The Idle Margin: Why Everyday Conservation Goes Undone

Muhammad Usman Naeem*

13th September 2026

Latest version here

Abstract

Vehicle idling—leaving the engine running while stopped—poses a recurring conservation puzzle: for sufficiently long stops, switching off saves drivers money on fuel while reducing emissions, so private and social incentives point in the same direction. Why, then, do drivers persist in idling? In partnership with a leading telematics firm in Pakistan, I collect 10-second GPS data on ignition status, speed, and location from 6,621 private vehicles. The average vehicle idles for 27 minutes per observed day. At baseline, drivers hold widespread misperceptions about whether warming up an engine extends engine life and how long a stop must last before shutting off saves fuel or reduces environmental harm. I use a randomised additive design varying private information, social information, payments, and an in-car reminder technology. Private information reduces daily idling from stops longer than 30 seconds by 1.8 minutes, or 9% of the control mean, and the effect persists after treatment ends. Estimated benefits are 3.7 times programme costs over the eight-week intervention, with about half accruing directly to drivers through fuel savings. Adding either social information or payments produces no detectable additional reduction. By contrast, activating the reminder device *increases* idling by 1.0 minute, with the increase concentrated in long spells—a pattern consistent with procrastination. The results show that everyday conservation failures can persist even when private and social incentives align—and that interventions meant to correct them can backfire.

Keywords: vehicle idling, air pollution, behavioural frictions, limited attention, information, reminders, field experiment.

JEL codes: R41, Q53, D91, Q58, C93, Q54.

*Tufts University, Department of Economics and The Fletcher School. Email: usman.naeem@tufts.edu. I am especially grateful to Abhijit Banerjee, Steve Cicala, Kyle Emerick, Michael Greenstone, Gabriel Kreindler, and Frank Schilbach for extensive guidance and feedback. I thank Maximilian Auffhammer, Augustin Bergeron, B. Douglas Bernheim, Emily Breza, Ujjayant Chakravorty, Stefano DellaVigna, Cheryl Doss, Meredith Fowlie, Laura Gee, Matthew Gentzkow, Douglas Gollin, Matthew Gudgeon, Rema Hanna, Allan Hsiao, Maulik Jagnani, B. Kelsey Jack, Asim Khwaja, Cynthia Kinnan, Botond Kőszegi, Jing Li, Margaret McMillan, Eoin McGuirk, Gilbert E. Metcalf, Joseph S. Shapiro, Adam Storeygard, Dmitry Taubinsky, Catherine Wolfram, and Sally Zhang, as well as seminar and workshop participants at Tufts, UC Berkeley, Stanford, and Harvard, for helpful comments. I thank Fatima Fida for excellent research assistance, Adnan Hassan for survey programming, and the study partner's teams for excellent implementation support. Funding was provided by the International Growth Centre and the Weiss Fund for Research in Development Economics. The study was approved by the Tufts Social, Behavioral, and Educational Research Institutional Review Board (STUDY00005690) and preregistered in the AEA RCT Registry (AEARCTR-0016449). The views expressed are my own and do not necessarily reflect those of the study partner; any errors are my own.

1 Introduction

Particulate air pollution is one of the world's largest threats to human health, with an especially severe burden in South Asia, home to roughly one-quarter of the global population (UN, 2024). In Pakistan, the average resident could live 3.3 years longer if particulate pollution met the World Health Organization guideline for $PM_{2.5}$ (Greenstone et al., 2025). Road transport is a major source, accounting for 43% of aggregate local air-pollutant emissions in Punjab—the country's most populous province—and 83% in Lahore—its most polluted city (FAO, 2018; Ilyas and Nissar, 2023).

Vehicle idling—leaving the engine running while stopped—may depart from the standard externality benchmark, in which privately optimal choices are socially inefficient because individuals do not internalise the costs they impose on others. For stops longer than 10 seconds, shutting off and restarting the engine generally saves fuel and reduces emissions relative to continued idling (Gaines and Rask, 2012), so private and social incentives can point in the same direction. Yet in the study data, idling in spells longer than this threshold accounts for nearly one-quarter of total engine-on time across major Pakistani cities, implying potentially large private and social gains from reducing idling. Some idling may reflect genuine private benefits, such as maintaining cabin comfort, but baseline beliefs also reveal substantial misperceptions about the private costs and benefits of idling. Idling may therefore persist not only because drivers fail to internalise its social costs, but also because they misperceive their own private costs and benefits (Carrico et al., 2009).

Vehicle idling therefore provides a useful setting to study a broader economic question: when conservation offers both private and social benefits but depends on repeated choices, what prevents people from conserving, and which scalable interventions can change behaviour?

I study this question in a large-scale randomised field experiment with 6,621 private vehicles, conducted in partnership with a leading telematics firm in Pakistan. The additive design varies private information, social information, payments, and an in-car reminder technology to distinguish among misperceptions, social motivation, private preferences, and inattention as explanations for idling. The sample spans the Islamabad Capital Territory and the provinces of Punjab and Sindh—jurisdictions that together account for over three-quarters of Pakistan's population (PBS, 2023). I use GPS trackers to collect data on ignition status, speed, and location at 10-second intervals, allowing me to measure idling passively rather than through self-reports. I classify a vehicle as idling when its ignition is on and its recorded speed is zero and aggregate the resulting spells to the vehicle-day level. The experiment comprises five treatment arms and a control arm.

During the eight-week intervention, drivers assigned to *Private Info* receive weekly messages explaining that modern engines do not require prolonged warm-up and that shutting off during stops longer than 10 seconds generally saves fuel and reduces engine wear. The messages also report each driver's recent weekly idling from stops longer than 30 seconds and translate reductions into personalised fuel savings. *Private & Social Info* adds information about the local health and global climate costs of idling. *Private & Social Info + Pay* additionally pays drivers for each minute by which this weekly idling falls below their baseline benchmark, using a per-minute rate calibrated to the estimated social cost of idling. The reminder arm adds *IdleGuard*, an in-vehicle reminder developed for this study that beeps after 30 seconds of idling and every 30 seconds thereafter while idling continues. A matched no-reminder arm receives identical information and payments but no active reminder, isolating the incremental effect of *IdleGuard*. The control arm receives weekly road-safety messages through the same communication channels.

First, baseline misperceptions are widespread. Among drivers, 76% believe that warming up the engine increases engine life, while only 39% believe that warm-up idling harms the environment. Drivers also substantially overestimate how long a stop must last before shutting off saves fuel or reduces environmental harm. Even using 30 seconds—rather than the 10-second benchmark—as a more conservative cutoff, 94% report a higher fuel-saving threshold and 85% a higher environmental-harm threshold; the corresponding mean reported thresholds are 6.6 and 5.0 minutes.

Second, *Private Info* reduces the primary outcome—daily idling from stops longer than 30 seconds—by 1.8 minutes per observed vehicle-day relative to the control arm, or 9% of the control post-period mean. It also reduces the daily number of such spells by 0.5, or 5% of the control mean, lowering idling on both the intensive and extensive margins. The other treatment arms also produce negative point estimates relative to the control: daily idling falls by 1.3 minutes under *Private & Social Info*, 1.9 minutes under *Private & Social Info + Pay*, and 1.7 minutes under the matched *Private & Social Info + Pay + No-Reminder* arm. The reduction under *Private & Social Info + Pay + Reminder* is smaller, at 0.7 minutes, and statistically indistinguishable from zero.

Third, the incremental effects provide little evidence that social information or payments reduce idling beyond the preceding arm. Adding social information increases daily idling from stops longer than 30 seconds by 0.5 minutes relative to *Private Info*; the estimate is statistically indistinguishable from zero, and its confidence interval rules out large additional reductions. Adding payments decreases idling by 0.6 minutes relative to *Private & Social Info*; this estimate is also statistically indistinguishable from zero, although its confidence interval does not rule out a meaningful additional reduction. By contrast, activating *IdleGuard* increases daily idling from stops longer than 30 seconds by 1.0 minute relative to the matched no-reminder arm, or 5% of the control post-period mean, and increases the daily number of such spells by 0.4, or 4% of the control mean. Because information and payments are held fixed, this comparison isolates the incremental effect of the real-time reminder.

Fourth, both the reduction from *Private Info* and the increase caused by *IdleGuard* are concentrated in sustained idling. Spells lasting more than two minutes account for 79% of the *Private Info* effect on daily idling from spells longer than 30 seconds and 81% of the incremental *IdleGuard* effect. The reminder's adverse effect therefore does not merely reflect behaviour shifting around its first beep at 30 seconds; it arises primarily in long idling spells.

Fifth, the mechanism evidence points to private misperceptions as an important constraint. *Private Info* moves both targeted endline beliefs in the intended direction and produces larger idling reductions among drivers with greater scope for belief correction, although differences across baseline-belief groups are imprecisely estimated. The weak average effect of social information also masks informative heterogeneity: adding it tends to strengthen the private-information effect among less-mistaken drivers but weaken it among those with stronger fuel-saving or environmental misperceptions. These patterns are consistent with broader social information diluting the targeted private-cost correction or conflicting with strongly held prior beliefs. Endline credibility responses provide suggestive support for the latter interpretation: social information is perceived as more credible among drivers with less contrary environmental priors, but not among those with more contrary priors. Payments likewise produce little additional reduction on average, suggesting limited residual price-responsive idling at the implemented incentive level.

Sixth, the evidence points to a two-stage explanation for why *IdleGuard* backfires. First, drivers appear to substitute away from preventive self-monitoring when they expect to rely on the device, allowing more idling spells to reach its 30-second reminder threshold. Consistent with this mechanism, activating *IdleGuard* increases the daily number of such spells relative to the matched no-reminder arm, with 87% of the increase arising in *traffic-like* idling. Second, drivers do not correct promptly once reminded: most of the resulting increase in idling occurs after the first beep, and 56% accrues beyond two minutes. Endline reports provide complementary evidence for both stages, with some drivers reporting greater comfort leaving the engine running because the beep would remind them and many reporting that they did not usually correct promptly after it sounded. This delayed corrective response is consistent with procrastination when switching off entails an immediate action cost but the benefits accrue subsequently (O'Donoghue and Rabin, 1999).

This paper makes three main contributions. First, it contributes to the economics of air pollution by studying idling as an important but comparatively neglected behavioural margin of vehicle emissions. Transport accounts for a substantial share of urban air pollution, yet field evidence on the behavioural choices that determine vehicle emissions remains limited (Jacobsen et al., 2023). Research on vehicle pollution has largely examined policies targeting externalities—including emissions standards, driving restrictions, and congestion pricing (Jacobsen et al., 2023; Davis, 2008; Kreindler, 2024; Cook et al., 2025)—while a smaller literature studies drivers' own fuel-conserving choices through speed and acceleration (Langer and McRae, 2014; Knittel and Tanaka, 2021). Idling is a distinct margin: for sufficiently long stops, switching off saves drivers fuel while reducing local and global emissions, so private and social incentives can align. To the best of my knowledge, this is the first large-scale field experiment to identify why drivers persist in idling and to test interventions designed to reduce it.

Second, the paper contributes to behavioural public economics and the design of corrective energy and environmental policy, which asks how policy should respond when private choices reflect both misperceptions and unpriced social costs (Allcott et al., 2014; Allcott and Taubinsky, 2015; Allcott and Sunstein, 2015; Heutel, 2015; Tsvetanov and Segerson, 2013; Allcott et al., 2015; Allcott and Mullainathan, 2010). A central empirical approach uses information to distinguish mistaken beliefs from residual preferences and thereby assess the scope for corrective policy (Allcott and Taubinsky, 2015). Existing information experiments have largely applied this logic to durable-purchase decisions—including lighting, vehicles, water heaters, and air conditioners—and generally find limited effects on actual purchases (Allcott and Taubinsky, 2015; Rodemeier and Löschel, 2025; Allcott and Knittel, 2019; Allcott and Sweeney, 2017; Newell and Siikamäki, 2014; Davis and Metcalf, 2016). I bring this logic to the repeated realised use of a durable that drivers already own. Private-cost information changes the beliefs it targets and persistently reduces idling, indicating that misperceptions can sustain inefficient use even after the durable-purchase decision has been made.

The additive design extends this approach by asking what explains the idling that remains once private-cost information is provided. Adding social information tests whether making external costs salient produces further conservation; adding payments tests whether residual idling responds to stronger private incentives; and activating a real-time reminder tests whether inattention at the moment of choice remains important. Social information and payments produce little additional reduction on average, whereas the reminder increases idling relative to the matched no-reminder arm. By sequentially targeting distinct candidate explanations for behaviour, the experiment provides evidence on which margins sustain privately and socially costly actions and how corrective policy should target them.

Third, the paper contributes to the economics of attention and reminders by showing that externally supplied attention can substitute for, rather than simply augment, people's own monitoring. Existing work studies how salience, reminders, planning prompts, and other attention aids affect behaviour, while related work in psychology shows that people may strategically offload prospective memory onto external cues (Chetty et al., 2009; Karlan et al., 2016; Carrera et al., 2018; Bronchetti et al., 2023; Damgaard and Gravert, 2018; Koch et al., 2024; Gilbert, 2015). I provide causal evidence from a behaviour-triggered reminder whose timing separates preventive monitoring before the cue from corrective action after it. Activating *IdleGuard* makes idling spells more likely to reach the reminder threshold, consistent with drivers substituting away from their own preventive monitoring, yet does not induce prompt correction once the reminder sounds. The result reveals a two-stage failure of external attention: reliance on the reminder weakens self-monitoring before the cue, while delayed action limits correction afterwards. This latter response is consistent with procrastination when acting entails an immediate cost (O'Donoghue and Rabin, 1999).

The rest of the paper is organised as follows. Section 2 describes the setting and experimental design; Section 3 presents the data and measurement; Section 4 develops the empirical framework; Section 5 reports the experimental results; Section 6 examines mechanisms; Section 7 discusses welfare and policy implications; and Section 8 concludes.

2 Setting and Experimental Design

2.1 Vehicle Idling in Pakistan

Road transport is a major source of local air pollution in Pakistan, accounting for an estimated 43% of aggregate air-pollutant emissions in Punjab and 83% in Lahore (FAO, 2018; Ilyas and Nissar, 2023). Within transport, drivers' everyday operating choices represent a potentially important margin of emissions abatement. I focus on vehicle idling—leaving the engine running while the vehicle is stationary.

