.** 2015. “Choice Overload: A Conceptual Review and Meta-Analysis.” *Journal of Consumer Psychology*, 25(2): 333–358.
- Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick.** 2002. “Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance.” *Tax Policy and the Economy*, 16: 67–113.
- Cohen, Jessica, and Pascaline Dupas.** 2010. “Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment.” *The Quarterly Journal of Economics*, 125(1): 1–45.
- Coppock, Alexander, and Donald P. Green.** 2016. “Is Voting Habit Forming? New Evidence from Experiments and Regression Discontinuities.” *American Journal of Political Science*, 60(4): 1044–1062.
- Daigneault, Pierre-Marc, Mathieu Ouimet, Alexandre Fortin-Chouinard, and Eriole Zita Nonki Tadida.** 2025. “How Effective Are Behavioral Interventions to Increase the Take-Up of Social Benefits? A Systematic Review of Field Experiments.” *Journal of Policy Analysis and Management*, 44(1).
- Dale, Allison, and Aaron Strauss.** 2009. “Don’t Forget to Vote: Text Message Reminders as a Mobilization Tool.” *American Journal of Political Science*, 53(4): 787–804.
- Deci, Edward L., and Richard M. Ryan.** 1985. *Intrinsic Motivation and Self-Determination in Human Behavior*. Springer.
- Deci, Edward L., and Richard M. Ryan.** 2000. “The “What” and “Why” of Goal Pursuits: Human Needs and the Self-Determination of Behavior.” *Psychological Inquiry*, 11(4): 227–268.

- DellaVigna, Stefano, and Devin Pope.** 2018. “What Motivates Effort? Evidence and Expert Forecasts.” *The Review of Economic Studies*, 85(2): 1029–1069.
- DellaVigna, Stefano, and Elizabeth Linos.** 2022. “RCTs to Scale: Comprehensive Evidence From Two Nudge Units.” *Econometrica*, 90(1): 81–116.
- De-loyde, Katie, Mark A. Pilling, Amelia Thornton, Grace Spencer, and Olivia M. Maynard.** 2025. “Promoting Sustainable Diets Using Eco-Labeling and Social Nudges: A Randomised Online Experiment.” *Behavioural Public Policy*, 9(2): 426–442.
- Dolan, Paul, and Matteo M Galizzi.** 2015. “Like Ripples on a Pond: Behavioral Spillovers and their Implications for Research and Policy.” *Journal of Economic Psychology*, 47: 1–16.
- Duckworth, Angela, Katherine Milkman, and David Laibson.** 2018. “Beyond Willpower: Strategies for Reducing Failures of Self-Control.” *Psychological Science in the Public Interest*, 19(3): 102–129.
- Duckworth, Angela L, Tamar Szabó Gendler, and James J Gross.** 2016. “Situational Strategies for Self-Control.” *Perspectives on Psychological Science*, 11(1): 35–55.
- Dykstra, Holly.** Forthcoming. “Patience Across Payday: The Role of Scarcity in Commitment Decisions.” *Journal of Political Economy Microeconomics*.
- Emens, Elizabeth F.** 2021. “Disability Admin: The Invisible Costs of Being Disabled.” *Minnesota Law Review*, 105(5): 2329–2377.
- Fonseca, Miguel A., and Shaun B. Grimshaw.** 2017. “Do Behavioral Nudges in Prepopulated Tax Forms Affect Compliance? Experimental Evidence with Real Taxpayers.” *Journal of Public Policy & Marketing*, 36(2): 213–226.
- Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C. Anna Spurlock, Annika Todd-Blick, and Peter Cappers.** 2021. “Default Effects and Follow-on Behavior: Evidence from an Electricity Pricing Program.” *The Review of Economic Studies*, 88(6): 2886–2934.
- Frey, Erin, and Todd Rogers.** 2014. “Persistence: How Treatment Effects Persist after Interventions Stop.” *Policy Insights from the Behavioral and Brain Sciences*, 1(1): 172–179.

