Working paper PAK-24109 June 2026

Time preferences and dishonesty in the public sector

Muhammad Yasir Khan
Priyoma Mustafi



DIRECTED BY LSE UNIVERSITY OF OXFORD

FUNDED BY UKaid from the British people

Time Preferences and Dishonesty in the Public Sector

Muhammad Yasir Khan* Priyoma Mustafi†

June 26, 2026

Abstract

Public sector workers who self-report their output frequently misreport when monitoring is imperfect, and the share of fabricated reports tends to rise toward the reporting deadline. We study whether this pattern is driven by workers' time preferences—specifically, present bias—using a field experiment with community health workers in Pakistan. We first measure individual present-bias parameters structurally using a Convex Time Budget task adapted to the field, and use these estimates to assign present-biased workers to time-structured financial incentive treatments. In the status-quo, we find that the probability of a false report is 35.2% in the control group, and that false reports concentrate heavily in the second half of the six-day work week—consistent with procrastination-driven backlog accumulation. In the experiment, present-biased workers are randomized to a front-loaded incentive, a uniform incentive, or a control group. Front-loaded incentives significantly increase claimed visits by two-thirds of a visit for present-biased workers, while a uniform incentive of equivalent expected value does not. There is suggestive evidence of misreporting decreasing: the likelihood a claimed visit is false goes down by 3 and 5 percentage points for both treatment groups, with the point estimates significant for only the uniform treatment. These findings indicate that incentives may reduce misreporting but timing of incentives likely does not matter.

JEL Codes: D9, J45, M52, O12
Keywords: present bias, time preferences, dishonesty, public sector, community health workers, incentives

*University of Pittsburgh. Email: yasir.khan@pitt.edu. This research was funded by the International Growth Centre. IRB approval was obtained at the University of Pittsburgh.

†Ahmedabad University. Email: priyoma.mustafi@ahduni.edu.in.

1 Introduction

In many developing countries the basic services of the state are delivered by frontline workers who observe their own output and self-report it to supervisors with neither the time nor the means to verify it (Chaudhury et al. 2006; Banerjee et al. 2008; Callen et al. 2023). When monitoring is imperfect, workers misreport: administrative records routinely overstate the activity that auditors can confirm, distorting both the services citizens receive and the data on Additionally misreporting is not constant over the reporting cycle. The share of reports that are false is low at the start of a work period and climbs steeply as the reporting deadline nears—a rising lie–time profile that the standard model of opportunistic cheating, in which a worker simply weighs the instantaneous gain from fabrication against the risk of detection, does not naturally explain (Nagin et al. 2002). Why does dishonesty escalate toward the deadline, and what does that escalation reveal about how workers make decisions over time? We ask whether the rising profile of dishonesty is driven, at least in part, by workers' present bias.

Whether present bias generates dishonesty on the job is, a priori, unclear, for three reasons. First, a rising lie–time profile is consistent with more than one mechanism. A present-biased worker with quasi-hyperbolic preferences (Laibson 1997; O'Donoghue and Rabin 1999) overweights the immediate cost of genuine effort, procrastinates, and lets a backlog of unfinished work accumulate; because real effort is increasingly costly to cram into the shrinking window before the deadline, she meets her quota by fabricating an ever-larger share of her reports as the deadline approaches. But the same rising profile can arise without any intertemporal distortion if the moral cost of fabrication falls as the deadline nears—lying under time duress feels excused—so that even a perfectly time-consistent worker fabricates more late in the cycle. The two forces—present bias and this "deadline pressure"—produce observationally similar trends, so documenting the trend only in misreporting is not enough; identifying present bias driven lying also requires the claimed work done over time. Second, although a substantial literature shows that present-biased workers procrastinate costly real-effort tasks (Augenblick et al. 2015; Kaur et al. 2015; DellaVigna and Malmendier 2006), none has traced the link from intertemporal preferences to dishonest self-reporting in a natural workplace where workers face repeated quotas under weak monitoring.[1] Third, even if present bias

[1]Laboratory work on dishonesty documents that people lie less than selfish payoff maximization predicts, attributing the gap to guilt, lying aversion, and a preference for being seen as honest (Gneezy 2005; Fischbacher and Föllmi-Heusi 2013; Abeler et al. 2019); these forces are time-invariant and do not predict a profile that rises toward a deadline. Field evidence on misreporting in public service has emphasized rational deterrence (Fisman and Miguel 2007; Olken 2007) and selection into the workforce (Hanna and Wang 2017), rather than the dynamics of intertemporal choice.

drives misreporting, it is not obvious how the timing of financial incentives—as opposed to their level—affects effort and honesty among workers who differ in their degree of present bias, which is precisely the margin a behaviorally targeted policy would have to exploit.

We organize these questions around a simple model of a worker who must file a fixed quota of reports by a deadline and who, each period, decides for each report still owed whether to *do* it, *postpone* it, or *lie*. Genuine effort is convex:real visits are progressively harder to cram into a single day so a worker who postpones accumulates a backlog that becomes prohibitively costly to clear honestly, and substitutes toward fabrication at the margin. The central comparative static suggests the falsification rate rises as the deadline approaches, and its slope is steeper the more present-biased the worker, above a common component that any deadline-pressure force would also produce. The model further signs the effect of incentives. Raising the relative price of early effort pulls genuine work forward and shrinks the backlog at its source, whereas a uniform per-report rate of the same expected value does not tilt the relative returns to early versus late effort, but raises the marginal return to honest completion where fabrication is most tempting—near the deadline.

To test them, we partner with the Department of Health in Khyber Pakhtunkhwa province of Pakistan to conduct a multi-phase randomized field experiment with Lady Health Workers (LHWs), the all-female cadre of community health workers who deliver preventive care door-to-door. The study has four phases. In Phase 1 we administer a baseline survey and record productivity benchmarks from workers' paper visit registers. In Phase 2 we elicit each worker's present-bias parameter structurally, using a Convex Time Budget (CTB) task (Andreoni and Sprenger 2012; Augenblick et al. 2015) adapted to the field: workers allocate a real-effort task,sending health-data SMS messages, across two adjacent workdays at experimentally varied intertemporal prices, from which we recover an individual present-bias parameter $\hat{\beta}$ following Andreoni et al. (2022). In Phase 3 we use these estimates to focus the experiment on *present-biased* workers ($\hat{\beta} < 1$), for whom the mechanism is active, and randomize them across three arms of equal expected payoff (Rs. 500, approximately $1.80): a *front-loaded* incentive that pays a higher per-visit rate for visits reported in the first three days of the work week than in the last three; a *uniform* incentive that pays a constant per-visit rate throughout; and a flat-bonus control.[2] In Phase 4, independent audit teams blind to treatment physically verify a randomly selected subset of each worker's claimed visits.

The LHW program offers organizational features that make this test possible. Each worker

[2]Non-present-biased workers ($\hat{\beta} \geq 1$) received a separate, non-pecuniary intervention that we analyze elsewhere. Field piloting indicated that the time-varying incentive induced gaming rather than additional effort among these workers, which is why we confined it to the present-biased stratum.

serves a single, fixed catchment area and, in the reporting campaign we study, must visit and record health data for a minimum of 36 households within a six-day work week—a well-defined quota and deadline, which are the essential ingredients of the mechanism. Workers submit paper visit reports to a Lady Health Supervisor who does not independently confirm them, so that a measurable quota meets imperfect monitoring. Workers operate independently in non-overlapping communities, limiting spillovers across treatment arms, and the weekly cycle is repeated, which lets us observe the within-week dynamics of effort and falsification that are the empirical basis for separating present bias from deadline pressure.

We measure behavior from two sources. Worker-reported visits come from the paper forms each LHW completes and submits to her supervisor, which our field teams then collect. For the audit, we randomly select one reported household per day—six per worker over the work week—and dispatch independent enumerators, with no knowledge of the worker's treatment assignment, to confirm in person whether the visit occurred, yielding an objective day-by-day measure of falsification. The experimental analysis focuses on the 176 present-biased workers identified by the structural elicitation.

We report three sets of findings. First, the rising profile of dishonesty is pronounced, and it is steeper for more present-biased workers. In the control group the probability that a claimed visit is false averages 0.33, but it is heavily back-loaded—0.16 in the first half of the week against 0.55 in the second, so that more than half of all late-week reports are fabricated. Regressing the falsification rate on $\hat{\beta}$, the day of the work week, and their interaction within the control group, the interaction is -0.089 ($p < 0.01$): each unit decrease in $\hat{\beta}$ (more present bias) steepens the within-week slope by 0.089 per day. Time-consistent workers also display a rising profile, identifying deadline pressure as a background force common to all workers; but the additional steepness among present-biased workers suggests that procrastination, and not merely calendar proximity to the deadline, drives late-cycle lying.

Second, the front-loaded incentive pulls genuine effort forward without manufacturing it. Present-biased workers in the front-loaded arm claim 0.65 more visits per day than the control group ($p < 0.01$), a 9.8% increase over the control mean of 6.594 daily visits, and the response is concentrated in the early week, exactly where the design intends to draw effort. Crucially, this additional output is not fabricated: the effect of the front-loaded arm on the probability that a report is false is small, negative, and statistically insignificant (-0.03). The uniform arm carries an identical Rs. 500 expected payoff but a flat intertemporal price and it does not move claimed visits to early period. This provides evidence for the possibility of the backlog mechanism—raising the relative price of early effort forestalls the late-week cramming into which a present-biased worker would otherwise default.

Third, the uniform incentive improves honesty: although it does not tilt the relative returns to early versus late effort, it lowers the probability of a false report in the *late* half of the work week by 0.13 ($p < 0.05$)—a roughly 24% reduction relative to the late-week control mean of 0.55—while leaving the early week statistically untouched; the front-loaded arm moves late-week falsification in the same direction (-0.06) but does not attain significance. The two incentives thus operate through different margins of the model: the front-loaded scheme reallocates effort toward the early week and shrinks the backlog at its source, whereas the uniform scheme raises the marginal attractiveness of honest completion at the deadline, where the temptation to fabricate peaks.

This paper makes three contributions. To the best of our knowledge, it provides the first field-experimental evidence that present bias is a mechanism for dishonest misreporting in the public sector. Existing models of workplace opportunism treat agents as rational cheaters who weigh the instantaneous gain from misreporting against expected detection costs, abstracting from intertemporal preferences (Nagin et al. 2002; Holmstrom and Milgrom 1991); laboratory evidence documents systematic departures from purely selfish cheating attributable to guilt, lying aversion, and image concerns (Gneezy 2005; Fischbacher and Föllmi-Heusi 2013; Abeler et al. 2019), but these forces are time-invariant and cannot generate a falsification profile that rises toward a deadline. We show that intertemporal preferences create a dynamic link between the timing of effort and the timing of fabrication, producing a rising lie–time profile whose slope increases with the worker's degree of present bias.