I distinguish three broad idling contexts: *warming*, before the vehicle first moves after ignition; *waiting*, such as during passenger pickup; and *traffic*, such as at stoplights or in congestion (Carrico et al., 2009). Modern engines do not require prolonged warming, and for stops longer than 10 seconds, shutting off and restarting generally saves fuel, reduces emissions, and limits engine wear (U.S. EPA, 2025; U.S. DOE, 2026; Gaines and Rask, 2012). In the study's baseline data, the median vehicle idles for approximately 20 minutes per observed day; across the study sample, idling spells longer than 10 seconds account for nearly 25% of total engine-on time.

2.2 Partner and Study Sample

I partnered with a leading vehicle-tracking company in Pakistan to conduct the experiment. Its trackers can be remotely configured to record ignition status, speed, and location at 10-second intervals, allowing idling spells to be measured passively and at high frequency. Such trackers are common in Pakistan's passenger-car market: banks and other financial institutions often require them for financed vehicles, while owners also install them for theft protection and vehicle recovery (TPL Trakker, 2024). During the recent expansion in automobile lending, bank-financed vehicles accounted for roughly 40–45% of new-car sales (Khan, 2021), so tracker-equipped cars are not confined to a narrow specialist population.

I applied four restrictions to the partner's universe of 72,076 active subscriptions. First, I retained passenger four-wheelers—sedans, hatchbacks, SUVs, and similar vehicles—and excluded motorcycles, rickshaws, tractors, trucks, and other commercial vehicles. This focuses the study on ordinary passenger cars, for which standard anti-idling guidance, including the 10-second shutoff rule, most directly applies. Second, I retained owner-driven vehicles and excluded company-linked cars, whose drivers may have weaker incentives to reduce fuel and maintenance costs. Third, to maximise continued GPS observability, I retained vehicles with either sufficient subscription time remaining through the study period or at least 18 months of prior tracking history, which indicated a greater likelihood of renewal. Fourth, I restricted the sample geographically using tracker installation city as a proxy for where the vehicle was typically driven, retaining vehicles in the Islamabad Capital Territory and the provinces of Punjab and Sindh. These restrictions yielded 42,101 eligible vehicles. Excluding 1,400 vehicles that had participated in an earlier pilot left a recruitment frame of 40,701, which I randomly ordered for recruitment.

I established a survey unit at the firm's in-house call centre to recruit participants by phone. Starting from the randomly ordered frame, the unit proceeded sequentially through the list until it reached the powered target of 6,000 drivers who had consented and completed the baseline survey. Recruiting sequentially from the randomised frame avoided fixing the sample in advance and falling short because of non-contact, non-consent, or incomplete interviews.

After providing informed consent to participate in a study of driving behaviour conducted jointly by the firm and Tufts University researchers, participants were told that they might be recontacted for further study information, weekly fuel-use reports, and an endline survey. The baseline survey collected demographic and contact information, including the WhatsApp number and email address used for treatment delivery, together with measures of driving and idling beliefs and social and environmental preferences. Participants received a small payment for completing the baseline survey, with additional incentives for the weekly fuel survey and an incentivised social-preference measure; all payments were disbursed at the end of the study.

Because treatment was delivered to an individual driver while GPS outcomes were measured at the vehicle level, enumerators were trained to identify and enrol the person who drove the vehicle most often. The baseline survey therefore recorded the main driver's contact information, aligning treatment delivery with the person responsible for most of the vehicle's observed driving behaviour.

Most vehicles in the 40,701-vehicle recruitment frame initially reported GPS data at intervals longer than 10 seconds. As recruitment progressed, I worked with the partner's diagnostics team to remotely reconfigure enrolled vehicles' trackers to report at 10-second intervals. Because temporary connectivity problems could delay these over-the-air updates, recruitment continued beyond the powered target rather than waiting for every enrolled tracker to be reconfigured. Recruitment ultimately yielded 10,491 completed baseline surveys. After applying the remaining respondent and GPS eligibility restrictions, the randomisation sample comprised 6,621 vehicles with complete baseline data and 10-second GPS reporting.

2.3 Experimental Treatments

The experiment includes five treatment arms and a control arm. Four treatment arms form an additive sequence in which private information is followed by social information, payments, and a real-time reminder. Comparisons along the first three arms identify the incremental effects of social information and payments. To identify the reminder effect, I add a fifth treatment arm that receives the same information and payments as the reminder arm but no active reminder. Drivers in both arms first completed the same incentivised demand elicitation for the reminder device, so their comparison isolates the incremental effect of activating the reminder while holding fixed any effect of the elicitation itself. Together, these contrasts distinguish among misperceptions, social motivation, private preferences, and attention as candidate determinants of idling. I also estimate the effect of each treatment arm relative to the control arm.

Across all six arms, participants received weekly messages through the partner's mobile app, WhatsApp Business account, and official email platform. Messages were sent each Monday for eight weeks beginning in early April 2026. Table 1 summarises the components included in each arm; the treatment content is described below.

Table 1: Experimental arms and treatment components

Arm	Treatment component			
	Private Info	Social Info	Pay	Reminder
Private Info	✓			
Private & Social Info	✓	✓		
Private & Social Info + Pay	✓	✓	✓	
Private & Social Info + Pay + Reminder	✓	✓	✓	✓
Private & Social Info + Pay + No-Reminder	✓	✓	✓	
Control				

Notes: Check marks indicate the components included in each arm. Comparisons along the first three treatment arms identify the incremental effects of social information and payments. Comparing the *Reminder* and *No-Reminder* arms identifies the incremental effect of activating the in-vehicle reminder. The control arm received general road-safety messages.

Private Info. Drivers in this arm received anti-idling information targeting *warming*, *waiting*, and *traffic* idling. The message explained that modern engines do not require prolonged warm-up and that shutting off during stops longer than 10 seconds generally saves fuel and reduces engine wear (U.S. EPA, 2025; U.S. DOE, 2026; Gaines and Rask, 2012). It also reported each driver's typical weekly idling from stops longer than 30 seconds and translated a one-minute reduction into personalised per-minute, weekly, and monthly fuel savings. The same message was delivered each week.

Private & Social Info. Drivers in this arm received the *Private Info* content plus information about the social costs of idling. The message explained that reducing *warming*, *waiting*, and *traffic* idling lowers local air pollution and CO_2 emissions. It translated the climate damages avoided by a one-minute reduction in idling into a personalised rupee value and reported corresponding weekly and monthly totals alongside the private fuel savings (Carleton and Greenstone, 2021; Rennert et al., 2022). The message also described the health harms from $PM_{2.5}$ (Greenstone et al., 2025) and included photographs of a smog-covered road in Lahore and a man stranded with his dog during a flood, with captions linking tailpipe $PM_{2.5}$ emissions to shorter lives and tailpipe CO_2 emissions to flooding.

Private & Social Info + Pay. Drivers in this arm received the *Private & Social Info* content plus a financial incentive. I calibrated a fixed Pigouvian-style payment of Rs. 4 per minute using the arm's average weekly idling from stops longer than 30 seconds and the associated estimated climate damages (Carleton and Greenstone, 2021; Rennert et al., 2022; Pigou, 1920; Baumol and Oates, 1971). Drivers earned Rs. 4 for each minute by which their weekly idling fell below their individual baseline benchmark. Each weekly message reported the payment rate, the maximum amount available, and, from the second week onward, the amount earned in the preceding week. Because payments were processed in arrears, drivers received a final payment notification on the Monday after the eight-week intervention ended. Across the three arms with a payment component, the average maximum weekly payment was Rs. 538, equivalent to 5.6% of estimated average weekly earnings for salaried participants in those arms (Pakistan Bureau of Statistics, 2025).

Private & Social Info + Pay + Reminder. Drivers in this arm received the *Private & Social Info + Pay* content and were offered *IdleGuard*, a GPS tracker developed for this study to deliver a real-time anti-idling reminder. After 30 seconds of continuous idling, *IdleGuard* emitted a beep and repeated the alert every 30 seconds while idling continued. Drivers were told that the device would remind them to switch off the engine during stops longer than 30 seconds. The reminder remained active throughout the eight-week intervention.

Private & Social Info + Pay + No-Reminder. Drivers in this arm received the same information and financial incentive as the *Reminder* arm but no active *IdleGuard* reminder. Comparing the two arms therefore isolates the incremental effect of activating the real-time reminder.

Control. Drivers in the control arm received general road-safety messages throughout the eight-week intervention, including reminders to wear a seatbelt and use indicators before changing lanes.

2.4 Randomisation

I randomised the 6,621-vehicle randomisation sample in two stages to construct the six experimental arms.

The first stage took place immediately after baseline and assigned vehicles to five groups: (1) *Private Info*, with 1,030 vehicles; (2) *Private & Social Info*, with 1,030 vehicles; (3) *Private & Social Info + Pay*, with 1,030 vehicles; (4) *Private & Social Info + Pay + Tech*, with 2,500 vehicles; and (5) *Control*, with 1,031 vehicles. I stratified the first-stage assignment by vehicle age (0–4 versus 5+ years), which is closely related to vehicle emissions (Jacobsen et al., 2023), and certainty about main-driver status (certain versus uncertain). Main-driver status was uncertain for 621 baseline respondents.

The second stage was conducted through treatment-briefing phone calls, which had been described during the baseline survey as possible follow-up calls providing additional information about the study. These calls began shortly after the first-stage randomisation and were conducted separately for each of the five first-stage groups. During calls to the 2,500 drivers assigned to the *Private & Social Info + Pay + Tech* group, participants completed an incentivised demand elicitation; after a randomly selected subset had their choices implemented, the remaining drivers entered a lottery assigning them to the *Private & Social Info + Pay + Reminder* or *Private & Social Info + Pay + No-Reminder* arm.

The procedure had two steps. First, drivers made an incentivised choice between a new tracker device provided at no cost and Rs. 10,000, using an elicitation approach in the spirit of the Becker–

DeGroot–Marschak mechanism (Becker et al., 1964). The new tracker was described as identical to the respondent's existing tracker, except that it would emit a gentle beep after the vehicle had remained stopped with the engine running for some time, reminding the driver to switch off the engine. Drivers were also asked how much they believed the device could reduce their weekly fuel expenditure and whether they would use it if provided free of charge.

Care was taken to minimise any pre-treatment behavioural response to the description of the device. In particular, enumerators did not disclose the precise interval after which the device would beep and did not use the terms "idling" or *IdleGuard*. To support incentive-compatible elicitation, enumerators received rigorous training, and call recordings were back-checked daily, following implementation practices used in related field work (Berkouwer and Dean, 2022). Respondents were instructed clearly, and twice, to state their true preference. They were informed that they had a small probability of receiving their chosen option and that stating a particular preference could not affect whether their choice was implemented, since implementation was determined by a random draw. Before the choice was locked in, each driver was told that it could not subsequently be changed and was asked whether they wished to retain that choice, encouraging careful consideration. The stated choice was implemented for a randomly selected 1% of drivers, corresponding to 25 participants, who were then removed from the experiment (Allcott et al., 2020).

Second, the remaining 2,475 drivers were told that a computerised lottery conducted during the call would determine whether they were assigned to receive the device. The lottery assigned 1,500 drivers to the *Private & Social Info + Pay + Reminder* arm and 975 drivers to the *Private & Social Info + Pay + No-Reminder* arm. Drivers assigned to the reminder arm were informed that a technician would contact them to arrange installation. Drivers assigned to the no-reminder arm were told that they would continue participating in the study as planned.

2.5 Implementation

Treatment briefing. Following the first-stage randomisation in early November 2025, the survey unit began conducting separate treatment-briefing calls for each of the five first-stage groups: *Private Info*, *Private & Social Info*, *Private & Social Info + Pay*, *Private & Social Info + Pay + Tech*, and *Control*. The calls served three purposes: (1) to provide the follow-up study information that participants had been told at baseline they might receive and to prepare them for the subsequent treatment-delivery phase; (2) to collect bank-account information from drivers assigned to first-stage groups with a payment component, allowing any anti-idling payments earned during treatment to be processed; and (3) within the *Private & Social Info + Pay + Tech* group, to conduct the demand elicitation and second-stage randomisation described above and to collect the contact and location information required for *IdleGuard* installation. To allow sufficient time for installation, drivers in the *Private & Social Info + Pay + Tech* group were prioritised for briefing, alongside a randomly selected 20% of drivers from each of the remaining first-stage groups. This sequencing allowed installations to begin while briefing was still under way. Treatment-briefing calls continued

Figure 1: Two-stage randomisation and construction of the final experimental arms

