

C A G E

813/2026
July 2026

Working Paper

Paying for power

Fiona Burlig

Anant Sudarshan

ISSN: 2978-0276

Grant number: ES/7504701/1

**UNIVERSITY
OF WARWICK**



**Economic
and Social
Research Council**

PAYING FOR POWER

Fiona Burlig and Anant Sudarshan*

June 17, 2026

Abstract

Developing country governments routinely attempt to collect revenue using threats they cannot systematically enforce. We study how citizens assess the credibility of such empty threats in the context of payment for electricity in Madhya Pradesh, India, where state-run utilities recovered only 60 cents per dollar of power supplied. Using two field experiments covering 30,000 households with high arrears, we show that the household response to a threat depends on the state's choice of messenger. The first experiment randomly exposes households to reminders, threats, and enforcement action without changing incentives, state capacity, laws, or information about debt. Legal threats delivered by local linesmen – state agents with a history of ignoring non-payment – have no effect. Yet identical notices sent by registered mail, bypassing linesmen, reduce arrears by 11.4 percent among recipients, a 241 percent return-on-investment. We hypothesize that choosing compromised messengers changes household beliefs about the state's credibility, implying dynamic effects that we test with the second experiment: a year later, we randomly mail a legal notice to previously-treated households. Past treatments affect future responses. Consumers originally visited by a linesman do not respond, while those not exposed to linesmen reduce arrears. Moreover, when we increase linesman credibility by requiring them to follow up on threats, this gap narrows. The experimental results are together consistent with a model in which consumers use the state's choice of messenger to infer the threat's credibility, and demonstrate that low-credibility state agents can render threats ineffective. Low-capacity governments may improve revenue collection by bypassing their agents.

Keywords: electricity; threats; messengers; state capacity; taxation

JEL Codes: O13; Q48; H26

*Burlig: Harris School of Public Policy and Energy Policy Institute (EPIC), University of Chicago, and NBER. Email: burlig@uchicago.edu. Sudarshan: Department of Economics, University of Warwick. Email: anant.sudarshan@warwick.ac.uk. We thank Susanna Berkouwer, Manasi Deshpande, Koichiro Ito, Ryan Kellogg, Mushfiq Mobarak, Peter Ganong, Michael Greenstone, Catherine Wolfram, Abhijeet Singh, E Somanathan, Lucie Gadenne, and seminar participants at Midwest Energy Fest, UC Energy Camp, the Coase Project, USC, the University of Warwick, LSE, UCSD, Cornell, and UCSB for helpful comments and suggestions. We gratefully acknowledge financial support from the Oak Foundation and the Griffin Applied Economics Incubator. Abhishek Deshwal, Animesh Jayant, Simran Kalra, Kabir Nagadia, Garrison Schlauch, Anjaney Singh, and Rathana Sudheer provided excellent research assistance. All remaining errors are our own. This project received IRB approval from the University of Chicago (Protocol No. IRB20-2127), and is registered on the AEA RCT registry (Identification No. AEARCTR-0008742).

“Off with her head!”

— *The Queen of Hearts*

“It’s all her fancy, that: they never executes nobody, you know.”

— *The Gryphon*

Lewis Carroll, *Alice’s Adventures in Wonderland* (1865)

1 Introduction

Revenue collection is an essential function of any state. It is particularly hard in low- and middle-income countries whose governments are responsible not only for taxation, but also for operating massive markets in goods such as electricity, water, telephone connections, and agricultural inputs. Unfortunately, this broad scope is rarely matched by state capacity, with policymakers often forced to rely on an under-resourced and corrupt state machinery that is ill-equipped to either monitor compliance or enforce penalties.¹

Faced with this reality, these governments find themselves levying taxes and sending bills, all the while relying on largely empty threats of penalties to induce compliance. In response, citizens comply sometimes, but not always. What makes an empty threat work in some settings but fail in others becomes a first-order question for the efficiency of revenue collection. In this paper we examine state threats in the context of a type of revenue collection of very large but underappreciated importance in poor countries: payment for power provided by state-run electricity utilities.

The setting for this paper is India, where accumulated losses in the electricity sector reached about 80 billion USD, or nearly 3 percent of GDP, in 2020–2021 (Anand et al. 2025).

We conducted our research in the state of Madhya Pradesh, home to nearly 90 million people,

1. Examples abound, cutting across policy areas: In Indonesia, Olken (2007) documents large gaps between audits and enforcement; in India, Duflo et al. (2018) find that environmental regulators rarely impose penalties even when they are aware of violations; in Pakistan, Khan, Khwaja, and Olken (2016) find that even highly incentivized tax collectors do not broadly target taxpayers. For that matter, wealthy countries also pass laws with limited monitoring or enforcement. A notable example comes from the US Internal Revenue Service, which in 2023 only audited 0.2 percent of individual income tax returns.

where for every dollar’s worth of power consumed, the government recovers less than 60 cents, largely due to unpaid bills and illegal connections. Using administrative data provided by the state electricity utilities, we show that accumulated arrears are large and increasing over time, and this is not a phenomenon driven by poverty alone. Payment rates are low even for richer households in the top decile of electricity consumption. The propensity of new consumers to pay their bills deteriorates as they gain experience with the utility, so that in equilibrium non-zero payments occur in less than 40 percent of billed months. Consistent with this, utility officials report in surveys that a lack of enforcement contributes to the non-payment problem. We describe the context and data sources in Section 2, and motivating evidence on payment behavior in Section 3.