Second, the paper contributes to the literature on self-control and present bias in labor markets, which documents that workers procrastinate costly real-effort tasks (Augenblick et al. 2015), pay for commitment devices that penalize missed interim targets (Kaur et al. 2015), and over-commit when they mispredict their own future effort (DellaVigna and Malmendier 2006). We extend this work in two ways. We show that procrastination carries a consequential externality—it inflates misreporting, imposing costs not only on the worker but on the principal and on the downstream policy decisions that rely on the data—and we develop a field strategy that uses structurally estimated individual $\hat{\beta}$ to tailor the *timing* of incentives. This extends the customization approach of Andreoni et al. (2022), who tailor the level of intertemporal incentives to measured discount factors, to the within-cycle timing structure of performance pay.

Third, the paper speaks to the literature on public sector performance in developing countries, which has sought to raise frontline effort through financial incentives (Muralidharan and Sundararaman 2011; Duflo et al. 2012; Khan et al. 2016), monitoring (Olken 2007),

and selection into service (Dal Bó et al. 2013; Hanna and Wang 2017).[3] Our findings add a dimension that this literature has largely overlooked: the *timing* of incentives, not only their level, shapes both effort and honesty, and the behaviorally optimal timing depends on the worker's present bias. By tracing the well-documented gap between reported and verified service provision (Chaudhury et al. 2006; Banerjee et al. 2008) partly to procrastination, we identify a behavioral source of the overstatement in administrative data; and through the contrast between the two incentive schemes and the time-consistent benchmark, we offer a framework that distinguishes present bias from deadline pressure—two forces that generate the same rising profile but call for different remedies, since backlog-reducing interventions (front-loaded pay, interim milestones) address the first while honesty reminders or heavier late-period auditing address the second.

The rest of the paper proceeds as follows. Section 2 develops the theoretical framework. Section 7 describes the study context. Section 8 outlines the experimental design and the elicitation of present bias. Section 9 presents the main results, and Section 10 concludes.

2 Theoretical Framework

We develop a model of workplace misreporting that embeds time-preferences in decisions to delay work generating a rising falsification profile, and characterizes how such lying responds to incentives. The model motivates the experimental design and the interpretation of results.

A worker must file a fixed quota of activity reports (household visits, site inspections, client calls) by a deadline. Each report can be earned by genuine effort or fabricated, and the verifier cannot tell them apart at the time of filing. Every period the worker therefore chooses, for each report still owed, whether to *do* it, *postpone* it, or *lie*. The model below shows how that choice depends on how many reports are left, how many periods remain, the cost of effort, and the cost of lying, and how *present bias* bends the whole profile. The base case delivers the central comparative static—lying rises as the deadline nears, and more steeply for present-biased workers—and the final section uses it to sign the effect of two pay schemes: a uniform piece rate and a high upfront (front-loaded) piece rate.

[3] See Finan et al. (2017) for a synthesis of the personnel economics of the developing-country state, and Callen et al. (2020) on how the quality of self-reported administrative data shapes policy decisions.

3 Environment

Time is discrete, $t \in \{1, \ldots, T\}$, with a deadline at T. The worker owes a quota of N reports by the end of period T. Let S_t be the *backlog* of reports still owed at the start of period t, so

$$S_1 = N, \qquad S_{t+1} = S_t - q_t, \qquad \sum_{t=1}^{T} q_t = N,$$

where $q_t \geq 0$ is the number of reports filed in period t. Each filed report is either *genuine* (n_t of them) or *fabricated* (ℓ_t of them), with

$$q_t = n_t + \ell_t, \qquad n_t, \ell_t \geq 0.$$

Genuine effort and fabrication carry the within-period costs

$$C(n_t) = \frac{c}{3} n_t^3 \quad (c > 0), \qquad \Lambda(\ell_t) = \frac{\psi}{2} \ell_t^2 \quad (\psi > 0). \tag{1}$$

The cubic effort cost has marginal cost $c\, n_t^2$ that climbs faster than fabrication's marginal cost $\psi\, \ell_t$ as quantities grow: real visits are progressively harder to cram into one period (fixed hours, travel, fatigue), so beyond a point fabrication is the cheaper margin. This convexity gap is a *soft capacity* on genuine effort—it bites smoothly rather than at a hard ceiling, and it is what turns a procrastinated backlog into lies.[4]

Present bias. Preferences are quasi-hyperbolic (β–δ) with $\delta = 1$ and present-bias factor $\beta \in (0, 1)$. The period-τ self weights her own current cost at 1 and every future cost at β, minimizing

$$g_\tau(q_\tau) \; + \; \beta \sum_{t > \tau} g_t(q_t), \tag{2}$$

where $g_t(\cdot)$ is the period flow cost defined in (5) below. A worker with $\beta \geq 1$ is *time-consistent*; $\beta < 1$ is *present-biased*. Setting $\delta = 1$ makes present bias the only intertemporal distortion. Workers are *naive*: each self believes her future selves will not be present-biased.

[4] The lying-cost slope ψ is held constant here to isolate present bias. Allowing it to *fall* as the deadline nears—$\psi_t = \psi_0 + \theta(T - t)$, "lying under time duress feels excused"—adds a second force that also tilts lying toward the deadline; it reinforces every result below.

4 The within-period decision: do, postpone, or lie

Take as given the number of reports q a worker files in some period. She splits them to minimize that period's cost,

$$\min_{n+\ell=q,\ n\geq 0} \ \frac{c}{3}\, n^3 + \frac{\psi}{2}\, \ell^2.$$

Lemma 1 (Within-period split and the lie rate)**.** *The cost-minimizing split equates marginal costs, $c\, n^2 = \psi\, \ell$, giving genuine effort*

$$n(q) = \frac{-\psi + \sqrt{\psi^2 + 4c\,\psi\, q}}{2c}, \qquad \ell(q) = q - n(q), \tag{3}$$

and a falsification (lie) rate

$$\rho(q) \ \equiv \ \frac{\ell(q)}{q} \ = \ 1 - \frac{n(q)}{q}. \tag{4}$$

The reduced flow cost is

$$g(q) = \frac{c}{3}\, n(q)^3 + \frac{\psi}{2}\, \ell(q)^2, \qquad g'(q) = c\, n(q)^2 = \psi\, \ell(q) > 0, \quad g''(q) > 0. \tag{5}$$

The lie rate $\rho(q)$ is strictly increasing in q, with $\rho \to 0$ as $q \to 0$ and $\rho \to 1$ as $q \to \infty$.

Proof. The objective is strictly convex in n on $n + \ell = q$, so the first-order condition $c\, n^2 = \psi(q - n)$ characterizes the optimum; solving the quadratic gives (3). Differentiating $c\, n^2 = \psi(q - n)$ yields $n'(q) = \psi/(2cn + \psi) \in (0, 1)$, so $\ell'(q) = 1 - n'(q) \in (0, 1)$ and $\ell(q)/q$ is increasing because ℓ is convex through the origin ($\ell(0) = 0$, ℓ' increasing). The limits follow from $n(q) \sim \sqrt{\psi q/c} \to \infty$ but $n(q)/q \to 0$. The envelope identity $g'(q) = c\, n(q)^2 = \psi\, \ell(q)$ is immediate from the binding constraint. □

Lemma 1helps isolate the three actions in the worker's choice set. *Doing* a report is a genuine visit n; *lying* is a fabrication ℓ; *postponing* is choosing a low q today and carrying the backlog $S_{t+1} = S_t - q_t$ forward. The lie rate depends only on the filing intensity q: the more reports a worker crams into a single period, the more of them she fabricates at the margin, because genuine effort has become expensive while fabrication has not. Present bias never enters the split directly; it acts on lying *only* by changing *when* reports are filed—i.e. by inflating q late in the cycle.

5 Base case: present bias, procrastination, and the lie–time profile

The worker chooses the filing path $\{q_t\}$ to minimize (2) subject to $\sum_t q_t = N$. We compare the time-consistent benchmark with the present-biased worker.

Time-consistent benchmark ($\beta = 1$). The first-order conditions are $g'(q_t) = \mu$ for a common multiplier μ, so a patient worker *smooths* filing: q_t and hence the lie rate $\rho(q_t) = R(q_t)$ are *flat* across the cycle. With a constant moral cost she has no reason to bunch, so she neither procrastinates nor escalates her lying. This is the no-present-bias reference profile.

Present-biased worker ($\beta < 1$). Each self over-weights her own immediate effort cost relative to the future, so she defers genuine work—"I will catch up later." Filing is back-loaded and the backlog mounts; the late selves then face heavy filing q_t and, by Lemma 1, fabricate an increasing share of it.

Proposition 1 (Present bias generates a rising lie–time profile)**.** *Let $\beta < 1$ and ψ constant. The naive present-biased worker postpones filing, so her filing intensity is increasing over the cycle, $q_1 < q_2 < \cdots < q_T$, and therefore her lie rate rises monotonically toward the deadline,*

$$\rho_t = R(q_t) \text{ is increasing in } t.$$

The escalation is steeper the lower is β: more impatient workers carry a larger late-cycle backlog and fabricate a larger share of it.

Proof Sketch. A naive period-τ self chooses q_τ believing every later self will file the patient (flat) continuation of the remaining backlog. Her first-order condition is $g'(q_\tau) = \beta\,\nu_\tau$, where ν_τ is the (believed, undistorted) marginal continuation cost of the backlog she carries. Because she discounts that continuation by $\beta < 1$ while bearing today's marginal cost undiscounted, she sets $g'(q_\tau) < \nu_\tau$ and so files *less* than the amount that would keep the path flat; the shortfall is rolled into $S_{\tau+1}$. Iterating, the realized backlog and realized filing both rise over t, and the last self clears whatever remains. Since $R(\cdot)$ is strictly increasing (Lemma 1), ρ_t rises with q_t. Lowering β deepens the per-period shortfall, raising q_t in the late periods and steepening the profile. □

The lie–time relationship and its dependence on present bias. Proposition 1 is yields two objects of interest.