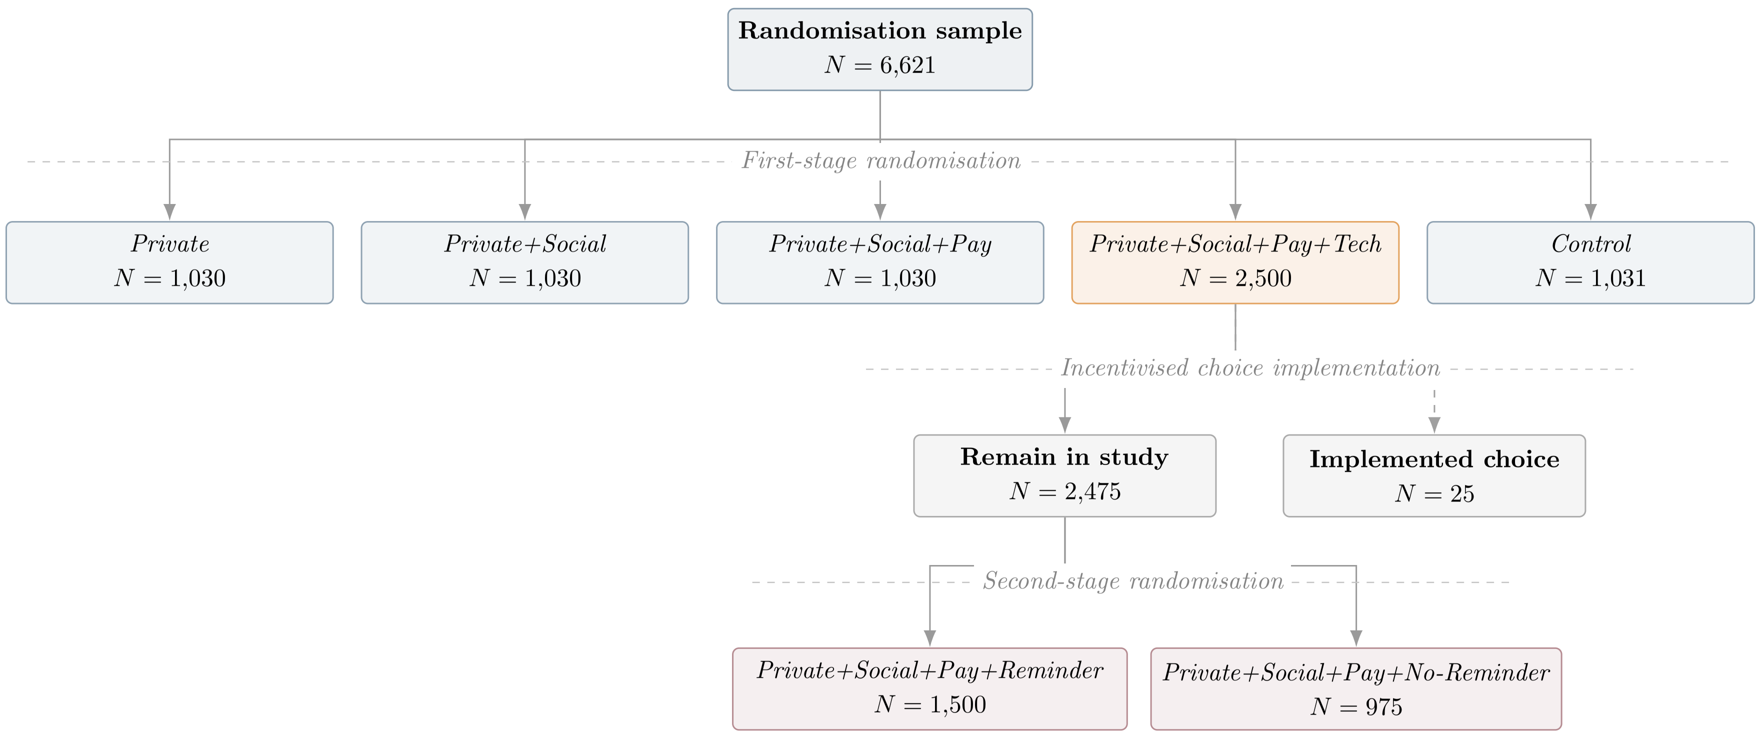

Notes: The first stage assigned the 6,621-vehicle randomisation sample to five groups, stratified by vehicle age and certainty about main-driver status. Drivers in the *Private + Social + Pay + Tech* group then chose between a new tracker device and Rs. 10,000. The incentivised choice was implemented for 25 randomly selected drivers, who were removed from the experiment; the remaining 2,475 were randomised between the *Private + Social + Pay + Reminder* and *Private + Social + Pay + No-Reminder* arms.

through mid-January 2026.

Drivers in the *Private Info* and *Private & Social Info* groups were told that, later in the study, they would begin receiving weekly messages about driving behaviour, including tips intended to improve their driving habits. They were informed that the messages would be sent every Monday for eight weeks through the partner's mobile app, WhatsApp, and email, and were asked to look out for them. Drivers in the *Control* group were told that they would receive general road-safety tips. To avoid any premature treatment response, idling was not mentioned at any stage of these calls. Indeed, the survey team itself had not been informed that the study focused on idling.

Drivers assigned to receive a payment component—those in *Private & Social Info + Pay* and, following the second-stage randomisation, those in *Private & Social Info + Pay + Reminder* and *Private & Social Info + Pay + No-Reminder*—were also informed that they would have an opportunity to earn money based on their driving behaviour. Idling was again not mentioned during these calls; the subsequent treatment messages instead explained how payments could be earned by reducing idling. Drivers were asked to provide bank-account or mobile-wallet details for receiving any payments earned. Those assigned to the *Private & Social Info + Pay + Reminder* arm were additionally asked to provide an address and contact details for the installation visit, as well as the city in which the vehicle was primarily driven. I used this information to coordinate *IdleGuard* installations in batches with the partner's nationwide field-force team as briefing calls progressed.

***IdleGuard* rollout.** Of the 1,500 drivers assigned to the *Private & Social Info + Pay + Reminder* arm, 1,129 (75%) accepted the *IdleGuard* offer. By August 2025, I had worked with the partner's diagnostics and laboratory teams to import and locally procure the required components and to develop, test in the laboratory, and field-test sufficient *IdleGuard* devices for the rollout. As briefing calls progressed, devices were dispatched in staggered batches from the partner's Karachi head office to regional field offices, with each shipment matched to the assigned drivers' reported installation cities and the technicians responsible for serving those locations. The rollout was coordinated through 30 regional dispatch locations and involved 80 field technicians serving drivers across 164 cities and towns. Field technicians received a per-installation incentive. In total, 678 devices were installed by the end of January 2026, corresponding to 60% of drivers who had accepted the offer.

Each installation took approximately 45–60 minutes. The technician first tested that the device was functioning correctly, including that it emitted the intended beep. A diagnostics lead at the head office then remotely disabled the beep using over-the-air technology while remaining on a call with the technician on site, allowing the lead to hear and confirm that the device had fallen silent. Drivers were therefore left with a fully functioning tracker whose reminder feature had been disabled and were informed that the beep would be activated only at a later launch date.

Treatment delivery. Although treatment briefing and *IdleGuard* installations had been completed by the end of January 2026, I delayed treatment delivery until after Ramadan because driving patterns could differ during the month. The eight-week intervention therefore began in early April

2026, simultaneously across all six experimental arms, and continued through the end of May. Of the 6,596 vehicles in the final six-arm experimental sample, 35 drivers had requested during treatment briefing not to receive further messages, leaving a treatment-delivery sample of 6,561 vehicles. To study the persistence of treatment effects and distinguish learning from continued message salience, I randomly assigned half of the *Private Info* arm to receive treatment for two additional weeks after the scheduled eight-week intervention, while treatment ended as scheduled for the remaining half.

Weekly treatment delivery involved three main activities. First, treatment messages were sent every Monday at 9:00 a.m. Pakistan Standard Time through the partner's client-facing mobile app, WhatsApp Business account, and official email platform. Second, beginning in the second treatment week, anti-idling payments earned during the preceding week were calculated and processed each Monday. Third, at treatment launch, the reminder feature of installed *IdleGuard* devices was activated remotely using over-the-air technology.

Each Sunday, I used GPS data from the preceding Monday–Saturday period to construct the driver-level statistics required to populate the message templates shown in Online Appendix Section A.2. I worked with the partner's information technology team to develop an automated system for transferring these data, populating the treatment templates, and delivering the resulting messages. WhatsApp and email messages were sent directly to enrolled main drivers using contact details collected at baseline and reconfirmed during briefing. Because the partner could not reliably identify the main driver's app user ID when multiple user IDs were linked to the same vehicle, app messages were sent to all user IDs associated with the randomised vehicle. In each of the eight common-treatment weeks, the system sent approximately 22,500 messages across the three channels and recorded message-level delivery and engagement information for each platform.

Sending app messages to multiple user IDs associated with the same vehicle was consistent with vehicle-level treatment assignment and outcome measurement. Contamination could arise, however, if a single user ID were linked to study vehicles assigned to different experimental arms. Such overlap was rare: in the baseline data, only 0.6% of app user IDs were linked to more than one study vehicle.

Using the account details collected during treatment briefing, I prepared weekly payment files alongside the treatment-message files and sent them to the partner's finance team for processing each Monday. I reconciled the resulting payment-status reports and resubmitted failed transactions for processing. Across the eight treatment weeks, drivers in the three treatment arms with a payment component earned an average of Rs. 78 per week, or 0.8% of the same estimated weekly earnings benchmark. The partner successfully processed 8,106 weekly payment transactions over the intervention, totalling approximately Rs. 2.05 million (USD 7,400).

To activate the *IdleGuard* reminder remotely, I shared the tracker details for all 678 successful installations with the partner's diagnostics team on the first day of treatment. The reminders had been activated on almost all devices by the following day; a small number came online later because temporary network problems delayed receipt of the remote activation command. Drivers could contact the partner's call centre by phone, email, or social media. During the eight-week treatment period, only 39 drivers requested that the reminder be deactivated because they found the beep annoying, corresponding to 6% of successful installations.

Fuel survey and endline. During the first seven treatment weeks, I administered an incentivised weekly fuel-expenditure survey through a WhatsApp chatbot using the WATI API. At 12:00 p.m. Pakistan Standard Time, the chatbot asked drivers to report their fuel expenditure for the previous day. Drivers who had not responded received a reminder at 8:00 p.m. the same day. Before launch, I randomly assigned the 6,596 drivers in the final six-arm experimental sample across Tuesday through Sunday, stratified by experimental arm. I re-randomised survey days each week so that reported expenditure would reflect a range of driving days rather than a fixed day of the week.

Following the intervention, I worked with the survey unit to conduct a phone endline survey beginning in mid-June 2026 and continuing through August 2026. The survey repeated selected baseline measures of beliefs and self-reported behaviour and collected new measures of message recall and engagement, perceived treatment credibility and usefulness, idling contexts, intentions to reduce future idling, and treatment-specific experiences with payments and *IdleGuard*. It also included incentivised choices regarding whether drivers wished to continue receiving the weekly messages.

2.6 Balance and Pre-Treatment Comparability

The five first-stage assignment groups are balanced across baseline survey and administrative characteristics. Because the second-stage randomisation created the final six experimental arms, I also assess their comparability in GPS observability and idling immediately before treatment launch. The estimated differences are generally small, and joint tests do not reject balance or pre-treatment comparability at conventional significance levels. See Online Appendix Tables C.1 and C.2.

3 Data and Measurement

3.1 GPS Tracker Data

Each of the 6,621 vehicles in the randomisation sample was equipped with a GPS tracker installed beneath the dashboard and powered by the vehicle's battery, allowing it to transmit data to the partner's servers even when the engine was off. The tracker's GPS receiver used signals from orbiting satellites to determine the vehicle's location (GPS.gov, n.d.), while an integrated cellular modem and subscriber identity module (SIM) card transmitted the resulting records to the partner's servers over the mobile network. Each row, or ping, in the raw GPS tracker data contains a timestamp, latitude, longitude, speed, ignition status, battery voltage, and other tracker variables.

The partner's trackers ordinarily reported data at a coarser frequency while vehicles were idling

Figure 2: Experimental timeline and implementation activities

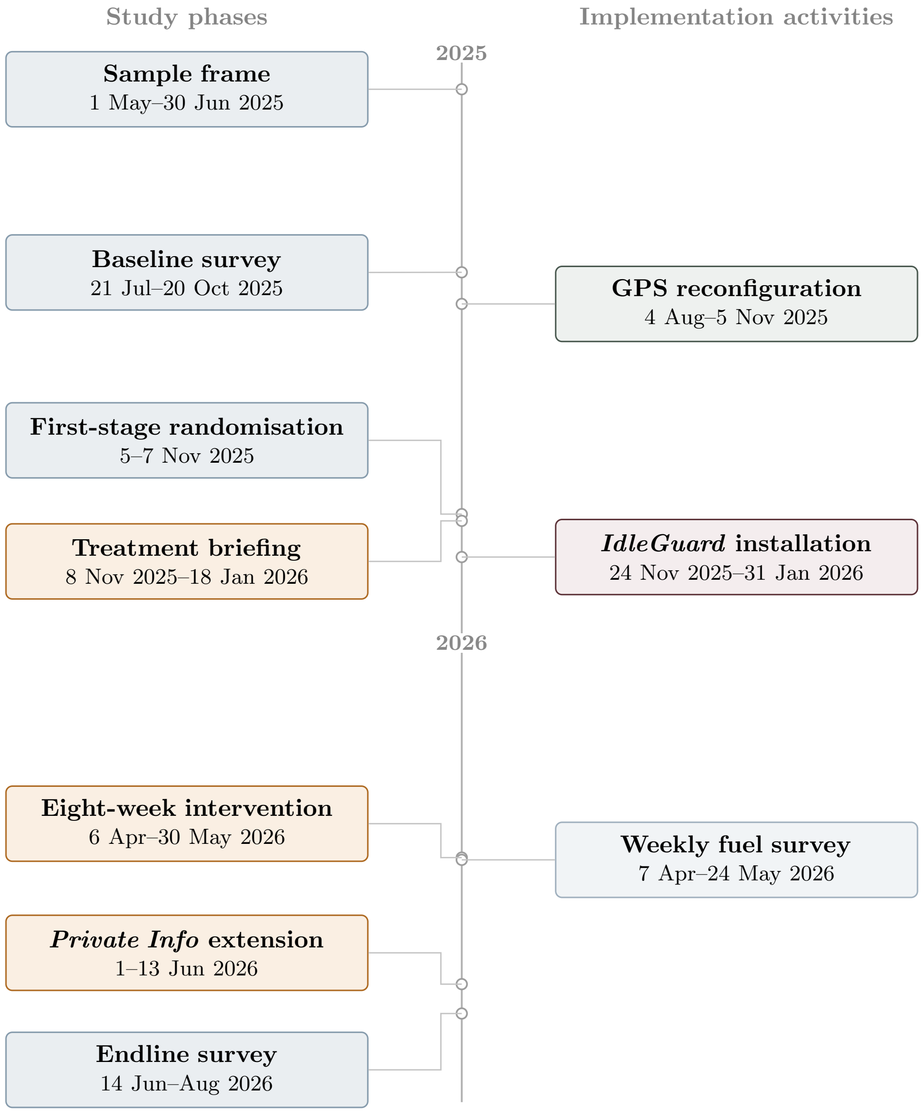

Notes: The central spine uses a linear calendar scale based on the start of each phase or activity; boxes are displaced only to avoid overlap. GPS trackers were reconfigured to report at 10-second intervals before randomisation, and *IdleGuard* installation was completed before treatment launch. Half of the *Private Info* arm was randomly assigned to receive messages for two additional weeks. The incentivised fuel survey was fielded through a WhatsApp chatbot during the first seven treatment weeks.

because its operational requirements placed greater emphasis on high-frequency reporting during movement. As baseline data collection progressed, I worked with the partner's diagnostics team to reconfigure the trackers installed in recruited vehicles to report at 10-second intervals while idling. A tracker was considered successfully reconfigured if, during the post-reconfiguration diagnostic period, at least 95% of its idling pings were no more than 15 seconds apart. The resulting data provide approximately 10-second-resolution records for observed vehicles from Monday through Saturday. Sundays were excluded from the weekly data extracts to allow sufficient time to process the preceding week's data and prepare treatment-delivery files for the following Monday. The GPS data cover the period from 17 November 2025 through 22 August 2026 and span 81 tracker-installation cities across Pakistan.