Our paper focuses on a key aspect of making threats – the identity of the messenger who communicates them. This element is critical because although governments can easily make new threats without changing monitoring or enforcement capacity², they must still choose who communicates them. Since the state is not a monolith, policymakers must pick messengers, often choosing local staff such as tax collectors, environmental inspectors, or utility linesmen. Electricity non-payment makes for a useful laboratory, because the amount owed is known to *both* the state *and* the consumer, allowing us to study threats absent information asymmetry. In most tax-avoidance settings, unpaid liabilities are not common knowledge.

Using two sequential field experiments, conducted in partnership with the state electricity utilities of Madhya Pradesh, we test whether these choices matter for the effectiveness of threats, and find that they can be highly consequential. The first experiment, whose design we describe in Section 5, used a sample of 30,000 electricity-consuming households with substantial unpaid debts. These consumers were randomly assigned to treatment arms, each

2. Elections and settlement schemes are prime opportunities for fresh threats. For instance, Madhya Pradesh announced a ‘one-time’ settlement scheme in 2025, promising to waive a portion of consumer electricity dues in exchange for partial lump sum payments. Announcements asked consumers to come forward early to benefit, while simultaneously promising to take strict action against those who continued to default. For such schemes to recover money, it is necessary that consumers believe the new threat.

involving a different interaction with the utility: SMS reminders and social comparisons, disconnections, and signed legal warnings, delivered either in-person by local linesmen or by registered post. These notices were pure threats: they noted the presence of substantial unpaid arrears, stated that the utility will punish non-payment, and mentioned the criminal laws being contravened by non-payment. Crucially, they were accompanied by no changes to the law, no increase in enforcement resources, no new incentives, and no change to the shared knowledge of dues held by the utility and consumers.

We find that when households receive legal notices by registered post from the head office, their arrears decline by 11.4 percent. Because registered post is relatively cheap, mailed notices generate a return on investment of 241 percent, even accounting for failed deliveries. Identical notices that were hand-delivered by the local utility linesman had an approximately zero effect on arrears, and we strongly reject equality between the two treatment effects, demonstrating the importance of the messenger in determining the effect of the threat. We discuss these results in more detail in Section 6.

To interpret these results we develop a simple model in Section 4 – households who owe money to the state utility update their beliefs about the credibility of an enforcement threat based on who the messenger is. When a threat is delivered by a messenger both the state and household know to be compromised – such as a utility linesman with a history of ignoring underpayment – the household infers that the threat is likely to be ‘empty’. When an identical threat is delivered in a manner that deliberately bypasses the compromised local agent, households are more likely to infer that the threat is credible. Households who reason in this manner will respond more strongly to threats delivered by post than to those delivered by the local linesman. If they additionally know that the state is capacity-constrained and must target enforcement action, they will also only partially pay down debt.

This model rationalizes the result of the first experiment but has additional implications that we test using a follow-up experiment on the same sample. If the choice of messenger changes consumer *beliefs* about state intent or credibility, then interactions in the past will

change how consumers respond to the same threat repeated in the future. Put differently, beliefs create path-dependence.

To test the dynamic predictions of our model, and thereby rule out alternative explanations, we ran a second experiment approximately a year later. In this second experiment, we randomly selected a subsample of consumers in each of Experiment 1’s treatment arms to receive a mailed legal notice, identically phrased to that used in one of the treatment arms in Experiment 1. This allows us to measure the effect of these notices, conditional on past interactions.

Consistent with the model, we find that consumers who had previously received a notice from a linesman did *not* change their arrears when receiving a similar notice by mail a year later. In contrast, households who had either never received a notice, or had received one by registered post, responded to the mailed threats in Experiment 2. In other words, exposure to a non-credible messenger affects both contemporaneous and future responses to threats. We also investigate what happens when the credibility of the linesman is increased. One of the arms of Experiment 1 does precisely this by asking linesmen to first deliver a notice and then follow up to carry out a disconnection (an experimentally induced increase in credibility). We find that in this case, the subsequent mailed threat in Experiment 2 *did* induce a payment response. This is consistent with our model where the identity of the messenger matters only insofar as their credibility is a signal about the state’s intent. We describe our results further in Section 7 and make the case that these dynamic relationships are hard to reconcile with alternative explanations such as salience, novelty of mailed notices, inherent dislike of in-person enforcement, or learning.