- Gillitzer, Christian, and Peer Ebbesen Skov.** 2018. “The Use of Third-Party Information Reporting for Tax Deductions: Evidence and Implications From Charitable Deductions in Denmark.” *Oxford Economic Papers*, 70(3): 892–916.
- Gneezy, Ayelet, Alex Imas, Amber Brown, Leif D Nelson, and Michael I Norton.** 2012. “Paying to Be Nice: Consistency and Costly Prosocial Behavior.” *Management Science*, 58(1): 179–187.
- Gneezy, Uri, and Aldo Rustichini.** 2000. “Pay Enough or Don’t Pay at All.” *Quarterly Journal of Economics*, 115(3): 791–810.
- Goldin, Jacob, Tatiana Homonoff, Rizwan Javaid, and Brenda Schafer.** 2022. “Tax Filing and Take-Up: Experimental Evidence on Tax Preparation Outreach and Benefit Claiming.” *Journal of Public Economics*, 206: 104550.
- Gollwitzer, Peter M., and Paschal Sheeran.** 2006. “Implementation Intentions and Goal Achievement: A Meta-Analysis of Effects and Processes.” *Advances in Experimental Social Psychology*, 38: 69–119.
- Guyton, John, Pat Langetieg, Day Manoli, Mark Payne, Brenda Schafer, and Michael Sebastiani.** 2017. “Reminders and Recidivism: Using Administrative Data to Characterize Nonfilers and Conduct EITC Outreach.” *American Economic Review*, 107(5): 471–475.
- Herd, Pamela, and Donald P. Moynihan.** 2018. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation.
- Hill, Seth J., and Thad Kousser.** 2016. “Turning Out Unlikely Voters? A Field Experiment in the Top-Two Primary.” *Political Behavior*, 38(2): 413–432.
- Hinojosa, Amanda S, William L Gardner, H Jack Walker, Claudia Coglisier, and Daniel Gullifor.** 2017. “A Review of Cognitive Dissonance Theory in Management Research: Opportunities for Further Development.” *Journal of Management*, 43(1): 170–199.
- Imas, Alex.** 2014. “Working for the “Warm Glow”: On the Benefits and Limits of Prosocial Incentives.” *Journal of Public Economics*, 114: 14–18.

- Inzlicht, Michael, Amitai Shenhav, and Christopher Y. Olivola.** 2018. “The Effort Paradox: Effort is Both Costly and Valued.” *Trends in Cognitive Sciences*, 22(4): 337–349.
- Jachimowicz, Jon, Shannon Duncan, Elke Weber, and Eric Johnson.** 2019. “When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects.” *Behavioural Public Policy*, 3(2): 159–186.
- Kahneman, Daniel.** 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kamenica, Emir.** 2008. “Contextual Inference in Markets: On the Informational Content of Product Lines.” *American Economic Review*, 98(5): 2127–49.
- Kaplan, Jacob.** 2023. “Predict the Race and Gender of a Given Name Using Census and Social Security Administration Data. Version 2.0.1.”
- Keller, Punam A., Bari Harlam, George Loewenstein, and Kevin G. Volpp.** 2011. “Enhanced Active Choice: A New Method to Motivate Behavior Change.” *Journal of Consumer Psychology*, 21(4): 376–383.
- Kirgios, Erika L., Edward H. Chang, Emma E. Levine, Katherine L. Milkman, and Judd B. Kessler.** 2020. “Forgoing Earned Incentives to Signal Pure Motives.” *Proceedings of the National Academy of Sciences*, 117(29): 16891–16897.
- Kristal, Ariella, and Ashley Whillans.** 2020. “What We Can Learn from Five Naturalistic Field Experiments That Failed to Shift Commuter Behaviour.” *Nature Human Behaviour*, 4(2): 169–176.
- Kruger, Justin, Derrick Wirtz, Leaf Van Boven, and Tracey W. Altermatt.** 2004. “The Effort Heuristic.” *Journal of Experimental Social Psychology*, 40: 91–98.
- Laibson, David.** 1997. “Golden Eggs and Hyperbolic Discounting.” *The Quarterly Journal of Economics*, 112(2): 443–477.
- Linos, Elizabeth, Allen Prohovsky, Aparna Ramesh, Jesse Rothstein, and Matthew Unrath.** 2022. “Can Nudges Increase Take-up of the EITC? Evidence from Multiple Field Experiments.” *American Economic Journal: Economic Policy*, 14(4): 432–452.