3.2 Measuring Idling

Ping-level idling classification. I classify a ping as idling when the vehicle's ignition status is on and its recorded speed is zero. Online Appendix B.1 validates ignition status as a strong proxy for whether the engine is running and documents the GPS-displacement correction used to remove false positives from slow movement.

Idling-spell construction. Within each vehicle-day and trip, I group consecutive idling pings into spells and aggregate their duration, excluding intervals spanning telemetry blackouts. Online Appendix B.2 provides complete construction details.

Vehicle-day outcome construction. I aggregate idling spells within each vehicle-day to construct daily idling duration and spell counts. The primary outcome is daily idling duration from spells longer than 30 seconds. The treatment informed drivers that shutting off the engine during stops longer than 10 seconds generally saves fuel and reduces emissions, while personalised feedback and payments were based on idling spells longer than 30 seconds (U.S. EPA, 2025; U.S. DOE, 2026; Gaines and Rask, 2012). I also construct outcomes using the 10-second threshold and total idling without a minimum spell-duration threshold. I screen implausible vehicle-days and manually audit extreme observations; results are unchanged under conservative winsorisation. Online Appendix B.3 provides complete construction and validation details.

3.3 Analysis Sample and Outcomes

Of the 6,596 vehicles in the final six-arm randomised sample, 6,451 appear at least once in the cleaned GPS tracker data over the full panel period, from 17 November 2025 to 22 August 2026. These vehicles constitute the GPS analysis panel. Online Appendix Figure A.1 summarises the construction of the experimental sample and GPS analysis panel. The unit of analysis is the vehicle-day, organised on a balanced vehicle-by-date grid. Because trackers report whether vehicles are parked, moving, or idling, behavioural outcomes are coded as zero when the corresponding activity does not occur. They remain missing only during temporary periods of tracker non-reporting caused by connectivity problems.

The primary outcome is daily idling duration from spells longer than 30 seconds. Secondary outcomes include total daily idling duration, daily idling duration from spells longer than 10 seconds, daily counts of spells longer than 10 and 30 seconds, daily moving minutes, and trip counts. I also examine effects across alternative spell-duration thresholds and bins and decompose idling into *warming-like*, *traffic-like*, and *waiting-like* contexts using the telemetry data. The first idling spell after ignition-on is classified as *warming-like* if the vehicle has not yet moved. A non-warming spell is classified as *traffic-like* if the vehicle moved within 120 seconds before the spell began and moves again within 300 seconds after its start. All remaining spells are classified as *waiting-like*.

3.4 Survey and Fuel-Expenditure Data

Baseline survey. The baseline survey collected demographic information on the vehicle's main driver, beliefs about idling and fuel-saving and environmental-harm thresholds, confidence in those beliefs, and stated and revealed measures of pro-social preferences, the latter elicited through an incentivised choice. I use these measures to assess statistical balance across assignment groups, characterise baseline misperceptions, and examine heterogeneity in treatment effects.

Weekly fuel-expenditure survey. The weekly fuel-expenditure survey asked the vehicle's main driver to report fuel expenditure on the previous day, with respondents instructed to report zero if none was incurred. I use this measure to estimate treatment effects on daily fuel expenditure. The survey-based measure provides a self-reported complement to engineering estimates of fuel use and emissions constructed from GPS inputs.

Endline survey. The endline survey re-elicited idling beliefs, fuel-saving and environmental-harm thresholds, and confidence in those beliefs to study belief updating. It also repeated the baseline stated measure of pro-social preferences to examine whether stated preferences changed over the study. The survey collected measures of message engagement, comprehension, perceived credibility, and incentivised revealed demand for continued treatment messages to study mechanisms. Finally, treatment-specific modules collected information to help interpret the incremental effects of social information and *IdleGuard*.

3.5 Data Quality and Attrition

Post-launch GPS data quality and observability. Behavioural outcomes remain missing during periods of tracker non-reporting, raising the possibility that estimated treatment effects partly reflect differential selection into the observed GPS sample. Following the GPS data-quality attrition check in Kreindler (2024), Online Appendix Table C.6 tests for treatment-arm differences in tracker reporting and GPS observability after treatment launch. The estimated differences are small and statistically insignificant, and the joint tests provide no evidence of differential post-launch data availability across study arms.

Survey attrition. Online Appendix Table C.7 tests for differential attrition in the weekly fuel-expenditure survey. Drivers in the three treatment arms with payments were more likely than those in the control arm to provide valid fuel-expenditure outcomes. I therefore do not report treatment effects on fuel expenditure for these arms relative to the control arm or the incremental effect of payments. Online Appendix Table C.8 finds no evidence of differential attrition for *Private Info* relative to the control arm or for the incremental effect of *IdleGuard*.

4 Empirical Framework

4.1 Treatment Logic and Predictions

Why do drivers privately choose to idle? Does idling reflect a mistake, a genuine preference, or inattention in the moment? And does concern for social harms also matter? The study uses a sequential additive treatment design to diagnose which of these candidate determinants are behaviourally important.

At a sufficiently long stop, the driver chooses between continuing to idle and switching off the engine to restart later. The driver chooses to idle if the utility from idling is at least as high as the utility from restarting:

$$U(\text{Idle}) \geq U(\text{Restart}).$$

The utility difference between the two actions can be written as

$$U(\text{Idle}) - U(\text{Restart}) \equiv B^p - C^p + D - \alpha C^s. \tag{1}$$

B^p denotes the genuine private benefit of idling relative to restarting. I normalise the private benefit of restarting to zero. The private benefits of idling include comfort and convenience, such as avoiding the hassle of switching off and restarting. Similarly, C^p denotes the private fuel and maintenance costs of idling relative to those of restarting. D captures behavioural distortions, including mistakes and inattention. Mistakes arise from failures of private-cost information, including beliefs that warming up benefits engine life and overestimation of the stop duration at which the fuel and maintenance costs of continued idling exceed those of restarting. Inattention means that, at the moment of choice, the driver fails to notice the relevant stop, track its duration, consider switching off, or act on that option. C^s denotes the social costs of idling relative to restarting, including damages from local air pollutants and greenhouse-gas emissions. Finally, α measures the weight the driver places on these social costs.

Private-cost information on warming requirements and on the fuel- and maintenance-saving thresholds for waiting and traffic idling, delivered through *Private Info*, targets the mistake component of D. Its effect is identified relative to the control arm and is expected to reduce idling. *Private & Social Info* adds information about the social costs of idling while holding private-cost information fixed. Its effect is therefore identified relative to *Private Info* and tests whether making C^s salient further reduces idling. Absent countervailing responses to the additional information, the predicted effect is more negative when drivers place greater weight, α, on these social costs.

Similarly, comparing *Private & Social Info + Pay* with *Private & Social Info* holds the information fixed and tests whether residual idling is responsive to financial incentives, as would be expected if it reflects genuine private value from idling. The predicted effect is negative. Finally, comparing *Private & Social Info + Pay + Reminder* with *Private & Social Info + Pay + No-Reminder* holds information and financial incentives fixed and tests whether residual idling reflects inattention at the moment of choice. The predicted effect is negative unless the reminder technology crowds out drivers' own monitoring.

4.2 Main Specification

The main estimation sample comprises 6,451 vehicles observed from 17 November 2025 through 30 May 2026, the final day of the eight-week treatment period, with the vehicle-day as the unit of observation. I use the following two-way fixed-effects difference-in-differences (TWFE-DiD) specification to estimate treatment effects for each treatment arm relative to the control arm. The same specification also underlies the incremental treatment effects defined in Section 4.3:

$$Y_{it} = \alpha_i + \delta_t + \sum_{g \neq P} \beta_g \left(\mathbb{1}\{G_i = g\} \times \text{Post}_t \right) + \varepsilon_{it}. \tag{2}$$

The vehicle and day are indexed by i and t, respectively. Y_{it} denotes the behavioural outcome for vehicle i on day t. α_i denotes vehicle fixed effects, which absorb time-invariant differences across vehicles, such as persistent differences in baseline idling or regular travel patterns. δ_t denotes day fixed effects, which absorb shocks common to all vehicles on a given day, such as weather or congestion. Because treatment assignment is random, the fixed effects are not required for causal identification; they improve precision by absorbing persistent vehicle-level and common day-level variation in idling. $\mathbb{1}\{G_i = g\}$ is an indicator for vehicle i's assignment to arm g, and Post_t equals one from 6 April 2026, when treatment delivery began for all arms, and zero otherwise. The control arm is the omitted category. Each β_g therefore compares the pre-to-post change in outcomes for vehicles assigned to treatment arm g with the corresponding change for vehicles assigned to the control arm. ε_{it} is the error term. The paper's primary diagnostic comparisons are the effect of *Private Info* relative to the control arm and the incremental treatment effects introduced in Section 4.3. Behavioural outcomes remain missing on unobserved vehicle-days rather than being coded as zero. Standard errors are clustered at the vehicle level throughout to account for serial correlation within vehicles over time.

4.3 Incremental Treatment Effects

The incremental effects of social information, payments, and *IdleGuard* are linear combinations of the treatment-arm coefficients from Equation 2. They compare *Private & Social Info* with *Private Info*, *Private & Social Info + Pay* with *Private & Social Info*, and *Private & Social Info + Pay + Reminder* with *Private & Social Info + Pay + No-Reminder*, respectively. The incremental effect of social information tests whether making social costs salient matters beyond private-cost information. The incremental effect of payments tests whether residual idling is responsive to financial incentives, as would be expected if it reflects genuine private value from idling. The incremental effect of *IdleGuard* tests whether residual idling reflects inattention at the moment of choice.

4.4 Dynamics and Persistence

I use the following treatment-by-week specification to examine how the treatment effects estimated in Equation 2 evolve over the eight treatment weeks. The estimation sample and time window are unchanged:

$$Y_{it} = \alpha_i + \delta_t + \sum_{g \neq P} \sum_{w=13}^{20} \beta_{gw} \left(\mathbb{1}\{G_i = g\} \times \mathbb{1}\{W_t = w\} \right) + \varepsilon_{it}, \tag{3}$$

$\mathbb{1}\{W_t = w\}$ is an indicator equal to one when day t falls in treatment week w. The control arm is the omitted category, and the twelve pre-treatment study weeks form the pooled reference period. Each β_{gw} therefore compares the change in outcomes for vehicles assigned to treatment arm g in treatment week w, relative to the pooled pre-treatment period, with the corresponding change for vehicles assigned to the control arm. All other terms are defined as in Equation 2.

At the end of the eight-week intervention, I randomly assigned the *Private Info* arm to *Early-Off* and *Late-On* subgroups. The *Early-Off* subgroup stopped receiving messages after 30 May 2026, while the *Late-On* subgroup continued receiving them for two additional weeks. To study the persistence of treatment effects across all treatment arms, I modify Equation 2 by replacing the single treatment-by-post interaction with treatment-arm interactions for the active-treatment period, the first two weeks after the common intervention ended, and the remaining ten post-treatment weeks. I also estimate a week-specific version over all twelve weeks after the common intervention ended to trace the evolution of post-treatment effects more finely. For the sharper comparison within the *Private Info* arm, I directly compare the randomly assigned *Early-Off* and *Late-On* subgroups during the two-week continuation window. I complement this comparison with a stacked switch-off specification reported in Online Appendix Table D.13.

5 Experimental Results

5.1 Baseline Idling and Beliefs

Daily idling levels are high before treatment, and baseline survey responses reveal widespread misperceptions about both the consequences of idling and when shutting off the engine becomes beneficial. As shown in Table 2, GPS coverage is high during the two weeks immediately preceding treatment launch: vehicles are observed on 87% of possible days. Over observed vehicle-days, total idling is also substantial, averaging 27 minutes per day. Idling from spells longer than 30 seconds accounts for more than 20 of these minutes. Although nearly all vehicles in the sample have modern engines that do not require prolonged warm-up, 76% of drivers believe that warming up the engine increases engine life (U.S. EPA, 2025; U.S. DOE, 2026). Only 39% believe that warming up the engine harms the environment. Policy guidance further indicates that restarting after stops longer than 10 seconds generally saves fuel and reduces emissions (U.S. EPA, 2025; U.S. DOE, 2026; Gaines and Rask, 2012), yet only 6% of drivers report a fuel-saving threshold of 30 seconds or less, and only 15% report an environmental-harm threshold of 30 seconds or less.

Table 2: Summary statistics for the experimental sample

	Median	Mean	SD	p10	p90	N
Panel A. Driver and vehicle profile						
Age (years)	40.00	41.02	11.40	28.00	57.00	6,596
Bachelor's degree or higher (=1)	1.00	0.66	0.47	0.00	1.00	6,589
Engine size ≤1000cc (=1)	0.00	0.49	0.50	0.00	1.00	6,483
Vehicle age ≥5 years (=1)	0.00	0.40	0.49	0.00	1.00	6,596
Panel B. Idling beliefs and preferences						
Stop duration very important for shutoff decision (=1)	0.00	0.21	0.40	0.00	1.00	6,596
Believes warm-up increases engine life (=1)	1.00	0.76	0.43	0.00	1.00	6,596
Believes warm-up harms environment (=1)	0.00	0.39	0.49	0.00	1.00	6,596
Fuel-saving shutoff threshold ≤30 seconds (=1)	0.00	0.06	0.24	0.00	0.00	6,466
Environmental-harm shutoff threshold ≤30 seconds (=1)	0.00	0.15	0.36	0.00	1.00	6,452
WWF-Pakistan donation out of Rs. 100	100.00	89.46	28.38	50.00	100.00	6,479
Panel C. GPS coverage and behaviour						
Share of days observed	1.00	0.87	0.32	0.00	1.00	6,596
Total idling minutes	19.99	27.11	34.74	5.54	54.83	5,846
Idling minutes from stops >30s	13.98	20.44	30.06	3.34	42.84	5,846
Idling spells from stops >30s	8.08	10.26	11.47	2.17	20.42	5,846
Trips	7.33	11.54	35.61	1.92	19.17	5,846

Notes: The unit of observation is the vehicle. The sample comprises the 6,596 vehicles in the final six-arm randomised sample, excluding the 25 removed before the second-stage *Reminder* versus *No-Reminder* randomisation. Panels A and B report baseline characteristics and beliefs; Panel C reports vehicle-level averages of GPS coverage and behaviour over the two clean weeks beginning on 23 and 30 March 2026. Share of days observed is computed over all possible vehicle-days, with unobserved days coded as zero; GPS behavioural outcomes are averaged over observed vehicle-days. Indicators equal zero or one.