- *Across periods (fixed worker).* The slope of lying in calendar time, $\rho_{t+1} - \rho_t$, is *positive* for a present-biased worker and *zero* for a time-consistent one. Lying is a late-cycle, near-deadline phenomenon, and it is a phenomenon of *procrastination*, not of any taste for dishonesty that varies with the calendar.
- *Across workers (the cross-partial).* The steepness of the profile is governed by β:

$$\frac{\partial}{\partial(1-\beta)}\big(\rho_T - \rho_1\big) > 0.$$

  The lie–time slope fans out as workers become more present-biased. Equivalently, the gap between a worker's profile and the flat time-consistent benchmark, $\Delta_t(\beta) = R(q_t(\beta)) - R(N/T)$, is positive, increasing in t, and increasing in $(1-\beta)$. This is the model's sharpest empirical content: the *interaction* of the deadline trend with elicited present bias is the present-bias fingerprint.

Comparative statics of the base case. From Lemma 1 and Proposition 1, total fabrication $L = \sum_t \ell_t$ moves intuitively in the primitives:

$$\frac{\partial L}{\partial(1-\beta)} > 0, \qquad \frac{\partial L}{\partial \psi} < 0, \qquad \frac{\partial L}{\partial c} > 0, \qquad \frac{\partial L}{\partial N} > 0.$$

More present bias, a cheaper effort technology (higher c), a softer lying penalty (lower ψ), or a heavier quota all raise lying.

6 Comparative statics: uniform vs. high upfront piece rates

We now attach money to reports and ask how the *timing* of pay changes lying, and how that depends on present bias. With T even, split the cycle into halves $H_1 = \{1, \ldots, T/2\}$ (the "first half") and $H_2 = \{T/2+1, \ldots, T\}$. Three schemes:

(Q) status quo: $w_t \equiv 0$; (U) uniform: $w_t \equiv x$; (F) front-loaded: $w_t = \begin{cases} 2x, & t \in H_1 \\ \frac{x}{2}, & t \in H_2. \end{cases}$

Scheme F pays a high upfront rate for first-half work and a low rate for second-half work; it is calibrated to cost the principal the same as U at the status-quo path. We maintain three assumptions.

Assumption 1 (Pay per filed report, paid on filing)**.** *Each report filed in period t pays w_t, received* in period t*. Because genuine and fabricated reports are indistinguishable at filing, both pay. The period-τ self minimizes $g_\tau(q_\tau) - w_\tau q_\tau + \beta \sum_{t>\tau} \big[g_t(q_t) - w_t q_t\big]$.*

Assumption 2 (Fixed task list)**.** *The quota N is a floor and a ceiling: the worker is paid for at most N reports and must file exactly N. There is no volume margin.*

Assumption 3 (Linear payments)**.** *The worker is risk-neutral in money; payments enter linearly.*

Two features of Assumption 1 do the work. Per filed report means money cannot buy honesty *within* a period—a rupee is earned whether the report is real or fake—so the split of Lemma 1 is untouched and pay can act only by relocating filings across periods. Paid on filing means the present-biased self values today's rate at 1 and a deferred rate at β, this assumption can be normalized.

Lemma 2 (Pay acts only on the filing path)**.** *Under Assumptions 1–3, for any path $\{q_t\}$ the within-period split is exactly that of Lemma 1 and the lie rate is $\rho_t = R(q_t)$. Piece-rate schemes change lying* only *through $\{q_t\}$.*

So a scheme reduces lying if and only if it moves filings out of the crammed late periods (where the marginal report is largely fabricated) into the lighter early periods (where it is largely genuine). How far it moves them, and how much each move helps, are the two terms below. We use the two-period version ($T = 2$, $H_1 = \{1\}$, $H_2 = \{2\}$) for closed forms.

The wedge: how far filings move. The status-quo path solves $g'(q_1^0) = \beta\, g'(q_2^0)$ with $q_2^0 = N - q_1^0$; since $\beta < 1$ this gives $q_1^0 < q_2^0$ (back-loading, as in Proposition 1). Under scheme s the first-period condition tilts to

$$g'(q_1^s) = \beta\, g'(N - q_1^s) + W_s\, x, \qquad W_U = 1 - \beta, \qquad W_F = 2 - \frac{\beta}{2}. \tag{6}$$

The *wedge* W_s measures how hard the scheme pushes filings into the first period. It decomposes as

$$W_F = \underbrace{(1-\beta)}_{\text{bias-harnessing } (=W_U)} + \underbrace{\left(1 + \tfrac{\beta}{2}\right)}_{\text{calendar spread}}.$$

The first term, common to U and F, is present bias turned against itself: paying x now rather than later is worth $x(1-\beta)$ to a worker who discounts the future extra, so a *uniform* rate already pulls work forward. The second term is the pure front-loading spread ($2x$ vs. $x/2$), which moves filings even for a time-consistent worker.

The gain: how much each moved filing helps. Reallocation maps into lying through the *marginal lie share*

$$\sigma(q) \;\equiv\; \frac{\partial \ell}{\partial q} \;=\; \frac{2c\,n(q)}{2c\,n(q)+\psi} \in (0,1),$$

the fraction of a marginal filed report that is fabricated; σ is increasing in q (crammed periods lie at the margin). Define the *reallocation gain* and a curvature term

$$G \;\equiv\; \sigma(q_2^0)-\sigma(q_1^0) \;\geq\; 0, \qquad D \;\equiv\; g''(q_1^0)+\beta\, g''(q_2^0) \;>\; 0.$$

G is the drop in total lying when one marginal report is moved from the late, crammed period to the early, light one; it is strictly positive precisely when the status-quo path is back-loaded, i.e. precisely when $\beta < 1$.

Proposition 2 (How much lying falls under each scheme)**.** *In the two-period model, the reduction in total fabrication relative to the status quo is, to first order in x,*

$$\Delta L_s \;\equiv\; L^Q - L^s \;=\; \frac{W_s(\beta)\,G(\beta)}{D}\,x \;+\; O(x^2), \qquad W_U = 1-\beta, \;\; W_F = 2-\tfrac{\beta}{2}.$$

Hence:

(i) ***(Uniform.)*** $\Delta L_U \;=\; \dfrac{(1-\beta)\,G}{D}\,x + O(x^2)$. *A uniform rate reduces lying* only *for present-biased workers, the more so the lower is β, and has* no *first-order effect at $\beta = 1$ (for a time-consistent worker with a binding quota the uniform rate is a lump sum and changes nothing).*

(ii) ***(Front-loaded.)*** $\Delta L_F = \dfrac{(2-\beta/2)\,G}{D}\,x + O(x^2)$. *The high upfront rate reduces lying for* every *worker with a back-loaded path, and strictly more than the uniform rate, $\Delta L_F > \Delta L_U$, whenever $G > 0$.*

(iii) ***(Ratio.)*** $\dfrac{\Delta L_F}{\Delta L_U} \;\to\; \dfrac{2-\beta/2}{1-\beta}$ *as $x \to 0$: about 5.5 at $\beta = 0.7$, 8 at $\beta = 0.8$, 15.5 at $\beta = 0.9$, and ∞ at $\beta = 1$. Front-loading dominates uniform most sharply for the* least *present-biased workers, because for them the uniform rate has almost nothing to harness.*

(iv) ***(Both rise with present bias.)*** *G is increasing in $(1-\beta)$, so present bias raises the lying reduction* twice*: through the wedge (payment-on-filing harnesses the bias) and through the gain (procrastination widens $q_2^0 - q_1^0$, enlarging G). The* level *of lying averted, ΔL_F, is largest for the most present-biased workers.*

Sketch. By Lemma 2, $L = \ell(q_1) + \ell(q_2)$ with $q_2 = N - q_1$, so $dL/dq_1 = \ell'(q_1) - \ell'(q_2) = \sigma(q_1^0) - \sigma(q_2^0) = -G$ at the status quo. Totally differentiating the equilibrium condition

(6) at $x = 0$ gives $dq_1^s/dx = W_s/D$ with $D = g''(q_1^0) + \beta g''(q_2^0)$. Chain rule: $\Delta L_s = -\,dL/dq_1 \cdot dq_1^s/dx \cdot x + O(x^2) = W_s G/D\, x + O(x^2)$. The signs of G and W_s in β, and the ratio, follow by inspection; $q_2^0 - q_1^0$ is increasing in $(1-\beta)$ from $g'(q_1^0) = \beta g'(q_2^0)$, and σ is increasing, so G is increasing in $(1-\beta)$. □

7 Context

7.1 Lady Health Workers

Community health workers play a key role in delivering preventive and basic health care in many countries around the world. According to the World Health Organization's Global Health Observatory, an estimated 4.5 million community health workers across 97 countries operate within the global healthcare system. Their role has received special attention since the 1970s, as low- and middle-income countries face a shortage of trained health professionals to promote preventive health care aimed at achieving sustainable development goals (Chaudhury et al. 2006).

In Pakistan, community health workers are considered the backbone of the preventive and primary healthcare system, especially in rural areas. These workers, all of whom are women,[5] function as a separate division of the Department of Health under the Lady Health Workers (LHW) program. Established as a special program in 1993, the LHW program now employs approximately 96,000 workers across the country. Since 2014, LHWs are considered full-time public sector employees, with job protections equivalent to those afforded to other members of the state bureaucracy.

Khyber Pakhtunkhwa province employs approximately 13,500 LHWs, making it one of the larger provincial programs. Each worker is affiliated with a Lady Health Supervisor (LHS) for administrative and reporting purposes, but is hired by the Department of Health to live and work in a clearly defined catchment area. Each catchment area is served by only one worker, thereby minimizing interaction between workers in their routine jobs. This feature limits spillovers across treatment arms and enables clean measurement of individual performance. A feature of LHWs in KP is that, administratively, they are hired either by the department directly or are retained as part of the Integrated Health Project.[6]

[5] According to the World Health Organization, women comprise 70% of workers in the health sector in 104 countries.

[6] This distinction does not affect their day to day job requirements, described below. The key difference between the two is how the LHWs are paid.

7.1.1 Duties of Lady Health Workers

Lady Health Workers are required to visit all the households in their catchment area on a regular basis, with the core duty of providing preventive and basic health care to residents at their doorstep. During these visits, workers record household-level health information, advise women on birth control and family planning, conduct antenatal checks for expectant mothers, and provide guidance on disease prevention and nutrition. LHWs also track the vaccination status of children in the household and refer cases requiring clinical attention to the nearest health facility.