- Lockhart, Jeffrey, Molly King, and Christin Munsch.** 2023. “Name-Based Demographic Inference and the Unequal Distribution of Misrecognition.” *Nature Human Behaviour*, 1–12.
- Loewenstein, George, Cindy Bryce, David Hagmann, and Sachin Rajpal.** 2015. “Warning: You are about to be Nudged.” *Behavioral Science & Policy*, 1(1): 35–42.
- Loewenstein, George, Joseph Price, and Kevin Volpp.** 2016. “Habit Formation in Children: Evidence from Incentives for Healthy Eating.” *Journal of Health Economics*, 45: 47–54.
- Madrian, Brigitte, and Dennis Shea.** 2001. “The Power of Suggestion: Inertia in 401k Participation and Savings Behavior.” *Quarterly Journal of Economics*, 116(4): 40.
- Mellström, Carl, and Magnus Johannesson.** 2008. “Crowding out in Blood Donation: Was Titmuss Right?” *Journal of the European Economic Association*, 6(4): 845–863.
- Mertens, Stephanie, Mario Herberz, Ulf Hahnel, and Tobias Brosch.** 2022. “The Effectiveness of Nudging: A Meta-Analysis of Choice Architecture Interventions Across Behavioral Domains.” *Proceedings of the National Academy of Sciences*, 119(1).
- Michelson, Melissa R., Neil Malhotra, Andrew Healy, Donald P. Green, Allison Carnegie, and Ali Adam Valenzuela.** 2012. “The Effect of Prepaid Postage on Turnout: A Cautionary Tale for Election Administrators.” *Election Law Journal*, 11(3): 279–290.
- Milkman, Katherine L., John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian.** 2012. “Following Through on Good Intentions: The Power of Planning Prompts.” *National Bureau of Economic Research*.
- Milkman, Katherine L., Julia A. Minson, and Kevin G. Volpp.** 2014. “Holding the Hunger Games Hostage at the Gym: An Evaluation of Temptation Bundling.” *Management Science*, 60(2): 283–299.
- Miller, Dale T., and Deborah A. Prentice.** 2016. “Changing Norms to Change Behavior.” *Annual Review of Psychology*, 67: 339–361.
- Nickerson, David W.** 2015. “Do Voter Registration Drives Increase Participation? For Whom and When?” *The Journal of Politics*, 77(1): 88–101.

- Nisa, Claudia F., Jocelyn J. Bélanger, Birga M. Schumpe, and Daiane G. Faller.** 2019. “Meta-Analysis of Randomised Controlled Trials Testing Behavioural Interventions to Promote Household Action on Climate Change.” *Nature Communications*, 10(1): 4545.
- Norton, Michael, Daniel Mochon, and Dan Ariely.** 2012. “The IKEA Effect: When Labor Leads to Love.” *Journal of Consumer Psychology*, 22(3): 453–460.
- O’Donoghue, Ted, and Matthew Rabin.** 1999. “Doing It Now or Later.” *American Economic Review*, 89(1): 87.
- Pereira, Rafael H.M., Renato S. Vieira, Fernando Bizzarro, Rogério J. Barbosa, Ricardo Dahis, and Daniel T. Ferreira.** 2023. “Free Public Transit and Voter Turnout.” *Electoral Studies*, 86: 102690.
- Resnjanskij, Sven, Jens Ruhose, Simon Wiederhold, Ludger Woessmann, and Katharina Wedel.** 2024. “Can Mentoring Alleviate Family Disadvantage in Adolescence? A Field Experiment to Improve Labor Market Prospects.” *Journal of Political Economy*, 132(3): 1013–1062.
- Rogers, Todd, and Katherine Milkman.** 2016. “Reminders Through Association.” *Psychological Science*, 27(7): 973–986.
- Rogers, Todd, Katherine Milkman, and Kevin Volpp.** 2014. “Commitment Devices: Using Initiatives to Change Behavior.” *The Journal of the American Medical Association*, 311(20): 2065–2066.
- Roth, Stefan, Thomas Robbert, and Lennart Straus.** 2015. “On the sunk-cost effect in economic decision-making: a meta-analytic review.” *Business research*, 8: 99–138.
- Royer, Heather, Mark Stehr, and Justin Sydnor.** 2015. “Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company.” *American Economic Journal: Applied Economics*, 7(3): 51–84.
- Schacter, Daniel L., Donna Rose Addis, and Randy L. Buckner.** 2007. “Remembering the Past to Imagine the Future: The Prospective Brain.” *Nature Reviews Neuroscience*, 8(9): 657–661.