The potential returns to reducing sustained idling are substantial. Idling in spells longer than 30 seconds accounts for nearly 18% of total engine-on time across the 12 pre-treatment weeks among vehicles in the analysis sample.

Baseline idling is concentrated in sustained spells rather than brief stops. Figure 3 assigns each idling spell to a mutually exclusive bin based on its total duration and reports the time contributed by each bin as a share of total idling time. Spells lasting more than two minutes alone account for 45% of all idling time, while spells longer than 30 seconds collectively account for roughly three-quarters. This concentration in sustained stops also supports the use of 30 seconds as the threshold for the paper's primary idling outcome.

Figure 3: Composition of baseline idling time by spell duration

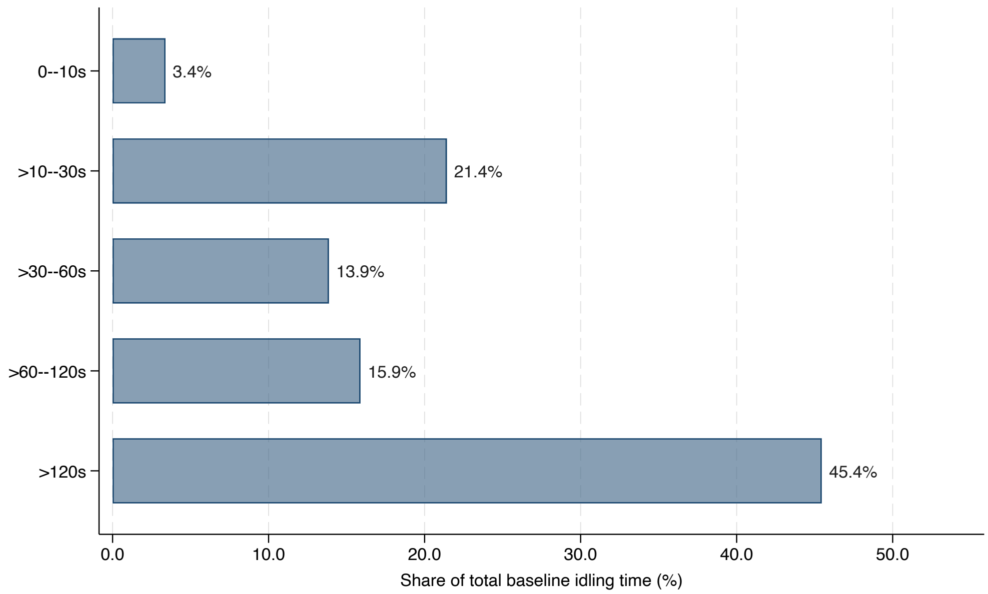

Notes: The figure reports the share of total baseline idling time contributed by mutually exclusive spell-duration bins. The sample comprises the 6,596 vehicles in the final six-arm randomised sample, excluding the 25 whose incentivised choices were implemented before the second-stage *Reminder* versus *No-Reminder* randomisation. Baseline idling is measured over observed GPS vehicle-days during the two clean weeks beginning on 23 and 30 March 2026.

5.2 Treatment Effects by Arm

Table 3 reports treatment effects for each treatment arm relative to the control arm. All five estimates for daily idling from stops longer than 30 seconds are negative. The largest reductions—for *Private Info*, *Private & Social Info + Pay*, and *Private & Social Info + Pay + No-Reminder*—range from 8% to 9% of the control post-period mean.

The *Private Info* arm reduces daily idling from stops longer than 30 seconds by 1.8 minutes, or 9% of the control mean, with the estimate statistically significant at the 1% level. The estimate is similar under a vehicle-level ANCOVA specification (Online Appendix Table D.1). It also reduces the daily number of such spells by 0.5, or 5% of the control mean, with the estimate statistically significant at the 5% level. The effects for *Private & Social Info* and *Private & Social Info + Pay* are broadly similar. The main exception is the smaller duration effect for *Private & Social Info*, which reduces daily idling from stops longer than 30 seconds by 1.3 minutes, or 6% of the control mean.

The estimated effect for *Private & Social Info + Pay + Reminder* is also negative, but it is the only treatment-arm estimate that is statistically indistinguishable from zero, with a 95% confidence interval of [−1.87, 0.41]. By contrast, *Private & Social Info + Pay + No-Reminder* produces larger reductions in both idling duration and spell counts. The next subsection examines this comparison directly, while Section 6 investigates why activating the reminder appears to weaken the effect of the otherwise identical no-reminder arm. Overall, private information meaningfully reduces idling on

Table 3: Treatment effects across behavioural margins

	Idling duration (minutes)			Idle spells (count)		Vehicle use	
	Total	$>$10s	$>$30s	$>$10s	$>$30s	Moving minutes	Trips
Private Info × Post	-1.993***	-1.952***	-1.809***	-1.067	-0.532**	-1.321	-0.325
	(0.737)	(0.719)	(0.654)	(0.705)	(0.229)	(1.739)	(0.559)
Private + Social Info × Post	-1.523**	-1.467**	-1.316**	-1.089	-0.479**	-1.707	-0.551
	(0.722)	(0.702)	(0.629)	(0.737)	(0.242)	(1.967)	(0.448)
Private + Social Info + Pay × Post	-2.087***	-2.029***	-1.873***	-1.058	-0.415*	-2.470	-0.779
	(0.756)	(0.737)	(0.673)	(0.716)	(0.223)	(1.793)	(0.563)
Private + Social Info + Pay + Reminder × Post	-0.826	-0.773	-0.730	-0.024	-0.066	-0.470	-0.341
	(0.665)	(0.643)	(0.582)	(0.637)	(0.206)	(1.629)	(1.123)
Private + Social Info + Pay + No-Reminder × Post	-1.876***	-1.854***	-1.725***	-0.915	-0.453**	-2.545	-0.153
	(0.718)	(0.699)	(0.631)	(0.660)	(0.222)	(1.709)	(0.457)
Control post-period mean	27.721	26.770	20.972	32.590	10.238	80.756	10.567
Observations	701441	701441	701441	701441	701441	701441	701441
Vehicles	6428	6428	6428	6428	6428	6428	6428
Vehicle FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Calendar-date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Joint p-value	0.023	0.020	0.017	0.280	0.051	0.568	0.661

Notes: The unit of observation is the vehicle-day. Each column reports treatment-arm-by-post coefficients from a regression absorbing vehicle and calendar-date fixed effects; the control arm is omitted. Post equals one from 6 April through 30 May 2026; the pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. Columns 1–3 report idling duration, Columns 4–5 idling-spell counts, and Columns 6–7 vehicle-use margins; thresholds refer to total spell duration. Outcomes are computed over observed GPS vehicle-days. The sample excludes the 25 vehicles removed before the second-stage *Reminder* versus *No-Reminder* randomisation. Standard errors are clustered by vehicle. The joint p-value tests whether all displayed coefficients equal zero. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

both the intensive and extensive margins, and these reductions are broadly preserved when social information or payments are added, but not when the real-time reminder is activated.

5.3 Incremental Treatment Effects

Figure 4 reports the incremental effects of social information, payments, and *IdleGuard*. Adding social information increases daily idling from stops longer than 30 seconds by 0.5 minutes relative to *Private Info*, but the estimate is statistically indistinguishable from zero. Its 95% confidence interval of $[-0.70, 1.69]$ rules out large additional reductions in idling. Adding payments decreases idling by 0.6 minutes relative to *Private & Social Info*, but this estimate is also statistically indistinguishable from zero. Its 95% confidence interval of $[-1.79, 0.68]$ does not rule out a meaningful additional reduction.

By contrast, activating *IdleGuard* increases daily idling from stops longer than 30 seconds by 1.0 minute relative to the matched no-reminder arm, equal to 5% of the control post-period mean, with the estimate statistically significant at the 10% level. It also increases the daily number of such spells by 0.4, equal to 4% of the control mean, with the estimate again statistically significant at the 10% level. Thus, social information and payments produce no detectable additional reductions,

Figure 4: Incremental treatment effects on idling minutes from stops longer than 30 seconds

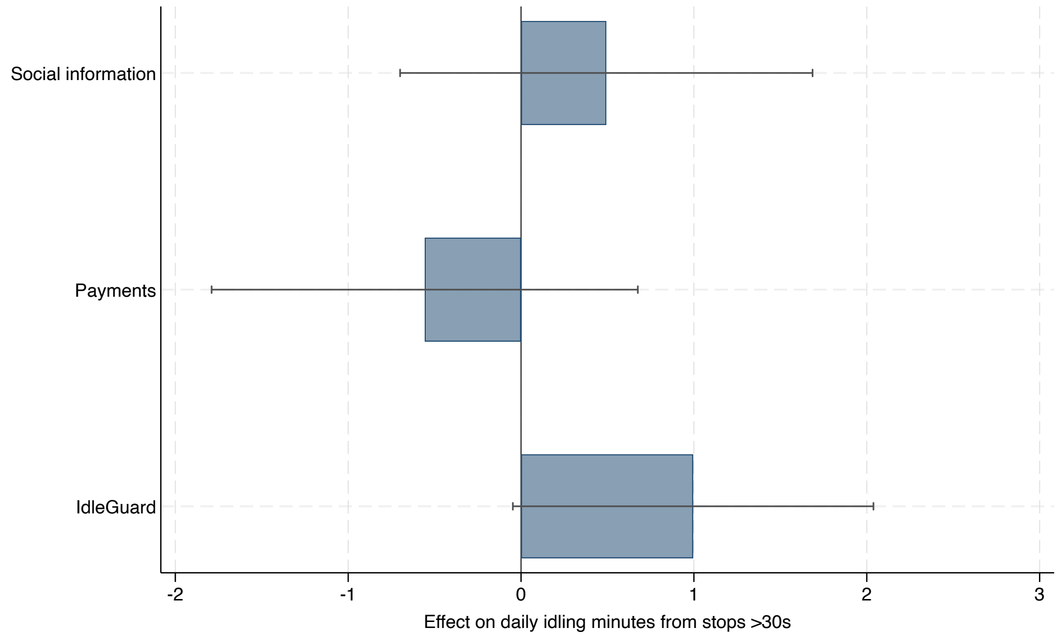

Notes: The figure reports the incremental effects of social information, payments, and *IdleGuard* on daily idling minutes from stops longer than 30 seconds. Positive coefficients indicate increased idling relative to the preceding arm or matched no-reminder arm. Horizontal bars report point estimates, and capped lines report 95% confidence intervals. Online Appendix Table D.2 reports the corresponding estimates, comparison definitions, specification, and inference.

while the real-time reminder weakens the otherwise identical no-reminder arm. Section 6 investigates why.

5.4 Treatment Dynamics

Average treatment effects reported in Section 5.2 may mask meaningful variation across treatment weeks. I therefore examine whether effects emerge immediately or build gradually, fade with repeated exposure, or remain stable over the intervention (Figure 5). The effects of *Private Info* appear early, remain broadly stable, and are statistically significant in most treatment weeks. The estimates for *Private & Social Info + Pay* are generally more negative than those for *Private Info* in the initial weeks. The estimated effects for *Private & Social Info + Pay + Reminder* gradually move towards zero and become positive in the latter weeks, whereas those for the matched *Private & Social Info + Pay + No-Reminder* arm remain negative throughout. Overall, the average treatment effects are not driven by isolated weeks, the weaker performance of the reminder arm becomes more pronounced over the intervention, and there is no general pattern of progressive treatment-effect decay. Online Appendix Figure D.6 extends this analysis to the full pre-treatment study-week path and shows no evidence of differential pre-treatment dynamics.

Figure 5: Dynamic treatment effects on idling minutes from stops longer than 30 seconds

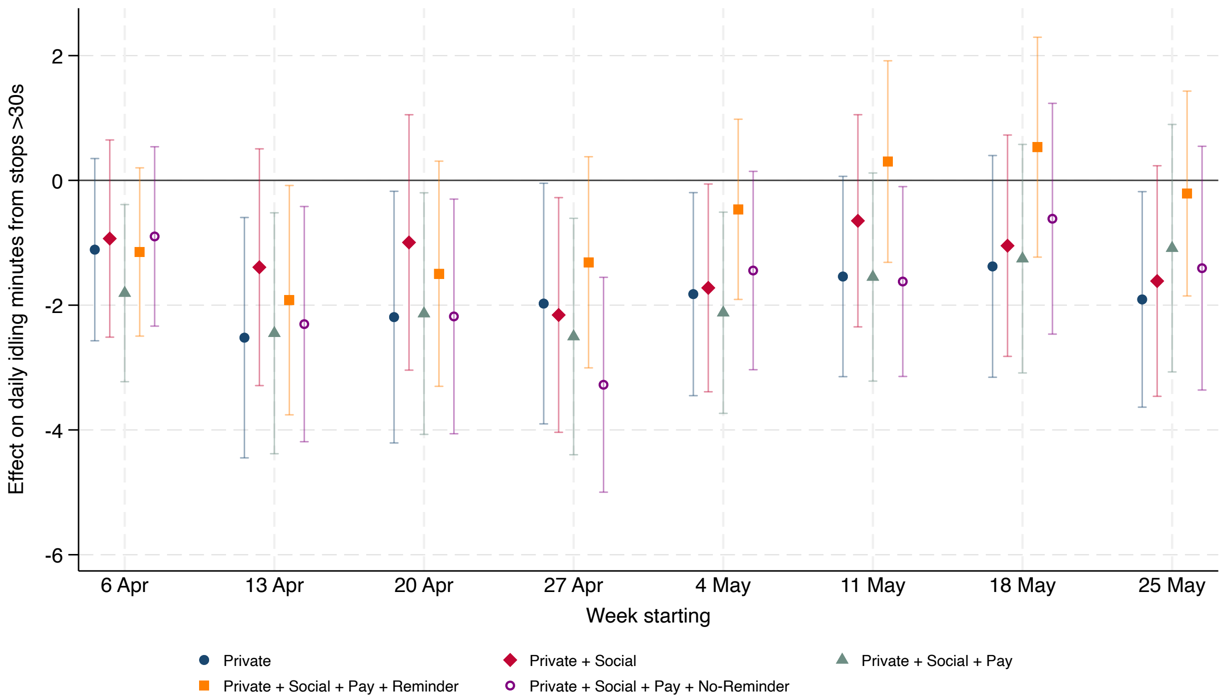

Notes: The figure reports week-specific treatment effects on daily idling minutes from stops longer than 30 seconds relative to the control arm. Points report treatment-arm-by-week coefficients, and capped lines report 95% confidence intervals. The 12 pre-treatment study weeks observed between 17 November 2025 and 4 April 2026 form the pooled reference period. Estimates are horizontally offset within each treatment week for readability. Regressions absorb vehicle and calendar-date fixed effects; outcomes are computed over observed GPS vehicle-days, and standard errors are clustered by vehicle. The sample excludes the 25 vehicles whose incentivised choices were implemented before the second-stage *Reminder* versus *No-Reminder* randomisation.