In this study, we focus on a specific reporting cycle: as part of a special campaign, the workers are required to visit and record health data for a minimum of 36 households with pregnant women and children within a six-day work week (Monday through Saturday). Workers fill out standardized paper-based forms recording the date and household of each claimed visit and submit these forms to their LHS at the end of the week.

7.1.2 Incentives, Monitoring, and Reporting

Lady Health Workers receive a fixed monthly salary from the Department of Health for their routine duties. Workers are hired to serve specific catchment areas and are rarely transferred. During this study, we recorded no instances of workers being transferred or leaving the study area.

Routine monitoring of LHWs is limited. While workers are expected to maintain records of their household visits, supervisors rarely cross-check these paper registers against independent verification of actual visits. The primary accountability mechanism is when the LHS decides to visit the catchment area of a worker and checks the paper-based record of households visited by the worker. This weak monitoring environment creates the conditions for moral hazard: workers can misreport their visits with low risk of detection in the absence of an independent audit mechanism.[7]

7.2 Khyber Pakhtunkhwa Province

Khyber Pakhtunkhwa is a province of approximately 35 million people in Pakistan's northwest. The province has a young population structure and faces significant challenges in key health indicators: the neonatal mortality rate stands at 42 deaths per 1,000 live births

[7]This is consistent with the broader pattern documented in the developing country public health literature, where reported and verified service provision diverge substantially (Chaudhury et al. 2006; Banerjee et al. 2008; Callen et al. 2023).

according to the Pakistan Demographic and Health Survey 2017–18. Despite an extensive network of community health workers and primary care facilities, preventive health services remain under-utilized, partly because of low worker productivity in the LHW program.

8 Experimental Design and Implementation Details

In this section, we outline the methodology of our experiment conducted between July 2024 and June 2026. The study proceeds in four phases: a baseline survey and productivity benchmark, time preference elicitation using the Convex Time Budget method, a randomized incentive experiment, and an independent audit of submitted reports.

8.1 Phase 1: Baseline Survey and Productivity Benchmarks

We began the study by administering a baseline survey to 734 LHWs. The survey collected information on workers' demographics, household composition, tenure, public service motivation, and behavioral preferences including risk preferences. In addition to the survey, we recorded initial productivity benchmarks for workers by extracting visit counts from their existing paper-based registers. These baseline productivity measures provide a reference for workers' typical effort levels prior to any experimental intervention and serve as a stratification variable for randomization.[8]

8.2 Phase 2: Time Preference Estimation via Convex Time Budget

To measure each worker's intertemporal discount factor and present-bias parameter, we conducted a Convex Time Budget (CTB) experiment (Andreoni and Sprenger 2012; Augenblick et al. 2015). The CTB method elicits time preferences by asking participants to allocate effort between two dates at experimentally varied intertemporal prices, recovering structural estimates of discounting parameters under assumptions about the cost of effort.

Real-effort task. To make the CTB ecologically valid for this population, we adapted the standard monetary CTB to a real-effort context following Augenblick et al. (2015). Workers were asked to allocate effort in the form of sending health information SMS messages—a task directly related to their routine duties—across two adjacent workdays, Day 1 (immediate)

[8]The final sample from baseline is 477 which reflects workers who (i) actively used the paper-based reporting system during the baseline productivity phase and (ii) were available and reachable by phone at the time of treatment delivery. We verify that selection into the final sample is uncorrelated with treatment assignment.

and Day 2 (one day later). Each LHW decided how many SMSs to send on each day, subject to an intertemporal budget constraint.

Budget sets and decisions. Workers faced allocation decisions of the form:

$$e_t + R \cdot e_{t+k} = m,$$

where e_t is the number of tasks allocated to the earlier date, e_{t+k} the number allocated to the later date, R is the experimentally varied task rate, and m is the budget. In our implementation, workers faced two budgets drawn uniformly from $m \in \{20, 30, 40\}$ and four task rates drawn uniformly from $R \in \{0.5, 0.75, 1, 1.25, 1.5\}$, generating a total of sixteen allocation decisions per worker. Each worker made eight decisions on Day 0 (the day of the baseline survey), allocating work between Day 1 and Day 2 at various prices, and eight decisions on Day 1 itself, again allocating between Day 1 and Day 2.[9]

Identification of time preferences. Comparing allocations made in advance (Day 0) with allocations made on the day the earlier date becomes immediate (Day 1) allows us to test for dynamic inconsistency. We recover the distribution of individual discount factors δ_i and present-bias parameters β_i structurally as follows:

Cost of effort. We assume that workers face increasing marginal costs of effort. Effort costs are represented by a power function, with curvature governed by the parameter γ. We estimate β_i for four values of the curvature parameter $\gamma \in \{1.5, 1.75, 2, 2.25\}$, and use the median estimate across specifications.

$$c(e) = e^{\gamma},$$

Quasi-hyperbolic discounting. We assume individuals engage in quasi-hyperbolic discounting (Laibson 1997; O'Donoghue and Rabin 1999) with δ as the standard discount factor and β as the present-bias parameter. For choices involving an immediate date (i.e. choices made on Day 1), future utility is discounted by $\beta\delta$. For choices involving only advance dates (i.e. choices made on Day 0), utility is discounted by δ alone.

For a given worker i, making allocation $(e_{1,i}, e_{2,i})$, the discounted disutility of effort can be written as

[9] We fix the delay between the two dates at $k = 1$ day, consistent with the within-week temporal horizon of the main experiment.

$$e_{1,i}^{\gamma} + \beta_i^{\mathbf{1}_{immediate=1}} \delta_i \cdot e_{2,i}^{\gamma}. \tag{7}$$

where

$$\mathbf{1}_{immediate=1} = \begin{cases} 1 & \text{allocation decision on Day 1 (Immediate),} \\ 0 & \text{allocation decision on Day 0 (Advance).} \end{cases}$$

The parameter β_i captures the degree of present bias for workers. If $\beta_i = 1$, the worker exhibit exponential discounting with discount factor δ_i, while if $\beta_i < 1$, the worker exhibits present bias, being less patient in immediate relative to advance decisions.

Completion probability. Following Andreoni et al. (2022), workers anticipate the possibility of failing to complete assigned work. Completion probabilities are assumed to take the form

$$p(e) = \frac{1}{1 + \alpha e},$$

where α is a completion parameter capturing how rapidly the probability of successful completion declines as workload for a day (e) increases. Larger values of α imply that workers perceive a greater risk of failing to complete larger task allocations.

Workers receive a completion bonus V conditional on successful completion. Consequently, when making allocation decisions they trade off expected completion benefits against discounted effort costs.

Structural estimation. The parameters α, β, and δ are estimated jointly using two-step GMM with standard errors clustered at the worker level.

Workers minimize the discounted costs of effort subject to the intertemporal budget constraint provided by their bonus contract. This yields a marginal condition, which forms the basis of the first GMM moment condition used in estimation.

$$\gamma e_{1,i}^{\gamma-1} - \frac{\beta_i^{\mathbf{1}_{immeadiate=1}} \delta_i}{R} \gamma e_{2,i}^{\gamma-1} = 0 \tag{8}$$

To identify the completion parameter α, we additionally exploit observed task completion outcomes. Assuming independence across the two dates, the probability of successfully

completing both is

$$\frac{1}{1+\alpha e_{1,i}} \cdot \frac{1}{1+\alpha e_{2,i}}, \tag{9}$$

which yields the second GMM moment condition.

Together, these moments identify the completion-risk parameter, α, which governs how completion probabilities decline with assigned workload. Using the estimated value of α, we recover individual discounting measures by rearranging the first-order condition. For advance decisions, made on Day 0, the first-order condition identifies an individual's long-run discount factor, δ_i. For immediate decisions, made on Day 1, the first-order condition identifies the composite discounting parameter, $\beta_i\delta_i$, which reflects both standard discounting and present bias. Thus, the individual present-bias parameter are recovered as

$$\beta_i = \frac{\beta_i\delta_i}{\delta_i}. \tag{10}$$

Estimated Beta. We restrict the baseline sample of 477 to workers for whom we can recover $\hat{\beta}_i$ in the range $[0.5, 2]$, yielding a sample of 471 workers with stable, well-identified estimates. Figure 1 presents the distribution of $\hat{\beta}$ in the final experimental sample. Workers with $\hat{\beta}_i < 1$ are classified as *present-biased*; workers with $\hat{\beta}_i \geq 1$ are classified as *non-present-biased*.

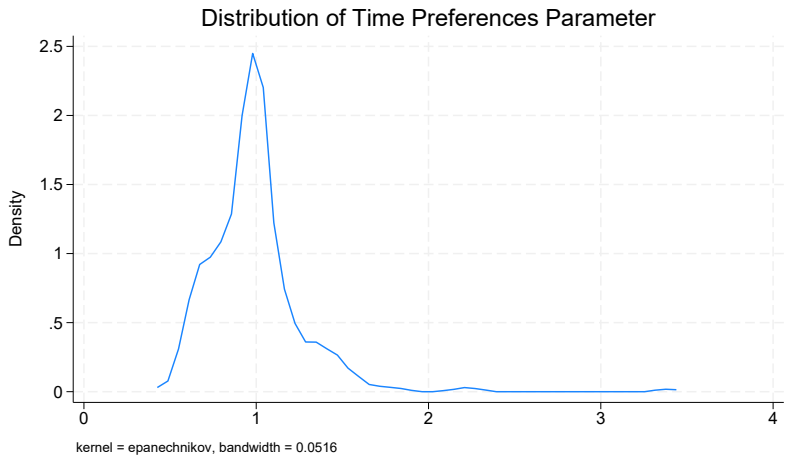

Figure 1: **Distribution of Estimated Present-Bias Parameter $\hat{\beta}$**

Notes: The figure plots the kernel density of the median present-bias estimate $\hat{\beta}$ across 471 workers for whom estimates fall in the range $[0.5, 2]$. Workers to the left of the dashed line at $\hat{\beta} = 1$ are classified as present-biased and constitute the present-biased experimental sample; workers to the right are classified as non-present-biased.

8.3 Phase 3: Pilot Testing

Prior to the main experiment, we conducted a pilot with 50 workers to assess the feasibility of different reporting cycle lengths. We tested both a 30-day and a 6-day work cycle. The 30-day cycle proved operationally challenging: collecting and verifying daily reports over an entire month created logistical difficulties for both workers and field teams, and within-month data were difficult to extract from hand-written registers in a timely manner. We therefore adopted a 6-day reporting week as the main experimental unit. The pilot also identified that non-present-biased workers in the upfront incentive condition showed gaming behavior—shifting work to the incentivized early period without a corresponding reduction in misreporting—which led us to restrict the time-varying incentive to the present-biased subsample. These 50 workers were excluded, thus the final available is 421. Of this sample, we conduct experiment with 176 workers who can be identified as present-biased, based on $\beta < 1$. Our model predicts the incentives will help reduce lying of such workers. We divided these workers into three groups as discussed below.