- Shaddy, Franklin, and Anuj K. Shah.** 2018. “Deciding Who Gets What, Fairly.” *Journal of Consumer Research*, 45(4): 833–848.
- Shaddy, Franklin, and Anuj K. Shah.** 2022. “When to Use Markets, Lines, and Lotteries: How Beliefs About Preferences Shape Beliefs About Allocation.” *Journal of Marketing*, 86(3): 140–156.
- Shu, Suzanne B., and Joann Peck.** 2011. “Psychological Ownership and Affective Reaction: Emotional Attachment Process Variables and the Endowment Effect.” *Journal of Consumer Psychology*, 21(4): 439–452.
- Sunstein, Cass R.** 2013. *Simpler: The Future of Government*. Simon and Schuster.
- Sunstein, Cass R.** 2017. “Nudges That Fail.” *Behavioural Public Policy*, 1(1): 4–25.
- Sunstein, Cass R.** 2021. *Sludge: What Stops Us From Getting Things Done and What To Do About It*. MIT Press.
- Sweeney, Martin, Peter John, Michael Sanders, Hazel Wright, and Lucy Makinson.** 2021. “Applying Behavioural Science to the Annual Electoral Canvass in England: Evidence from a Large-Scale Randomised Controlled Trial.” *Electoral Studies*, 70: 102277.
- Thaler, Richard H., and Shlomo Benartzi.** 2004. “Save More Tomorrow™: Using Behavioral Economics to Increase Employee Saving.” *Journal of Political Economy*, 112(S1): S164–S187.
- Tiefenbeck, Verena, Thorsten Staake, Kurt Roth, and Olga Sachs.** 2013. “For Better or for Worse? Empirical Evidence of Moral Licensing in a Behavioral Energy Conservation Campaign.” *Energy Policy*, 57: 160–171.
- van Dijk, Wilco W., Sjoerd Goslinga, Bart W. Terwel, and Eric van Dijk.** 2020. “How Choice Architecture Can Promote and Undermine Tax Compliance: Testing the Effects of Pre-populated Tax returns and Accuracy Confirmation.” *Journal of Behavioral and Experimental Economics*, 87: 101574.

- Verplanken, Bas, and Suzanne Faes.** 1999. “Good Intentions, Bad Habits, and Effects of Forming Implementation Intentions on Healthy Eating.” *European Journal of Social Psychology*, 29(5-6): 591–604.
- Whillans, Ashley, Joseph Sherlock, Jessica Roberts, Shibeal O’ Flaherty, Lyndsay Gavin, Holly Dykstra, and Michael Daly.** 2021. “Nudging the Commute: Using Behaviorally Informed Interventions to Promote Sustainable Transportation.” *Behavioral Science & Policy*, 7(2): 27–49.
- Yoeli, Erez, Moshe Hoffman, David G. Rand, and Martin A. Nowak.** 2013. “Powering Up with Indirect Reciprocity in a Large-Scale Field Experiment.” *Proceedings of the National Academy of Sciences*, 110: 10424–10429.