5.5 Persistence after Treatment Ends

Table 4 reports treatment effects for each treatment arm relative to the control arm across three periods: the eight treatment weeks, the first two post-treatment weeks, and the remaining ten post-treatment weeks. With the exception of *Private & Social Info + Pay + Reminder*, the estimates remain negative over the remaining ten post-treatment weeks, although their magnitude and precision vary across arms.

In the first two post-treatment weeks, *Private Info* reduces daily idling from stops longer than 30 seconds by 2.0 minutes, or 9% of the control mean in the same period, with the estimate statistically significant at the 1% level. This reduction is similar to the estimated effect during the treatment period. The post-treatment *Private Info* estimates pool vehicles whose messages ended after eight weeks with those randomly assigned to receive messages for two additional weeks. The direct continuation and switch-off tests reported in Online Appendix Tables D.12 and D.13 are imprecisely estimated and provide no detectable evidence that continuing the messages for two additional weeks further reduced idling or that idling rebounded when messages stopped. Over the remaining ten post-treatment weeks, the pooled *Private Info* estimate remains negative at 1.1 minutes per day, but is no longer statistically significant. The effect of *Private & Social Info + Pay + No-Reminder* also remains negative at 1.8 minutes per day, or 8% of the contemporaneous control mean, and is statistically significant at the 5% level. For the *Early-Off* subgroup whose messages ended after the scheduled eight-week campaign, the corresponding pooled twelve-week post-switch-off estimate remains negative and is reported in Online Appendix Table D.14.

Online Appendix Table D.11 reports the persistence of the incremental treatment effects. Prior assignment to *IdleGuard* is associated with 1.5 minutes more idling relative to the matched no-reminder arm in the first two post-treatment weeks, equal to 7% of the control mean in the same period, with the estimate statistically significant at the 10% level. This difference grows to 2.1 minutes over the remaining ten post-treatment weeks, equal to 10% of the control mean, and is statistically significant at the 1% level.

Taken together, the treatment effects show substantial but heterogeneous persistence across the twelve weeks after the common intervention ended. The estimates for *Private Info* and *Private & Social Info + Pay + No-Reminder* remain negative relative to the control arm over the remaining ten post-treatment weeks. By contrast, the estimate for *Private & Social Info + Pay + Reminder* is close to zero over this period. The resulting positive Reminder–No-Reminder difference persists after *IdleGuard* is no longer active, consistent with more durable behaviour change in the no-reminder arm than in the otherwise identical reminder arm.

5.6 Where in the Idling Distribution Do Effects Arise?

Baseline idling is concentrated in sustained spells. As shown in Figure 3, spells lasting more than two minutes account for 45% of total baseline idling time, while spells longer than 30 seconds collectively account for roughly three-quarters. I therefore examine where along the spell-duration distribution the treatment effects arise. This pattern is informative for both the behavioural interpretation of the results and the policy relevance of the interventions.

Online Appendix Table D.5 shows that treatment effects relative to the control arm are concentrated in spells lasting more than two minutes. For *Private Info*, approximately 79% of the total effect on idling from spells longer than 30 seconds comes from the $>$ 120-second bin: 1.43 of the 1.81 minutes reduced per vehicle-day. The other treatment arms broadly follow the same pattern. The corresponding shares are approximately 71% for *Private & Social Info*, 88% for *Private & Social Info + Pay*, 85% for *Private & Social Info + Pay + Reminder*, and 83% for *Private & Social Info + Pay + No-Reminder*. Consistent with this concentration, the joint test rejects the null of no treatment effects in the $>$ 120-second bin ($p = 0.026$), while the corresponding tests for the shorter-duration bins do not reject.

For the incremental treatment effects, Online Appendix Table D.6 shows that approximately 81% of the increase in idling caused by activating *IdleGuard*, relative to the matched no-reminder arm, arises in the $>$ 120-second bin. In sum, the treatment effects are concentrated in longer idling

Table 4: Persistence of treatment effects on idling minutes from stops longer than 30 seconds

	Idling minutes from stops >30s		
	Treatment weeks	First 2 post weeks	Remaining post weeks
Private Info × Post	-1.870***	-2.025***	-1.142
	(0.660)	(0.782)	(0.778)
Private + Social Info × Post	-1.329**	-0.755	-0.785
	(0.631)	(0.914)	(0.691)
Private + Social Info + Pay × Post	-1.910***	-1.335	-1.020
	(0.672)	(0.873)	(0.769)
Private + Social Info + Pay + Reminder × Post	-0.750	0.122	0.337
	(0.581)	(0.771)	(0.631)
Private + Social Info + Pay + No-Reminder × Post	-1.698***	-1.420	-1.768**
	(0.631)	(0.885)	(0.709)
Control period mean	20.972	22.032	22.070
Observations	1,076,021	1,076,021	1,076,021
Vehicles	6,429	6,429	6,429
Vehicle FE	Yes	Yes	Yes
Calendar-date FE	Yes	Yes	Yes
Joint p-value	0.015	0.043	0.027

Notes: The unit of observation is the vehicle-day, and the outcome is daily idling minutes from stops longer than 30 seconds. Each column reports treatment-arm-by-period coefficients from a regression absorbing vehicle and calendar-date fixed effects; the control arm is omitted. The pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. The treatment period runs from 6 April through 30 May 2026, followed by two post-treatment weeks from 1 through 13 June and ten from 15 June through 22 August 2026. The *Private Info* coefficients pool *Early-Off* and *Late-On* vehicles in both post-treatment periods. Outcomes are computed over observed GPS vehicle-days, and standard errors are clustered by vehicle. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

spells, and the increase in idling caused by activating *IdleGuard* does not appear to reflect merely mechanical bunching around its first beep at 30 seconds. Instead, the increase arises primarily in sustained idling. I examine the mechanisms underlying this pattern in the next section.

6 Mechanisms

6.1 Evidence of Belief Correction

Did information work through beliefs? Table 2 documents widespread misperceptions about both the consequences of idling and when shutting off the engine becomes beneficial. Table 5 directly tests whether *Private Info* corrected the two private-cost beliefs targeted by the treatment. Relative to the control arm, *Private Info* reduces by 6.3 percentage points the share of respondents who believe that warming up the engine increases engine life, from a control mean of 81.1%. It also increases by 5.2 percentage points the share reporting that shutting off the engine during a stop of 30 seconds or less saves fuel, from a control mean of 9.2%. Both estimates are statistically significant at the 1% level and adjust for the corresponding baseline belief.

Table 5: Effects of *Private Info* on endline private-cost beliefs

	Endline private-cost beliefs	
	Believes warm-up increases engine life	Fuel-saving shutoff threshold $\leq$ 30 seconds
Private Info	-0.063***	0.052***
	(0.020)	(0.016)
Control mean	0.811	0.092
Observations	1628	1558
Baseline outcome control	Yes	Yes

Notes: The unit of observation is the vehicle. The sample comprises vehicles assigned to *Private Info* or *Control* with a usable endline outcome and the corresponding baseline belief. Column 1 equals one if the respondent believes that warming up the engine increases engine life. Column 2 equals one if the respondent reports a fuel-saving shutoff threshold of 30 seconds or less. Each column compares *Private Info* with *Control* and adjusts for the corresponding baseline belief. A negative estimate in Column 1 and a positive estimate in Column 2 indicate belief correction. Endline availability does not differ significantly between the two groups for either outcome. Robust standard errors are shown in parentheses. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

These effects show that the information treatment moved posterior beliefs in the intended direction. Consistent with this interpretation, Online Appendix Table E.2 shows larger behavioural reductions among drivers whose baseline beliefs left more scope for correction, although the differences across belief groups are imprecisely estimated. Taken together with the reduction in idling reported in Section 5.2, the evidence indicates that *Private Info* changed both the targeted beliefs and behaviour. It does not, however, establish that individual-level belief updating fully mediates the behavioural response.

6.2 Why Do Social Information and Payments Add Little?

Although the incremental effect of social information is statistically indistinguishable from zero on average, this masks informative heterogeneity by baseline fuel-saving and environmental misperceptions. Table 6 reports effects on daily idling minutes from stops longer than 30 seconds by whether the respondent's baseline fuel-saving shutoff threshold was below or at least 300 seconds. It shows three comparisons: (1) the effect of *Private Info* relative to the control arm, (2) the effect of *Private & Social Info* relative to the control arm, and (3) the incremental effect of social information. Among less-mistaken drivers, who reported a threshold below 300 seconds, *Private Info* reduces idling relative to the control arm, while *Private & Social Info* produces a larger reduction that is statistically significant at the 5% level. Adding social information therefore does not offset the private-information effect and may strengthen it, although the incremental effect is imprecisely estimated. More-mistaken drivers respond about twice as strongly to *Private Info* as less-mistaken drivers, consistent with belief updating. Among these drivers, however, adding social information increases idling relative to *Private Info*, reversing much of the reduction generated by private information. The incremental effect of social information differs significantly across the two belief groups at the 5% level. This pattern is consistent with a dilution mechanism: adding the broader social-information layer weakens the targeted private-cost correction precisely where that correction is most valuable.

Online Appendix Table E.5 reveals a similar pattern for beliefs about the shutoff threshold for environmental harm. The two environmental-belief groups respond similarly to private information alone. Once environmental information is added, however, drivers with more contrary baseline beliefs respond substantially less, a pattern suggestive of belief conflict. I probe this interpretation using respondents' confidence in their threshold beliefs. If belief conflict drives this pattern, social information should weaken the private-information effect most among drivers whose mistaken environmental beliefs are held with greater confidence. Consistent with this prediction, the increase in idling caused by adding social information is larger among drivers with strong environmental misperceptions—those who reported a threshold of at least 300 seconds and were confident or very confident in that response.

Endline evidence sharpens these interpretations. Adding social information does not measurably improve either of two environmental beliefs relevant to its content: whether engine warm-up harms the environment or whether shutting off within 30 seconds reduces environmental harm (Online Appendix Table E.7). Direct endline measures of message overload, action relevance, and demand for continued messages provide no independent support for dilution (Online Appendix Table E.8). By contrast, among respondents who report reading the treatment messages, adding social information reduces low-credibility ratings among drivers with less contrary environmental priors, but not among those with more contrary priors; the difference in effects is marginally statistically significant ($p = 0.059$; Online Appendix Table E.9). This pattern provides suggestive evidence that more contrary priors limit the credibility of social information, consistent with the belief-conflict interpretation.

The incremental effect of payments is also statistically indistinguishable from zero, although its point estimate is negative. To reduce concerns that this result reflects limited power, Online Appendix Table E.10 pools the two non-*IdleGuard* payment arms. The pooled estimate remains modest and statistically insignificant. Conceptually, payments were designed to test whether idling that remained after information provision reflected genuine private benefits from continued idling. Taken together, the results suggest that private information addresses the principal binding constraint, leaving limited scope for payments to reduce idling further at the implemented incentive level.

Table 7 supports this interpretation. It defines environmental scepticism using beliefs about the environmental harm from engine warm-up and the corresponding shutoff threshold, and examines heterogeneity in the effect of *Private Info* and the incremental effects of social information and payments. Among less-sceptical drivers, *Private Info* reduces idling, while adding social information modestly strengthens this effect. Consistent with limited residual price-responsive idling at the implemented incentive level, adding payments yields little further reduction. Among more-sceptical drivers, however, adding social information increases idling relative to *Private Info*, consistent with the belief-conflict mechanism described above. This creates scope for payments to substitute for

Table 6: Social-information effects by baseline fuel-saving threshold beliefs

	Idling minutes from stops >30s
Panel A. Fuel-saving threshold below 300 seconds	
Private Info	-1.204
	(0.803)
Private & Social Info	-2.061**
	(0.849)
Social information	-0.856
	(0.949)
Panel B. Fuel-saving threshold at least 300 seconds	
Private Info	-2.310***
	(0.791)
Private & Social Info	-0.649
	(0.689)
Social information	1.661**
	(0.798)
Panel C. At least 300 seconds minus below 300 seconds	
Private Info	-1.106
	(0.898)
Private & Social Info	1.411*
	(0.855)
Social information	2.517**
	(1.240)
Control post-period mean, threshold below 300s	20.858
Control post-period mean, threshold at least 300s	21.097
Observations	325308
Vehicles	2976
Vehicle FE	Yes
Calendar-date FE	Yes

Notes: The unit of observation is the vehicle-day, and the outcome is daily idling minutes from stops longer than 30 seconds. Panels A and B report effects for drivers whose baseline fuel-saving shutoff threshold was below and at least 300 seconds, respectively; Panel C reports the corresponding at-least-300-seconds-minus-below-300-seconds differences. Within each panel, rows report the effects of *Private Info* and *Private & Social Info* relative to the control arm and the incremental effect of social information. Regressions absorb vehicle and calendar-date fixed effects. Post equals one from 6 April through 30 May 2026; the pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. Outcomes are computed over observed GPS vehicle-days. The sample excludes the 25 vehicles removed before the second-stage *Reminder* versus *No-Reminder* randomisation. Standard errors are clustered by vehicle. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

ineffective social motivation: adding payments largely restores the reduction achieved by private information alone. The limited average gain from payments therefore appears to reflect less a general ineffectiveness of financial incentives than the fact that there is usually little residual idling for them to address.