8.4 Treatments for Present Biased Workers

In each treatment condition, workers were contacted by phone two days before the start of their reporting week. During this phone call, the enumerator delivered the treatment script and explained how to complete the paper report forms. Workers were told that field teams would collect their completed forms from the LHS at the end of the week. In all treatment conditions, the expected payoff to the worker is Rs. 500.

8.4.1 Front-Loaded Incentive Treatment

The front-loaded treatment is designed to raise the relative price of early effort for present-biased workers, incentivizing them to front-load their genuine work and thereby reduce the late-cycle backlog that the model predicts drives misreporting. Workers in this treatment received a base participation payment of Rs. 200, plus a per-visit bonus linked to the timing of their reported visits. The per-visit bonus was Rs. 12.5 for each visit reported in the first three days of the work week (Monday through Wednesday) and Rs. 4.2 for each visit reported in the last three days (Thursday through Saturday), for a maximum of 36 visits paid.

The treatment script was delivered in Urdu and Pashto. The translated Urdu version reads:

Front-Loaded Script *(translated from Urdu):* "For completing this activity and submitting the report to your LHS, you will receive a bonus of Rs. 200. But you can also earn additional bonus beyond this Rs. 200. We will give you an additional bonus for visiting each household

and collecting data—for up to 36 households. For households you visit and collect data from during the *first* three days of the work week, you will receive Rs. 12.5 per visit. For households you visit and collect data from during the *last* three days of the work week, you will receive Rs. 4.2 per visit. This means that if you complete all 36 visits in the first three days, you will receive Rs. 450 in addition to your base Rs. 200—for a total of Rs. 650. Based on your data, we will prepare a report for the health department. Please complete this activity over six days and submit your report to your LHS."

Mathematically, the payment to worker i in the front-loaded treatment is:

$$w_i^F = 200 + 12.5 \cdot \min(n_i^1,\, 36) + 4.2 \cdot \max\big(\min(n_i^1 + n_i^2,\, 36) - n_i^1,\, 0\big)\,,$$

where n_i^1 is the number of visits reported in the first three days and n_i^2 in the last three days, with total reported visits capped at 36 for payment. Under a uniform distribution of visits across the week (18 per half), the expected payment equals Rs. 500, maintaining payoff equivalence with the other treatment arms.[10]

8.4.2 Uniform Incentive Treatment

The uniform incentive treatment offers a constant per-visit payment throughout the work week. Workers received a base participation payment of Rs. 200 plus Rs. 8.3 per reported visit, for a maximum of 36 visits (maximum total: Rs. 499, approximately Rs. 500). This treatment tests the effect of standard pay-for-performance incentives, holding incentive level constant at the same expected payoff as the front-loaded treatment but removing the intertemporal price differential.

Uniform Script *(translated from Urdu):* "For completing this activity and submitting the report to your LHS, you will receive a bonus of Rs. 200. But you can also earn additional bonus beyond this Rs. 200. For any household you visit and collect data from during the activity week, you will receive approximately Rs. 8.3 per visit, for up to 36 visits. Completing all 36 visits would earn you Rs. 300 in addition to your Rs. 200 base—for a total of Rs. 500. Based on your data, we will prepare a report for the health department. Please complete this activity over six days and submit your report to your LHS."

[10]The per-visit rates were calibrated so that expected payoff equals Rs. 500 when visits are distributed evenly across the two halves of the work week, which is the optimal allocation for a time-consistent worker facing a uniform bonus. The ratio of early to late per-visit rates ($12.5/4.2 \approx 3$) was chosen to meaningfully shift the relative cost of early versus late effort while remaining operationally tractable.

8.4.3 Control

Workers assigned to the control group received a flat bonus of Rs. 500 for completing at least 36 household visits and submitting their paper report to the LHS. The control script contained no information about the importance of data and no performance-linked payment.

Control Script *(translated from Urdu):* "Please remember that during this activity you must visit at least 36 households over the six-day work week. For participating in this activity and submitting the report to your LHS, you will receive a bonus of Rs. 500. Based on this data, we will prepare a report for the health department. Please complete this activity over six days and submit your report to your LHS."

8.5 Phase 4: Independent Audit of Reports

To distinguish genuine visits from fabricated ones, we implemented an independent household audit following the submission of each worker's paper report. For each worker, an audit team randomly selected one reported household visit per day of the work week—six households per worker in total. Audit teams were deployed immediately upon collection of the paper forms from the LHS to minimize recall bias.[11] Independent audit enumerators—with no knowledge of the worker's treatment assignment—visited each selected household in person and asked residents to confirm or deny that the LHW had visited on the reported date.

8.6 Randomization, Sample, and Timeline

Workers were randomized at the individual level stratified by the center to which they are attached, acting as a cluster. Within the present-biased stratum (approximately 176 workers), workers were randomized with equal probability into the front-loaded incentive, uniform incentive, and control conditions.

The study was conducted from October 2025 through June 2026. Baseline surveys and CTB elicitation were completed in October and November 2025, followed by pilot testing in November and December 2025. Main experiment treatment calls were delivered in January 2026, with independent audits conducted in the days immediately following each reporting week.

[11] In cases where a worker did not report a visit on a particular day, we oversampled from the remaining days and weight these observations accordingly in the analysis.

8.7 Data Sources

We use data from three sources to measure worker behavior.

Worker-reported visit data. Each LHW filled out a standardized paper-based form recording the date and household identifier of every reported visit during the six-day work cycle. Workers submitted these forms to their LHS at the end of the week, and our field teams collected them from the LHS. These forms provide the daily time series of self-reported visit counts.

Audit data. For each worker, independent audit teams verified one randomly selected reported household per day (six households per worker). For each selected household, enumerators recorded whether the resident confirmed the LHW's visit. These data yield the probability that a reported visit is genuine and are used to construct both visit-level falsification indicators and worker-level estimates of total false reports.

Baseline worker survey. At baseline, we collected information on workers' demographics (age, education, tenure, catchment area size), behavioral preferences, and prior-month productivity. These data are used to verify randomization balance.

Table 1: Summary statistics

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
age	1994.542	6.167	1980	2010	301
total_visits	122.05	26.093	25	221	301
experience	18.422	6.754	2	30	301
donation	49.166	26.11	0	100	301
avg_psm	4.153	0.378	2.8	5	301
avg_big5	3.566	0.356	2.55	4.933	301
median_beta	0.977	0.229	0.478	1.824	305
Present Biased Sample					
age	1993.701	5.985	1980	2010	174
total_visits	122.58	28.502	40	221	174
experience	19.086	6.454	2	29	174
donation	44.098	24.307	0	100	174
avg_psm	4.185	0.374	3.35	5	174
avg_big5	3.579	0.347	2.667	4.933	174
median_beta	0.805	0.126	0.478	0.999	176

8.8 Randomization Balance

Table 2 reports the balance of pre-treatment characteristics across the three treatment arms in the present-biased sample. For each of the seven baseline characteristics—worker age, prior-month visit count, years of experience, dictator-game giving, public service motivation (PSM), the Big Five personality index, and the estimated present-bias parameter $\hat{\beta}$—the table reports the arm-specific mean together with a joint orthogonality test of equality across the three arms. The randomization is well balanced: none of the seven characteristics is significantly imbalanced, as every joint orthogonality p-value exceeds 0.10, the smallest being 0.17 for public service motivation. Reassuringly, the estimated present-bias parameter itself is statistically indistinguishable across arms ($p = 0.89$), confirming that workers' time-preference type is orthogonal to treatment assignment within the present-biased stratum.

Table 2: **Randomization Balance: Present-Biased Sample**

	Age	*Monthly Visits*	*Experience* in Years	*Dictator Game Donation*	*PSM*	*Big 5*	*Median* $\hat{\beta}$
High Upfront Incentive	1995.293	125.296	16.426	31.795	4.163	3.634	0.844
	(1.926)	(5.501)	(1.821)	(5.775)	(0.156)	(0.061)	(0.042)
Uniform Incentive	1994.275	125.837	17.178	32.264	4.053	3.620	0.853
	(1.959)	(5.936)	(1.910)	(5.955)	(0.158)	(0.066)	(0.041)
Control	1993.668	126.435	17.253	33.676	4.154	3.552	0.844
	(2.150)	(5.832)	(2.078)	(5.914)	(0.149)	(0.062)	(0.044)
Joint orthogonality p-value	0.32	0.97	0.72	0.91	0.17	0.31	0.89
# of Workers	174	174	174	174	174	174	176

Notes: This table tests if random assignment is orthogonal to the value of the baseline characteristics $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

8.9 The Rising Profile of Dishonesty

We begin by documenting the time trend in falsification, which is the central empirical regularity that motivates the theoretical model. Panel A of Figure 2 shows the mean probability of a false report by day of the work week (Monday through Saturday) in the control group. The falsification rate rises steeply from roughly 0.10–0.15 in the first two days of the work week to above 0.55 by Friday and Saturday. Panel B shows the corresponding time trend in total claimed visits per day; claimed visits also rise toward the end of the week, consistent with workers cramming their reported output near the deadline.

The model predicts that the slope of this rising profile should be *steeper* for more present-biased workers—for whom the backlog channel is more pronounced—above the common deadline-pressure component. Figure 3 separates the control group by estimated $\hat{\beta}$ stratum. Present-biased workers (red dashed line, $\hat{\beta} < 1$) show a notably steeper rise in falsification

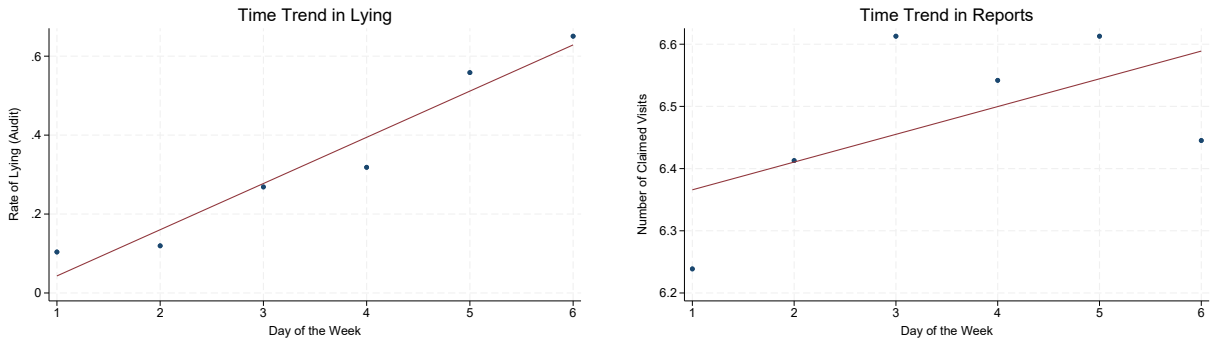

Figure 2: **Time Trend in Falsification and Claimed Visits (Control Group)**

Notes: Panel A (left) plots the mean probability of a false report by day of the work week among control group workers (both present-biased and non-present-biased controls). Panel B (right) plots the mean number of claimed visits per day. The x-axis shows the day of the work week (1 = Monday, 6 = Saturday). Fitted using a binscatter estimator.

than non-present-biased workers (blue solid line, $\hat{\beta} \geq 1$), particularly in the second half of the work week when the procrastination-induced backlog becomes binding.

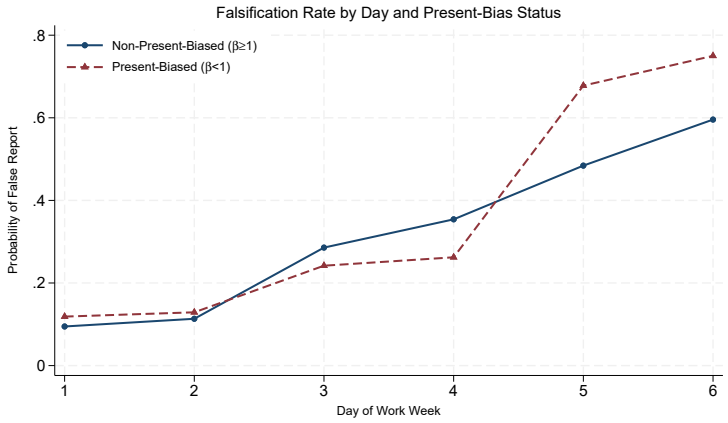

Figure 3: **Falsification Profile by Present-Bias Status (Control Group)**

Notes: The figure plots the mean probability of a false report by day of the work week for control group workers, separately for present-biased ($\hat{\beta} < 1$; red dashed line) and non-present-biased ($\hat{\beta} \geq 1$; blue solid line) workers. The divergence in slopes toward the deadline is consistent with the present-bias channel predicted by Proposition 3.

Table 3 formalizes this pattern with a regression of the falsification rate on $\hat{\beta}$, day of week, and their interaction, estimated on the control group only. The key coefficient is $\hat{\beta} \times$ Day of Week. Column 1 (full control sample) estimates this interaction at -0.089 ($p < 0.01$): for each unit decrease in $\hat{\beta}$ (more present bias), the slope of the rising profile increases by 0.089 per day. This is a large and precisely estimated effect—a worker with $\hat{\beta} = 0.5$ has a within-week slope that is 0.045 higher per day than a time-consistent worker, implying an excess falsification probability of 0.22 accumulated over the five-day interval from day 1 to day 6. Column 3 (non-present-biased control) shows among the strongest effect (-0.089; $p < 0.01$),

consistent with identifying the interaction from within-stratum variation among workers for whom $\hat{\beta} \geq 1$: even within this group, workers with values of $\hat{\beta}$ closer to 1 exhibit steeper rising profiles. The common component of the time trend (the coefficient on day of week without the interaction) is large and positive, confirming the presence of deadline pressure as a background force for all workers.

Table 3: **Present-Bias Heterogeneity in the Rising Falsification Profile (Control Group)**

Dep Var: Pr(False Report = 1)	(1)	(2)	(3)
	Full Control	PB Control	NPB Control
Present Bias ($\hat{\beta}$)	0.266**	0.086	0.298**
	(0.108)	(0.411)	(0.143)
Day of Week	0.209***	0.211**	0.209***
	(0.030)	(0.085)	(0.043)
$\hat{\beta}$ × Day of Week	-0.089***	-0.091	-0.089**
	(0.028)	(0.105)	(0.036)
Control Mean	0.332	0.352	0.320
# Observations	930	355	575

Notes: Each column regresses the audit-based false report indicator on the estimated present-bias parameter $\hat{\beta}$, the day of the work week (1–6), and their interaction, among control group workers only. Column 1 pools both present-biased and non-present-biased controls; column 2 restricts to present-biased controls ($\hat{\beta} < 1$); column 3 restricts to non-present-biased controls ($\hat{\beta} \geq 1$). All regressions absorb Lady Health Supervisor fixed effects. Significant negative interactions ($\hat{\beta}\times$ Day) indicate that workers with lower $\hat{\beta}$ (more present-biased) show steeper rising profiles of falsification. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

9 Main Analysis and Results

This section presents the main experimental results. Throughout, we restrict attention to the present-biased subsample—the 176 workers for whom the structural elicitation returns $\hat{\beta} < 1$—because the behavioral mechanism the paper tests, and the time-structured incentives we designed to counter it, apply specifically to workers whose genuine effort is distorted by procrastination.[12] We proceed in four steps. We first lay out the estimation strategy. We then report treatment effects on effort and on misreporting, at both the worker-day and worker levels. We next decompose these effects by the timing of effort within the work

[12]The non-present-biased workers received a separate, non-pecuniary treatment (the task importance message) rather than the time-structured financial incentives, and are analyzed separately. Pooling the two subsamples would conflate two distinct behavioral channels and two distinct interventions, and the pilot evidence in Section 8.4 indicates that the time-varying incentive induces gaming rather than effort among non-present-biased workers, which is why it was confined to the present-biased stratum.

week, which is where the theory makes its sharpest predictions. Finally, we ask whether the treatment effects scale with the measured degree of present bias.

9.1 Estimation Equation

To estimate the causal effect of the time-structured incentives on the present-biased sample, we estimate

$$Y_{ijt} = \alpha_0 + \alpha_1 \,\text{Upfront}_j + \alpha_2 \,\text{Uniform}_j + \gamma_t + \delta_c + \varepsilon_{ijt}, \tag{11}$$

where Y_{ijt} is an outcome for worker j (or for household i in the catchment area of worker j) on day t of the six-day work week. Upfront_j and Uniform_j are indicators for assignment to the front-loaded and uniform incentive arms, respectively, so that each coefficient is measured relative to the omitted control arm, which received a flat Rs. 500 bonus with no performance link. γ_t denotes day-of-week fixed effects, which absorb the common within-week dynamics of reporting and falsification—most importantly the steep deadline-pressure trend documented in Section 8.9—so that treatment effects are identified from between-worker variation within a given day and supervisor cluster rather than from the shared calendar profile. δ_c denotes Lady Health Supervisor cluster fixed effects, which net out differences in local monitoring intensity, catchment characteristics, and audit logistics across supervisory areas. For worker-level analyses that collapse the data across days, we drop the day-of-week fixed effects and retain only the supervisor fixed effects. Because randomization was conducted at the worker level and each worker contributes multiple worker-day observations, we cluster standard errors at the worker level throughout. Given the balance of pre-treatment characteristics within the present-biased stratum reported in Table 2, α_1 and α_2 recover the average causal effects of the two incentive structures. Because both arms were calibrated to the same Rs. 500 expected payoff, the contrast between α_1 and α_2 isolates the effect of the *timing* of the incentive, holding its *level* fixed—the central comparison of the experiment.

9.2 Treatment Effects on Effort and Misreporting

We begin with the two primary behavioral outcomes: the intensity of effort, measured by the number of household visits a worker claims per day, and the honesty of reporting, measured by the audit-verified probability that a claimed visit is fabricated. Table 4 reports these effects at the worker-day level, the finest unit at which both effort and falsification are observed. Column 1 takes the number of claimed visits per day as the dependent variable. Columns 2 and 3 take the probability that a report is false as the dependent variable, differing in how they treat households the audit could not resolve: column 2 uses the strict

audit-based indicator, defined only for the households an audit team was able to locate and classify, while column 3 uses an extended measure that additionally codes as false those reported households the auditors could not locate, on the conservative presumption that an untraceable household was fabricated. All three specifications absorb day-of-week and supervisor fixed effects, and the foot of the table reports the control-group mean of each outcome so that every coefficient can be read against the relevant counterfactual.

Table 4: **Effects on Worker Behavior: Visits and False Reports**

	Visits per Day	*Pr(False Report = 1)*	*Pr(False Report = 1)*
High Upfront Incentive	0.65***	-0.03	-0.02
	(0.14)	(0.03)	(0.03)
Uniform Incentive	0.09	-0.05	-0.06*
	(0.14)	(0.03)	(0.03)
Control Mean	6.594	0.352	0.365
# Observations	1057	1057	1083

Notes: Regressions at the worker-day level estimated on the present-biased sample ($\hat{\beta} < 1$, audited observations) using equation (11). The dependent variable is the number of claimed visits per day in column 1, the strict audit-verified probability that a report is false in column 2, and an extended falsification measure in column 3 that additionally codes as false those reported households the audit teams could not locate (treating untraceable households as fabricated). All regressions absorb day-of-week and Lady Health Supervisor fixed effects. "Control Mean" is the mean of the dependent variable in the control arm. Standard errors clustered at the worker level in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The front-loaded incentive produces a substantial and statistically significant increase in effort. Workers assigned to the front-loaded arm claim 0.65 additional visits per day relative to the control group ($p < 0.01$, column 1), an increase of 9.8% over the control mean of 6.59 daily visits. This is an economically meaningful response: with a binding quota of 36 visits over the six-day week, an effect of this size moves the typical worker from just meeting her quota to comfortably exceeding it, and the worker-level collapse in Table 5 confirms the sign, with total claimed visits over the cycle rising by 2.18 (column 1). The decisive comparison, however, is with the uniform arm. The uniform incentive carries an identical Rs. 500 expected payoff but flattens the intertemporal price of effort, and it has no detectable effect on claimed visits whatsoever (0.09; $p > 0.10$). That an incentive of exactly the same expected value moves effort only when its price schedule is tilted toward the early week is the pattern predicted by the backlog mechanism of Proposition 2: raising the relative price of early effort draws genuine work forward and forestalls the late-week cramming into which a present-biased worker would otherwise default.[13]

[13] An additional margin of misreporting is concordance between the days of actual versus report days of work. Figure 5 presents the concordance between the visit date reported by the worker and the visit date recalled by the audited household among confirmed visits, suggesting good temporal accuracy in reports that are verified as genuine.