6.3 Why Does *IdleGuard* Increase Idling?

IdleGuard was designed to restore attention at the moment of choice, but it appears to have triggered a competing pathway: reduced self-monitoring. Activating *IdleGuard* increases daily idling from stops longer than 30 seconds by roughly one minute relative to the matched no-reminder arm, equal to about 5% of the control post-period mean. The estimate is statistically significant at the

Table 7: Information and payment effects by environmental scepticism

	Idling minutes from stops >30s
Panel A. Not environmentally sceptical	
Private Info	-1.462**
	(0.740)
Social information	-0.366
	(0.753)
Payments	0.432
	(0.722)
Panel B. Environmentally sceptical	
Private Info	-2.591***
	(0.985)
Social information	2.563**
	(1.151)
Payments	-3.053**
	(1.500)
Panel C. Sceptical minus not sceptical	
Private Info	-1.128
	(1.013)
Social information	2.929**
	(1.375)
Payments	-3.486**
	(1.665)
Control post-period mean, not environmentally sceptical	20.388
Control post-period mean, environmentally sceptical	22.160
Observations	663203
Vehicles	6076
Vehicle FE	Yes
Calendar-date FE	Yes

Notes: The unit of observation is the vehicle-day, and the outcome is daily idling minutes from stops longer than 30 seconds. Environmentally sceptical drivers did not report that warming up the engine harms the environment and reported an environmental-harm shutoff threshold of at least 300 seconds; unsure warm-up responses are excluded. Panels A and B report effects for non-sceptical and sceptical drivers, and Panel C reports sceptical-minus-non-sceptical differences. Rows report the effect of *Private Info* relative to the control arm and the incremental effects of social information and payments. Regressions absorb vehicle and calendar-date fixed effects. Post equals one from 6 April through 30 May 2026; the pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. Outcomes are computed over observed GPS vehicle-days. The sample excludes the 25 vehicles removed before the second-stage *Reminder* versus *No-Reminder* randomisation. Standard errors are clustered by vehicle. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

10% level (Online Appendix Table D.2). Heterogeneity and telemetry analyses provide internally consistent evidence that substitution away from self-monitoring may explain this effect.

Who drives the increase? Table 8 shows that it is concentrated among drivers with high baseline overall idling and high baseline *warming-like* idling. Table 9 also shows larger adverse point estimates among drivers with greater pre-treatment demand for the device: those who chose the device over cash, said they would use it if offered for free, or expected it to generate positive fuel savings. The corresponding differences relative to lower-demand drivers are statistically indistinguishable from zero.

Taken together, the increase appears strongest among drivers who had more idling to monitor and who expressed greater demand for, or confidence in, the reminder technology. This pattern is difficult to reconcile with the device simply being irrelevant. Instead, it is consistent with drivers reducing their own monitoring when they expected to rely on *IdleGuard*, thereby weakening part of the behavioural response observed in the matched no-reminder arm.

IdleGuard beeped after 30 seconds of idling and repeated the alert every 30 seconds while the spell persisted. If the device substituted for drivers' own monitoring, more spells should persist through the first 30 seconds and cross the reminder threshold under *IdleGuard* than under the matched no-reminder arm. This is what I find. Online Appendix Table E.14 shows that the no-reminder arm reduces the daily number of spells longer than 30 seconds by 0.5 relative to the control arm, with the estimate statistically significant at the 5% level. Activating *IdleGuard* increases this outcome by 0.4 spells relative to the matched no-reminder arm, with the estimate statistically significant at the 10% level. Approximately 87% of this increase arises in *traffic-like* idling. The pattern is consistent with preventive self-monitoring in the no-reminder arm being weakened by the device.

The threshold-crossing difference persists after the reminder is removed. Over the twelve post-treatment weeks, vehicles previously assigned to the no-reminder arm have 0.6 fewer spells longer than 30 seconds per day than the control arm, while the corresponding estimate for vehicles previously assigned to *IdleGuard* is close to zero (Online Appendix Table E.15). The resulting prior-*IdleGuard*-assignment difference is 0.9 spells per day, is statistically significant at the 1% level, and is concentrated primarily in *traffic-like* idling. Because the reminder was no longer active during this period, the persistence of the gap is consistent with more durable preventive self-monitoring in the no-reminder arm.

Because the primary outcome counts the full duration of every spell longer than 30 seconds, rather than only time beyond 30 seconds, the increase caused by *IdleGuard* could arise mechanically from more spells crossing the threshold. For example, if a driver in the no-reminder arm switches off after 20 seconds while an *IdleGuard* driver continues to 35 seconds, the outcome records zero for the former and the full 35 seconds for the latter, even though their behaviour differs by only 15 seconds. Online Appendix Table E.16 shows that this mechanical feature is not the main driver: only about 19% of the increase in idling caused by *IdleGuard* comes from the first 30 seconds of spells crossing the threshold, while about 81% comes from idling beyond 30 seconds. The decomposition also shows that drivers do not generally correct immediately after the first beep by switching off or moving. More than half of the full gap—56%—comes from idling beyond two minutes.

Online Appendix Table E.17 confirms that the increase is concentrated in sustained spells. Spells lasting longer than two minutes account for 81% of the full increase in idling caused by *IdleGuard*. Activating *IdleGuard* increases daily idling from these spells by 0.8 minutes relative to the matched no-reminder arm, equal to about 7% of the corresponding no-reminder post-period mean, with the estimate statistically significant at the 10% level. It also increases their daily number by 0.13 spells, equal to about 6% of the no-reminder post-period mean, with the estimate statistically significant at the 5% level.

Endline responses provide complementary, though descriptive, evidence for both stages. Among

Table 8: Incremental *IdleGuard* effects by baseline idling intensity

	Idling minutes from stops >30s	
	Overall	Warming-like
Low baseline	-0.134	-0.383
	(0.459)	(0.513)
High baseline	2.051**	2.288**
	(0.957)	(0.950)
Difference: high minus low	2.185**	2.671**
	(1.062)	(1.080)
No-Reminder post-period mean, low baseline	9.776	12.889
No-Reminder post-period mean, high baseline	30.061	28.080
Observations	254782	254782
Vehicles	2200	2200
Vehicle FE	Yes	Yes
Calendar-date FE	Yes	Yes

Notes: The unit of observation is the vehicle-day, and the outcome is daily idling minutes from stops longer than 30 seconds. The sample comprises vehicles assigned to the *Reminder* and matched *No-Reminder* arms. Each column reports the incremental effect of *IdleGuard* for vehicles with low and high baseline idling. Column 1 defines baseline intensity using the vehicle-level average of total idling over the two clean pre-treatment GPS weeks beginning on 23 and 30 March 2026; Column 2 uses the corresponding average of *warming-like* idling. Low and high denote vehicles at or below and above the median, respectively. Difference rows report high-minus-low effects. Regressions absorb vehicle and calendar-date fixed effects. Post equals one from 6 April through 30 May 2026; the pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. Outcomes are computed over observed GPS vehicle-days, and standard errors are clustered by vehicle. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

drivers assigned to *IdleGuard* who reported whether the reminder changed how comfortable they were leaving the engine running, 42% said that it made them more comfortable doing so because the beep would remind them, consistent with reduced preventive self-monitoring. In a separate question, 60% of drivers who reported how they responded after the beep said that they did not usually correct promptly, consistent with incomplete corrective action. By contrast, random assignment to *IdleGuard*, relative to the matched no-reminder arm, has no detectable effect on endline beliefs about the amount of engine warm-up required. Online Appendix Tables E.19 and E.20 report these results.

Taken together, the evidence suggests a two-stage response: activating *IdleGuard* weakens preventive self-monitoring, allowing more spells to cross the 30-second threshold, while corrective action after the reminder is delayed, causing those spells to persist. The remaining question is why *IdleGuard* drivers "snooze" after being reminded. Section 6.4 develops a procrastination model, drawing on O'Donoghue and Rabin (1999), that provides a natural explanation for this delayed response.

6.4 Why Do *IdleGuard* Drivers Snooze?

The telemetry evidence above suggests that drivers "delegate" monitoring to the device: activating *IdleGuard* makes idling spells more likely to reach the first beep relative to the matched no-reminder

Table 9: Incremental *IdleGuard* effects by pre-treatment demand for the device

	Idling minutes from stops >30s		
	Device choice	Would use if free	Expected fuel savings
Lower demand	0.305	1.158	0.129
	(0.836)	(1.184)	(1.528)
Higher demand	1.791**	1.133*	1.910**
	(0.793)	(0.642)	(0.793)
Difference: higher minus lower	1.486	-0.025	1.781
	(1.151)	(1.347)	(1.721)
No-Reminder post-period mean, lower demand	20.486	19.990	19.984
No-Reminder post-period mean, higher demand	19.470	19.925	20.780
Observations	221962	222801	160666
Vehicles	2021	2028	1461
Vehicle FE	Yes	Yes	Yes
Calendar-date FE	Yes	Yes	Yes

Notes: The unit of observation is the vehicle-day, and the outcome is daily idling minutes from stops longer than 30 seconds. The sample comprises vehicles assigned to the *Reminder* and matched *No-Reminder* arms. Each column reports the incremental effect of *IdleGuard* for drivers with lower and higher pre-treatment demand for the device. In Column 1, lower and higher demand denote choosing cash and the device, respectively, in the incentivised briefing-stage elicitation. In Column 2, they denote reporting that the driver would not or would use the device if offered for free; in Column 3, they denote reporting zero or positive expected fuel savings. Difference rows report higher-minus-lower-demand effects. Regressions absorb vehicle and calendar-date fixed effects. Post equals one from 6 April through 30 May 2026; the pooled pre-treatment period comprises 12 study weeks observed between 17 November 2025 and 4 April 2026. Outcomes are computed over observed GPS vehicle-days. Each column includes drivers with a valid response to the corresponding measure. The sample excludes the 25 vehicles removed before the second-stage *Reminder* versus *No-Reminder* randomisation. Standard errors are clustered by vehicle. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

arm. Yet even after being reminded, drivers do not generally switch off the engine or move promptly enough to prevent these spells from becoming sustained. I explain this post-reminder delay using the simple immediate-cost procrastination model developed by O'Donoghue and Rabin (1999), adapted here to a setting with repeated opportunities to act.

Setup. The model begins conditional on an idling spell having reached a reminder beep; it therefore explains post-reminder delay rather than whether the spell reaches the first beep. Time is discrete, and each period corresponds to the 30 seconds between successive *IdleGuard* beeps. Let t denote the instant at which a beep sounds and the driver chooses whether to act, and suppose the driver expects the stop to end at time T. The driver can switch off immediately or continue idling and reconsider at the next beep. Switching off at t entails an immediate one-time action cost $k > 0$, capturing the cognitive and physical effort of responding to the reminder, operating the ignition, and later restarting the vehicle. It also immediately begins generating a constant net benefit flow. I denote by $b > 0$ the net benefit generated over each subsequent 30-second interval, combining fuel and maintenance gains less the flow cost of reduced comfort or convenience. For notational convenience, the flow generated over the interval beginning at t is indexed by its endpoint $t + 1$. The formal exposition focuses on waiting and traffic idling, where the relevant action is switching off; the same timing logic applies to warming idling, where the relevant action is instead to begin moving.

Time-consistent benchmark. Consider first a time-consistent driver with exponential discount factor $\delta \in (0, 1]$. Relative to leaving the engine on for the remainder of the stop, the value at time t of switching off immediately is

$$V_t^{TC}(\text{switch now}) = -k + \sum_{s=t+1}^{T} \delta^{s-t} b. \tag{4}$$

Switching off therefore has positive conventional net present value whenever

$$\sum_{s=t+1}^{T} \delta^{s-t} b > k. \tag{5}$$

This condition can hold over a sufficiently long stop even when the net benefit from any single 30-second interval is small relative to the action cost.

Present-biased choice. Now suppose the driver has quasi-hyperbolic (β, δ) preferences, where $\beta \in (0, 1]$ captures present bias. The action cost incurred at t receives full weight, while payoffs accruing after the decision—including the net-benefit flow generated over the ensuing 30-second interval—are discounted by β. At time t, switching off immediately is therefore valued at

$$V_t(\text{switch now}) = -k + \beta \sum_{s=t+1}^{T} \delta^{s-t} b. \tag{6}$$

This asymmetry may make delaying action by one beep attractive. If the driver plans to switch off at the next beep, the action cost is shifted to $t + 1$ and is therefore discounted by both β and δ, while the net benefit from the intervening 30-second interval is forgone. The value of this plan at time t is

$$V_t(\text{switch next beep}) = -\beta\delta k + \beta \sum_{s=t+2}^{T} \delta^{s-t} b. \tag{7}$$

The two options yield the same benefits from $t + 2$ onward. Comparing Equations 6 and 7, the driver switches off immediately if and only if

$$\begin{aligned} -k + \beta\delta b &\geq -\beta\delta k, \\ \beta\delta b &\geq (1 - \beta\delta)k. \end{aligned} \tag{8}$$

The left-hand side is the value at time t of the net-benefit flow generated over the first 30-second interval after switching off, which is forgone by waiting. The right-hand side is the current-self gain from postponing the action cost: paying k immediately rather than valuing the same cost at only

$\beta\delta k$ next period. The driver therefore switches off immediately only when the former is at least as large as the latter.

Condition 5 can hold while Condition 8 fails: switching off is worthwhile over the remainder of the stop, even though the current self prefers postponing it until the next beep. Intuitively, the driver delays action not because switching off is undesirable overall, but because moving the action cost out of the present is attractive enough to outweigh the loss of one 30-second net-benefit flow.

Repeated deferral. Repeated delay follows naturally for a naïve present-biased driver. At time t, the driver plans to switch off at the next beep but incorrectly expects the future self to carry out that plan. When $t + 1$ arrives, however, the action cost is again immediate, and the same comparison recurs. A sequence of intended one-period delays can therefore produce a sustained idling spell. A sophisticated driver anticipates this future temptation and may instead switch off earlier to avoid repeated deferral. Because the telemetry data reveal actions rather than beliefs about future behaviour, they cannot distinguish naïve from sophisticated present bias.