Crucially, the extra claimed visits under the front-loaded incentive are not manufactured. The estimated effect of the front-loaded arm on the probability that a report is false is small, negative, and statistically insignificant under both the strict audit measure (-0.03; column 2) and the conservative extended measure that treats untraceable households as fabricated (-0.02; column 3). Read jointly with the visit result, this is the key honesty finding for the present-biased sample: the front-loaded incentive raises the volume of claimed visits by 9.8% while leaving the falsification rate statistically unchanged, which implies that the additional claims are predominantly *genuine* rather than fabricated. The intervention therefore expands real service delivery and not merely the paper record of it, and one that distinguishes a behaviorally targeted incentive from a blunt pay-for-report scheme that would be expected to inflate the count and the lies in equal measure. The uniform incentive, by contrast, doesn't raise effort but lowers the pooled falsification rate (a significant -0.06 in column 3). As we show in Section 9.3, this pooled effect masks a sharp honesty effect that is confined to the late week and is washed out here by averaging against an unaffected early week.

Table 5 aggregates the data to the worker level and expresses the results in the natural units of the policy problem: the total number of visits a worker claims over the cycle, and the estimated number of those claims that are fabricated, where the latter is constructed as total claimed visits multiplied by the worker's audit-based falsification rate. The figures for the control arm are themselves a central descriptive result. The control worker claims 37.9 visits over the week and fabricates an estimated 13.3 of them, so that roughly 36% of all reported activity under the status quo is false—a baseline that quantifies the scale of the misreporting problem this paper sets out to explain. Both incentive arms shift estimated fabrication in the honest direction: the front-loaded arm lowers estimated false visits by 0.57 and the uniform arm by 1.61, the latter a 12% reduction relative to the control mean of 13.3. Neither worker-level estimate reaches conventional significance (standard errors of 1.35 and 1.37), which is to be expected: collapsing 1057 worker-day observations into 175 worker-level observations discards most of the within-week variation that identifies the effects, and with it most of the statistical power. The worker-day and timing specifications, which retain that variation, recover the same patterns more precisely.

9.3 Timing of Effort and Misreporting

The pooled effects in Tables 4 and 5 average over the work week, but the theory makes sharp predictions about *when* within the week each incentive should bite, and these timing predictions are the most direct test of the present-bias mechanism. Present bias operates through

Table 5: **Effects on Total Claimed and Estimated False Visits (Worker Level)**

	Total Claimed Visits	*Estimated False Visits*
High Upfront Incentive	2.18**	-0.57
	(0.93)	(1.35)
Uniform Incentive	0.98	-1.61
	(0.94)	(1.37)
Control Mean	37.881	13.279
# Observations	175	175

Notes: Regressions at the worker level using the collapse of audited observations, absorbing Lady Health Supervisor fixed effects. Estimated false visits = total claimed visits × Pr(false). Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the filing path: a procrastinating worker lets genuine work accumulate and is then forced into heavy, fabrication-prone filing in the final days, so the falsification rate climbs toward the deadline (Proposition 2). The front-loaded incentive is designed to defuse this backlog by raising the relative price of early effort, and should therefore lift early-week claimed visits and depress late-week fabrication. The uniform incentive leaves the intertemporal price flat and so does not alter the relative returns to early versus late effort — yet Table 6 shows a significant increase of 1.38 in late-week visits (column 2), likely reflecting a level effect: per-visit pay makes genuine effort more valuable throughout the week and amplifies workers' pre-existing tendency to concentrate visits late. At the same time, by raising the marginal return to every genuine visit, it makes honest completion more attractive than fabrication at the margin, an effect the model locates precisely in the late week where the temptation to fabricate is greatest. Figure 4 provides the visual counterpart to these predictions, plotting the mean number of claimed visits (Panel A) and the mean falsification rate (Panel B) by day of the work week for each arm in the present-biased sample.

Panel A shows that the front-loaded arm lifts claimed visits above the control profile through the first half of the week, consistent with the incentive drawing genuine effort forward, while the uniform arm tracks the control profile closely throughout——-consistent with the absence of an intertemporal price differential—-before diverging modestly upward in the late half. Panel B shows that falsification in the front-loaded arm falls below the control profile in the second half of the week, exactly the interval in which procrastination-induced backlog would otherwise push the falsification rate to its peak, although the day-level series is noisy.

Table 6 formalizes the timing analysis by splitting each outcome into an early half (days 1–3, Monday–Wednesday) and a late half (days 4–6, Thursday–Saturday), and three patterns stand out. First, the effort response to the front-loaded incentive is concentrated early:

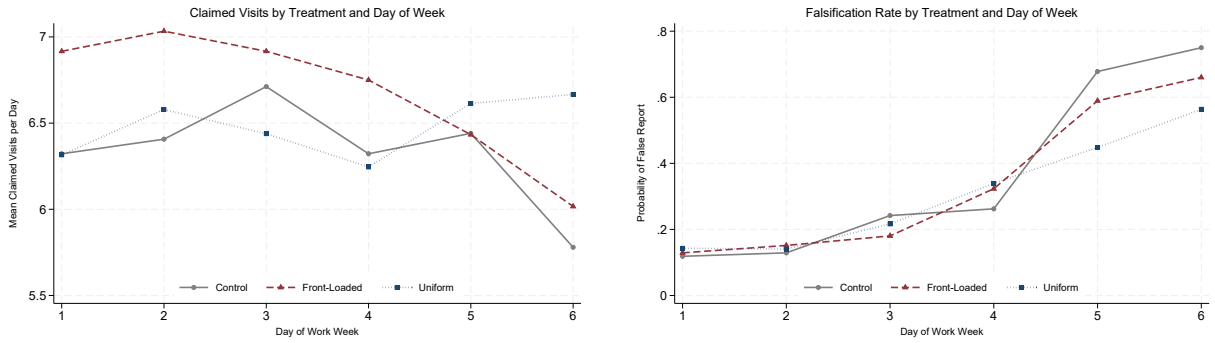

Figure 4: **Claims and Falsification by Treatment and Day of Week (Present-Biased Sample)**

Notes: Panel A (left) plots the mean number of claimed visits per day for each treatment arm (present-biased sample). Panel B (right) plots the mean probability of a false report per day for each treatment arm. Day 1 = Monday, Day 6 = Saturday.

claimed visits rise by 1.43 in the early half (column 1) against an increase of only 0.98 in the late half (column 2), the asymmetry predicted by the front-loading mechanism, with the difference between the two halves individually significant ($p<0.05$). Second, the falsification estimates confirm that the deadline is where the action is. In the control arm the probability of a false report is just 0.164 in the early half but 0.552 in the late half (columns 3–4)—more than half of all late-week reports are fabricated—which is the within-week fingerprint of backlog-driven lying that motivates the whole study. Against this backdrop the uniform incentive delivers the single sharpest honesty result in the experiment: it lowers the late-half probability of a false report by 0.13 ($p < 0.05$, column 4), a roughly 24% reduction relative to the late-week control mean of 0.552, while leaving the early half statistically untouched (an insignificant +0.02, column 3). The front-loaded arm moves the late-week falsification rate in the same direction (−0.06, column 3) but does not attain significance. Third, translated into visit counts (columns 5–6), the late-half reduction in estimated fabrication is 1.65 visits for the uniform arm and 0.60 for the front-loaded arm—both economically sizable against a late-half control mean of 10.3 fabricated visits, though imprecisely estimated at the worker level. Taken together, the timing decomposition shows the two incentives operating through different margins of the model: the front-loaded incentive primarily reallocates *genuine* effort toward the early week and so shrinks the backlog at its source, whereas the uniform incentive primarily raises the marginal attractiveness of *honest* completion in the late week, where fabrication is otherwise concentrated. The two interventions thus map onto the two arguments of the falsification rate in Lemma 1: the front-loaded arm works through the filing path q_t, the uniform arm through the relative payoff to genuine completion at the deadline.

Table 6: **Effects by Timing: Early vs. Late in the Work Week**

	Claimed Visits		*Pr(False Report = 1)*		*Estimated False Visits*	
	(1)	(2)	(3)	(4)	(5)	(6)
High Upfront Incentive	1.43**	0.98	0.00	-0.06	0.14	-0.60
	(0.69)	(0.74)	(0.04)	(0.05)	(0.96)	(1.24)
Uniform Incentive	-0.15	1.38*	0.02	-0.13**	0.15	-1.65
	(0.70)	(0.75)	(0.04)	(0.05)	(0.97)	(1.25)
Control Mean	19.441	18.441	0.164	0.552	3.017	10.312
# Observations	175	173	545	512	175	173
Time Period	Early	Later	Early	Later	Early	Later

Notes: Regressions at the worker × half-week level. Early = days 1–3 (Monday–Wednesday); Late = days 4–6 (Thursday–Saturday). Present-biased sample, audited observations only. LHS fixed effects absorbed. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

9.4 Dose-Response with $\hat{\beta}$

A natural further test is whether the treatment effects scale with the measured degree of present bias: the backlog mechanism predicts that workers with lower $\hat{\beta}$, who procrastinate more, should benefit most from an incentive engineered to counter procrastination. Table 7 probes this dose–response by augmenting equation (11) with a demeaned continuous $\hat{\beta}$ and its interactions with the two treatment indicators, with claimed visits per day in column 1 and the probability of a false report in column 2. The main effect of $\hat{\beta}$ on daily claimed visits is large, negative, and precisely estimated (-2.17; $p < 0.01$), indicating that within the present-biased group more present-biased workers (lower $\hat{\beta}$) mechanically claim more visits per day on average—consistent with these workers relying more heavily on the reporting margin. The treatment-by-$\hat{\beta}$ interactions are positive for both arms (1.37 for front-loaded and 0.91 for uniform) but are imprecisely estimated and far from significant, and the treatment main effects themselves become unstable relative to Table 4 once the interactions are included. We read this instability as a power limitation rather than as evidence on the mechanism. By construction the present-biased subsample restricts $\hat{\beta}$ to the narrow interval below one (Table 1 reports a standard deviation of just 0.13 in this group), leaving little identifying variation with which to separate the treatment main effects from their interaction with $\hat{\beta}$; the treatment dummies, $\hat{\beta}$, and their products are correspondingly collinear, and the continuous dose–response is uninformative in the current sample. We therefore refrain from interpreting the magnitude or sign of the interaction terms, and regard a sharper test of the dose–response—which exploits the full continuous distribution of $\hat{\beta}$—as awaiting the completion of data collection.

Table 7: **Dose-Response with Estimated Present Bias (Present-Biased Sample)**

	Visits per Day	*Pr(False) = 1*
Front-Loaded	-0.451	-0.030
	(0.887)	(0.202)
Uniform	-0.628	-0.198
	(0.944)	(0.229)
$\hat{\beta}$ (demeaned)	-1.949**	-0.234
	(0.781)	(0.176)
Front-Loaded $\times$ $\hat{\beta}$	1.372	0.003
	(1.093)	(0.247)
Uniform $\times$ $\hat{\beta}$	0.911	0.182
	(1.157)	(0.278)
Control Mean	6.594	0.352
# Observations	1057	1057

Notes: The specification augments equation (11) with a demeaned continuous $\hat{\beta}$ and its interactions with treatment dummies. Present-biased sample, audited observations. Day-of-week and LHS fixed effects absorbed. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

9.5 Interpretation

The results for the present-biased sample line up with the two-force framework of Section 2 in a specific and internally consistent way. The front-loaded incentive—which raises the relative price of early effort—acts on the present-bias channel: it pulls genuine visits into the early week (Table 6, columns 1–2) and raises total genuine effort by 9.8% without inflating the falsification rate (Table 4). This is the signature of an intervention that shrinks the procrastination-induced backlog rather than one that simply purchases more reports, honest or otherwise. The uniform incentive, which holds expected payoff fixed but strips out the intertemporal price differential, does not alter the relative returns to early versus late effort — yet appears to amplify late-week genuine visits (Table 6, column 2), likely as a level effect. It still improves honesty exactly where the model says it should — cutting the late-week falsification rate by roughly a quarter (Table 6, column 4).

That two incentives of identical expected value leave such different behavioral footprints is the central message of the experiment. For present-biased workers, the *timing* of performance pay, not merely its level, governs both how much genuine effort they supply and how honestly they report it: tilting the same Rs. 500 toward the early week buys real effort, while spreading it evenly buys late-week honesty, and the flat-bonus control buys neither. A policymaker who cares about the integrity of self-reported administrative data therefore has a lever that operates through the structure of the incentive rather than its cost. Two

caveats temper these conclusions. First, the worker-level and dose–response estimates are imprecise, reflecting a present-biased sample of 176 workers and, for the dose–response, the narrow support of $\hat{\beta}$ within that sample; the worker-day and timing specifications, which preserve within-week variation, carry the statistical weight of the argument. Second, because the analysis is confined to present-biased workers by design, these magnitudes speak to the population for whom the mechanism is active and should not be extrapolated to workers whose rising falsification profile is driven by deadline pressure alone. We expect the completion of data collection to sharpen the worker-level and dose–response estimates without altering the qualitative pattern established here.

10 Conclusion

This paper provides the first field experimental evidence that present bias is a mechanism for misreporting in the public sector, and demonstrates that the timing of financial incentives can address this behavioral source of dishonesty. Using a Convex Time Budget elicitation to measure individual present-bias parameters, and a subsequent multi-arm randomized experiment targeting treatments to workers' behavioral types, we show that the rising profile of falsification within a work week is partly driven by procrastination-induced backlog accumulation—a signature of present bias—and not solely by deadline pressure acting on the moral cost of fabrication.

The practical implication is that public sector incentive design should account not just for the level of performance pay but for its timing. A front-loaded incentive scheme that rewards early-week effort at a higher per-unit rate than late-week effort—while holding expected payoff constant—significantly increases genuine effort among present-biased workers. Standard uniform incentives of equivalent expected value do not produce similar gains, suggesting that the intertemporal price dimension of performance contracts matters independently of the incentive level.

References

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond, "Preferences for truth-telling," *Econometrica*, 2019, *87* (4), 1115–1153.

Andreoni, James and Charles Sprenger, "Risk preferences are not time preferences," *American Economic Review*, 2012, *102* (7), 3357–76.

_ , Michael Callen, Karrar Hussain, Muhammad Yasir Khan, and Charles Sprenger, "Using preference estimates to customize incentives: an application to polio vaccination drives in Pakistan," *Journal of the European Economic Association*, 2022, p. jvac068.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger, "Working over time: Dynamic inconsistency in real effort tasks," *The Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

Banerjee, Abhijit V, Esther Duflo, and Rachel Glennerster, "Putting a band-aid on a corpse: incentives for nurses in the Indian public health care system," *Journal of the European economic association*, 2008, *6* (2-3), 487–500.

Bó, Ernesto Dal, Frederico Finan, and Martín A Rossi, "Strengthening state capabilities: the role of financial incentives in the call to public service," *The Quarterly Journal of Economics*, 2013, *128* (3), 1169–1218.

Callen, Michael, Saad Gulzar, Ali Hasanain, Muhammad Yasir Khan, and Arman Rezaee, "Data and policy decisions: Experimental evidence from Pakistan," *Journal of Development Economics*, 2020, p. 102523.

_ , _ , _ , _ , and _ , "The political economy of public sector absence," *Journal of Public Economics*, 2023, *218*, 104787.

Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F Halsey Rogers, "Missing in action: teacher and health worker absence in developing countries," *Journal of Economic perspectives*, 2006, *20* (1), 91–116.

DellaVigna, Stefano and Ulrike Malmendier, "Paying not to go to the gym," *American Economic Review*, 2006, *96* (3), 694–719.

Duflo, Esther, Rema Hanna, and Stephen P Ryan, "Incentives work: Getting teachers to come to school," *American economic review*, 2012, *102* (4), 1241–1278.

Finan, Frederico, Benjamin A. Olken, and Rohini Pande, "The personnel economics of the developing state," in Abhijit Vinayak Banerjee and Esther Duflo, eds., *Handbook of Economic Field Experiments*, Vol. 2, North-Holland, 2017, pp. 467–514.

Fischbacher, Urs and Franziska Föllmi-Heusi, "Lies in disguise—an experimental study on cheating," *Journal of the European Economic Association*, 2013, *11* (3), 525–547.

Fisman, Raymond and Edward Miguel, "Corruption, norms, and legal enforcement: evidence from diplomatic parking tickets," *Journal of Political Economy*, 2007, *115* (6), 1020–1048.

Gneezy, Uri, "Deception: The role of consequences," *American Economic Review*, 2005, *95* (1), 384–394.

Hanna, Rema and Shing-Yi Wang, "Dishonesty and selection into public service: Evidence from India," *American Economic Journal: Economic Policy*, 2017, *9* (3), 262–290.

Holmstrom, Bengt and Paul Milgrom, "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design," *Journal of Law, Economics and Organizations*, 1991, *7*, 24.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, "Self-control at work," *Journal of Political Economy*, 2015, *123* (6), 1227–1277.

Khan, Adnan Q, Asim I Khwaja, and Benjamin A Olken, "Tax farming redux: experimental evidence on performance pay for tax collectors," *The Quarterly Journal of Economics*, 2016, *131* (1), 219–271.

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

Muralidharan, Karthik and Venkatesh Sundararaman, "Teacher performance pay: Experimental evidence from India," *Journal of political Economy*, 2011, *119* (1), 39–77.

Nagin, Daniel S, James B Rebitzer, Seth Sanders, and Lowell J Taylor, "Monitoring, motivation, and management: the determinants of opportunistic behavior in a field experiment," *American Economic Review*, 2002, *92* (4), 850–873.

O'Donoghue, Ted and Matthew Rabin, "Doing it now or later," *American economic review*, 1999, *89* (1), 103–124.

Olken, Benjamin A, "Monitoring corruption: evidence from a field experiment in Indonesia," *Journal of Political Economy*, 2007, *115* (2), 200–249.

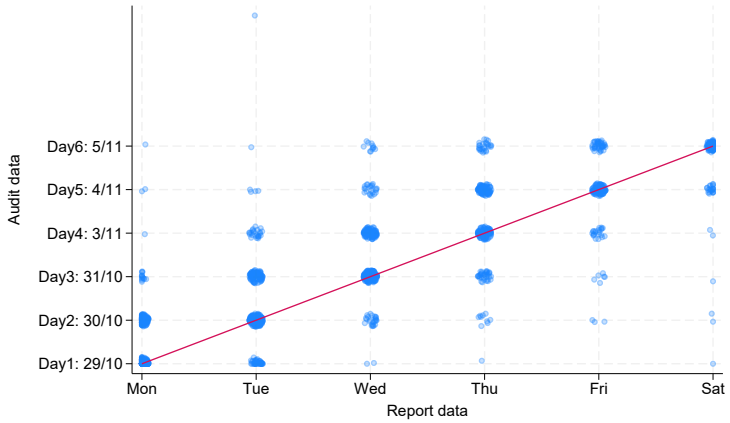

Figure 5: **Concordance Between Reported and Audit-Recalled Visit Date**

Notes: Scatter plot of the visit date recorded in the worker's paper report (x-axis) against the date of most recent LHW visit recalled by the household at the time of audit (y-axis), for households that confirmed a visit. The 45-degree line indicates perfect concordance. Jitter applied to aid visualization.

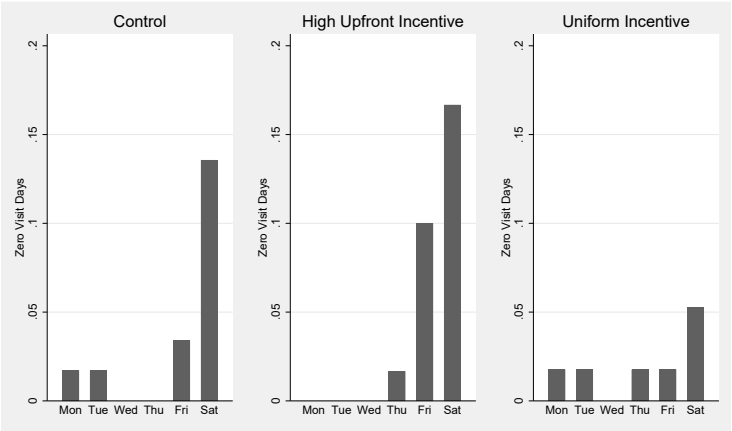

Figure 6: **Number of Workers Reporting Zero Visits By Day**



theigc.org