Connection to the evidence. The model links the two stages of the *IdleGuard* effect. Activating the device makes idling spells more likely to reach the 30-second reminder threshold, consistent with weaker preventive self-monitoring. Once that threshold is reached, the reminder does not generally prompt immediate correction: most of the resulting increase in idling accrues after 30 seconds, and more than half accrues beyond two minutes. Repeated present-biased deferral therefore offers a parsimonious explanation for how a reminder can restore attention without inducing prompt action. The telemetry evidence is consistent with this mechanism but does not identify present bias or distinguish naïve from sophisticated drivers.

7 Welfare and Policy Implications

7.1 Physical Benefits

I translate the effect of *Private Info* relative to the control arm—a reduction of 1.952 daily idling minutes from stops longer than 10 seconds per observed vehicle-day—into fuel savings and avoided CO_2 and $PM_{2.5}$ emissions. I calculate cumulative impacts over four evaluation horizons: the eight-week treatment period; the 20-week period from treatment launch through the twelve observed post-treatment weeks; and two one-year scenarios with faster and slower decay after treatment ends. Online Appendix Table F.2 reports the resulting fuel savings and avoided emissions per treated vehicle. Although the observed-period savings are modest for an individual vehicle, they become substantial when aggregated across many vehicles. These physical impacts form the basis for the cost–benefit analysis below.

Table 10: Cost–benefit results per treated vehicle

Evaluation horizon	Gross benefit	Programme cost	Net benefit	Benefit–cost ratio
8-week treatment period	1,375.55	375.68	999.87	3.66
20-week observed period	3,200.13	375.68	2,824.45	8.52
1 year: faster decay	4,264.47	375.68	3,888.79	11.35
1 year: slower decay	5,632.89	375.68	5,257.21	14.99

Notes: Benefits and costs are reported in PKR per treated vehicle. Gross benefits comprise avoided private fuel expenditure, global climate damages, and local $PM_{2.5}$-related health damages. Programme cost is the per-vehicle cost of app, email, and WhatsApp delivery over the eight-week campaign. The 20-week horizon includes twelve observed post-treatment weeks. The one-year scenarios assume linear decay to zero either 26 weeks after treatment ends or by campaign week 52. Net benefit equals gross benefit minus programme cost; the benefit–cost ratio equals gross benefit divided by programme cost. The underlying fuel-use, emissions, exposure, valuation, and cost inputs are documented in the Online Appendix.

7.2 Monetised Benefits and Cost-Effectiveness

I monetise the private fuel savings and avoided CO_2 and $PM_{2.5}$ emissions and compare the resulting benefits with programme costs. Gross benefits comprise avoided fuel expenditure, global climate damages, and local $PM_{2.5}$-related health damages. Programme costs comprise app, email, and WhatsApp delivery over the eight-week messaging campaign. Table 10 reports gross benefits, programme costs, net benefits, and benefit–cost ratios per treated vehicle across the four evaluation horizons. The resulting net benefits are positive under every horizon considered, with benefit–cost ratios ranging from 3.66 over the treatment period to 14.99 under the slower-decay one-year scenario.

7.3 Scalability and Policy Design

Digital messaging campaigns are inexpensive to scale. GPS trackers are widespread among cars in Pakistan because banks and other financial institutions commonly require them for financed vehicles as a means of protecting the underlying asset. Bank financing accounted for an estimated 40–45% of new-car sales during the recent expansion in automobile lending, so many passenger vehicles already generate telematics data (Khan, 2021). This existing infrastructure allows personalised anti-idling messages to be constructed from vehicle data that are already collected and delivered at low marginal cost. Less personalised anti-idling information could also be distributed more broadly through WhatsApp or text messages without requiring vehicle-level telemetry.

I illustrate the potential scale of such a campaign by applying the per-vehicle estimates to alternative shares of Pakistan's registered passenger-vehicle stock. Under the preferred illustrative reach assumption of 25%, corresponding to approximately 1.22 million vehicles, the slower-decay one-year scenario implies fuel savings of about 11.5 million litres and avoided emissions of approximately 28,800 tonnes of CO_2 and 10.7 tonnes of $PM_{2.5}$. The corresponding estimated national net benefit is approximately PKR 6.4 billion, or USD 23.1 million. These calculations illustrate the potential scale of the programme rather than forecasting realised coverage or aggregate effects.

8 Conclusion

This paper studies why people persist in a repeated conservation behaviour for which reducing activity can generate both private and social gains. In a large-scale field experiment with 6,621 vehicles in Pakistan, I use an additive design to distinguish among misperceptions, social motivation, private preferences, and inattention as explanations for vehicle idling. Correcting private-cost misperceptions reduces idling by 1.8 minutes per vehicle-day, or 9% of the control mean, and the effect persists after treatment ends. Adding social information or payments produces little additional reduction. By contrast, activating a real-time in-vehicle reminder increases idling relative to an otherwise identical no-reminder arm. The effects of both private information and the reminder are concentrated in sustained idling spells.

The results point to two distinct behavioural mechanisms. Private information changes the beliefs it targets and reduces idling, indicating that misperceptions can sustain inefficient behaviour even in the repeated use of a durable that people already own. The reminder generates a different response. Drivers appear to rely on *IdleGuard* in place of their own preventive monitoring, allowing more idling spells to reach the reminder threshold, and then fail to correct promptly once reminded. Most of the resulting increase occurs well after the first beep, consistent with procrastination when acting entails an immediate cost while the benefits accrue subsequently. Externally supplied attention can therefore substitute for self-monitoring rather than simply augment it.

These findings have implications for the design of corrective environmental policy. The private-information intervention is inexpensive to deliver, generates benefits that exceed programme costs, and can be scaled using existing digital and telematics infrastructure. More broadly, the results caution against treating persistent harmful behaviour as evidence that stronger incentives or additional intervention layers are necessarily required. Effective policy instead depends on identifying what sustains the behaviour—whether mistaken beliefs, genuine private benefits, weak social motivation, or failures of attention and follow-through—and targeting that margin directly. In repeated conservation settings, a simple intervention aimed at the relevant distortion may outperform a more intensive one.

References

Allcott, Hunt and Cass R. Sunstein, "Regulating Internalities," *Journal of Policy Analysis and Management*, 2015, *34* (3), 698–705.

_ **and Christopher Knittel**, "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments," *American Economic Journal: Economic Policy*, 2019, *11* (1), 1–37.

_ **and Dmitry Taubinsky**, "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 2015, *105* (8), 2501–2538.

_ **and Richard L. Sweeney**, "The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment," *Management Science*, 2017, *63* (1), 21–39.

_ **and Sendhil Mullainathan**, "Behavior and Energy Policy," *Science*, 2010, *327* (5970), 1204–1205.

_ **, Christopher Knittel, and Dmitry Taubinsky**, "Tagging and Targeting of Energy Efficiency Subsidies," *American Economic Review*, 2015, *105* (5), 187–191.

_ **, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow**, "The Welfare Effects of Social Media," *American Economic Review*, 2020, *110* (3), 629–676.

_ **, Sendhil Mullainathan, and Dmitry Taubinsky**, "Energy Policy with Externalities and Internalities," *Journal of Public Economics*, 2014, *112*, 72–88.

Baumol, William J. and Wallace E. Oates, "The Use of Standards and Prices for Protection of the Environment," *The Swedish Journal of Economics*, 1971, *73* (1), 42–54.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak, "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 1964, *9* (3), 226–232.

Berkouwer, Susanna B. and Joshua T. Dean, "Credit, Attention, and Externalities in the Adoption of Energy Efficient Technologies by Low-Income Households," *American Economic Review*, 2022, *112* (10), 3291–3330.

Bronchetti, Erin T., Judd B. Kessler, Ellen B. Magenheim, Dmitry Taubinsky, and Eric Zwick, "Is Attention Produced Optimally? Theory and Evidence from Experiments with Bandwidth Enhancements," *Econometrica*, 2023, *91* (2), 669–707.

Carleton, Tamma and Michael Greenstone, "Updating the United States Government's Social Cost of Carbon," Working Paper 2021-04, Becker Friedman Institute for Economics at the University of Chicago November 2021.

Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky, "The Limits of Simple Implementation Intentions: Evidence from a Field Experiment on Making Plans to Exercise," *Journal of Health Economics*, 2018, *62*, 95–104.

Carrico, Amanda R., Paul Padgett, Michael P. Vandenbergh, Jonathan Gilligan, and Kenneth A. Wallston, "Costly myths: An analysis of idling beliefs and behavior in personal motor vehicles," *Energy Policy*, 2009, *37* (8), 2881–2888.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cook, Cody, Aboudy Kreidieh, Shoshana Vasserman, Hunt Allcott, Neha Arora, Freek van Sambeek, Andrew Tomkins, and Eray Turkel, "The Short-Run Effects of Congestion Pricing in New York City," NBER Working Paper 33584, National Bureau of Economic Research 2025.

Damgaard, Mette Trier and Christina Gravert, "The Hidden Costs of Nudging: Experimental Evidence from Reminders in Fundraising," *Journal of Public Economics*, 2018, *157*, 15–26.

Davis, Lucas W., "The Effect of Driving Restrictions on Air Quality in Mexico City," *Journal of Political Economy*, 2008, *116* (1), 38–81.

_ **and Gilbert E. Metcalf**, "Does Better Information Lead to Better Choices? Evidence from Energy-Efficiency Labels," *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (3), 589–625.

FAO, "Remote Sensing for Spatio-Temporal Mapping of Smog in Punjab and Identification of the Underlying Causes Using GIS Techniques (R-SMOG)," Technical Report, Food and Agriculture Organization of the United Nations 2018.

Gaines, Linda and Eric Rask, "Which Is Greener: Idle, or Stop and Restart?," Argonne National Laboratory and Alternative Fuels Data Center fact sheet 2012. Accessed 28 June 2026.

Gilbert, Sam J., "Strategic Offloading of Delayed Intentions into the External Environment," *The Quarterly Journal of Experimental Psychology*, 2015, *68* (5), 971–992.

GPS.gov, "GPS," https://www.gps.gov/gps. U.S. Government; accessed 19 July 2026.

Greenstone, Michael, Tanushree Ganguly, Christa Hasenkopf, Nishka Sharma, and Hrishikesh Chandra Gautam, "Air Quality Life Index 2025 Annual Update," Technical Report, Energy Policy Institute at the University of Chicago 2025. Air Quality Life Index.

Heutel, Garth, "Optimal Policy Instruments for Externality-Producing Durable Goods under Present Bias," *Journal of Environmental Economics and Management*, 2015, *72*, 54–70.

Ilyas, Hassan and Hira Nissar, "Sectoral Emission Inventory of Lahore," Technical Report, The Urban Unit, Government of Punjab 2023.

Jacobsen, Mark R., James M. Sallee, Joseph S. Shapiro, and Arthur A. van Benthem, "Regulating Untaxable Externalities: Are Vehicle Air Pollution Standards Effective and Efficient?," *The Quarterly Journal of Economics*, 2023, *138* (3), 1907–1976.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman, "Getting to the Top of Mind: How Reminders Increase Saving," *Management Science*, 2016, *62* (12), 3393–3411.

Khan, Aamir Shafaat, "Car Loans Hit Record High of Rs308bn in FY21," *Dawn*, July 2021. Accessed July 16, 2026.

Knittel, Christopher R. and Shinsuke Tanaka, "Fuel Economy and the Price of Gasoline: Evidence from Fueling-Level Micro Data," *Journal of Public Economics*, 2021, *202*, 104496.

Koch, Alexander K., Dan Mønster, and Julia Nafziger, "Spillover Effects of Reminder Nudges in Complex Environments," *Proceedings of the National Academy of Sciences*, 2024, *121* (17), e2322549121.

Kreindler, Gabriel, "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications," *Econometrica*, 2024, *92* (4), 1233–1268.

Langer, Ashley and Shaun McRae, "Step on It: Approaches to Improving Existing Vehicles' Fuel Economy," 2014. Working paper.

Newell, Richard G. and Juha V. Siikamäki, "Nudging Energy Efficiency Behavior: The Role of Information Labels," *Journal of the Association of Environmental and Resource Economists*, 2014, *1* (4), 555–598.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

Pakistan Bureau of Statistics, "Labour Force Survey 2024–25," Technical Report, Pakistan Bureau of Statistics, Government of Pakistan, Islamabad September 2025.

PBS, "Seventh Population and Housing Census 2023," https://www.pbs.gov.pk/content/final-results-census-2023 2023. Pakistan Bureau of Statistics. Accessed 29 June 2026.

Pigou, Arthur Cecil, *The Economics of Welfare*, London: Macmillan and Co., 1920.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Müller, Richard J. Plevin, Adrian E. Raftery, Hana Ševčíková, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff, "Comprehensive Evidence Implies a Higher Social Cost of CO2," *Nature*, 2022, *610*, 687–692.

Rodemeier, Matthias and Andreas Löschel, "Information Nudges, Subsidies, and Crowding Out of Attention: Field Evidence from Energy Efficiency Investments," *Journal of the European Economic Association*, 2025, *23* (4), 1397–1430.

TPL Trakker, "Advanced Bank Safety System," https://tpltrakker.com/industry/banking-2/ 2024. Accessed July 16, 2026.

Tsvetanov, Tsvetan and Kathleen Segerson, "Re-evaluating the Role of Energy Efficiency Standards: A Behavioral Economics Approach," *Journal of Environmental Economics and Management*, 2013, *66* (2), 347–363.

UN, "World Population Prospects 2024," https://population.un.org/wpp/ 2024. United Nations, Department of Economic and Social Affairs, Population Division. Accessed 29 June 2026.

U.S. DOE, "Light-Duty Vehicle Idle Reduction Strategies," https://afdc.energy.gov/conserve/idle-reduction-light 2026. Alternative Fuels Data Center. Accessed 29 June 2026.

U.S. EPA, "What You Can Do to Reduce Pollution from Vehicles and Engines," https://www.epa.gov/transportation-air-pollution-and-climate-change/what-you-can-do-reduce-pollution-vehicles-and 2025. Accessed 29 June 2026.