We make three main contributions to existing research. First, we highlight the importance of threats for revenue recovery. Economists have studied how to reduce tax delinquency through removing information asymmetries (Pomeranz 2015; Naritomi 2019; Kleven et al. 2011); reducing compliance costs (Okunogbe and Pouliquen 2022); increasing penalties and enforcement resources (Aparicio, Carrillo, and Shahe Emran 2011; Chalfin and McCrary

2017; Dwenger et al. 2016); changing staff incentives (Khan, Khwaja, and Olken 2016); and even changing tax rates (Bergeron, Tourek, and Weigel 2024). In contrast, threats need not involve changes in penalties, information, action sets, or incentives. Yet in settings with low state capacity, even empty threats can be critically important, determining the baseline compliance levels upon which we may layer incremental improvements. We show that mailed threats reduce arrears and this intervention appears robust, repeatable, and highly cost-effective. Mailed legal notices even outperform physical disconnections. They also outperform SMS reminders and nudges – which have little or no effect – contrasting with prior work that has used peer comparisons to accelerate payment of overdue taxes, improve reporting, and reduce electricity consumption (Hallsworth et al. 2017; Alm, Bloomquist, and McKee 2017; Sudarshan 2017).

Second, by documenting that the credibility of a threat’s messenger matters for its effectiveness, we build on prior work on the importance of agents. Qualitative political scientists have documented the complex ways in which “street-level bureaucrats” shape how formal laws and institutional structures translate into outcomes on the ground (Lipsky 1980; Lotta et al. 2022). Economists have also studied the choice of agents in influencing policy outcomes, but largely in settings outside revenue collection and with interventions where changing or bypassing agents creates substantive changes in information, action sets, or payoffs. Muralidharan and Sundararaman (2013) show that teachers on short-term contracts have different incentives and produce different outcomes than permanent staff. Muralidharan, Niehaus, and Sukhtankar (2016) and Banerjee et al. (2020) show that e-governance measures can improve how public money is spent by changing the information that is visible and the actions bureaucrats can take. Fellner, Sausgruber, and Traxler (2013), Hallsworth et al. (2017), and Neve et al. (2021) test enforcement messages and nudges in developed country settings but vary the information content not the messenger. A literature on agricultural extension and health workers (BenYishay and Mobarak 2019; Kondylis, Mueller, and Zhu 2017; Banerjee et al. 2021) finds that peer or community-based messengers can be more persuasive than

government agents. These are settings involving the adoption of beneficial technologies, not compliance with threats. Balán et al. (2022) demonstrate that people with local knowledge can serve as effective tax collectors because they can better target households who are likely to pay. In contrast, we show that when local linesmen are not allowed to choose their targets, using them as messengers weakens the threat’s effect.

Closer to this work, Ortega and Scartascini (2020) compare letters, emails, and in-person communication to delinquent taxpayers in Colombia. They find in-person visits are effective in inducing payments but at a high cost. In that study, visits involved new staff. In our case, the cheaper impersonal option is far more effective, because we examine a setting where the agent used as an in-person messenger (i.e. the utility linesman) already interacts with consumers, and may have low credibility. This is crucial context in developing countries with low baseline compliance, because adding new staff may be cost-prohibitive, while existing enforcers may have low credibility. Our experimental results suggest that raising the credibility of staff or bypassing them both make threats more effective, but the latter can be cheaper and easier. We also demonstrate that the initial choice of a messenger can impact the outcome of multiple interactions: a threat delivered by mail outperforms the same threat delivered by a low-credibility linesman *both* contemporaneously *and* dynamically.

Finally, we contribute to the literature on utility service delivery in developing countries. We demonstrate that ability to pay is not the sole driver of non-payment, but rather that perceived government credibility matters, building on prior work that has investigated social norms (Burgess et al. 2020) and liquidity constraints (Jack and Smith 2020). Perhaps the most important finding in this paper is that in a setting with low state capacity, identical threats can range from highly effective to completely useless depending on who delivers them.

2 Context and data

Struggling electric utilities are a major source of financial losses for developing country governments, as energy is typically supplied by the state rather than by the private firms common in rich countries (Besant-Jones et al. 2004). In low-and middle-income countries, loss-making and often bankrupt state-run utilities recover only 35 percent of the average cost of power, with consumer non-payment playing a key role (Burgess et al. 2020). In turn, these utilities are often forced to ration supply and reduce quality (McRae 2015; Burgess et al. 2020), with power shortages directly hindering growth (Allcott, Collard-Wexler, and O’Connell 2016).

In India, the location of this research, the accumulated losses of state electricity utilities were about 80 billion USD in 2020–2021, approximately 3 percent of national GDP. The year-on-year revenue shortfall increased by 8.2 billion USD from the previous year, equal to 12.4 percent of all personal income tax collected in that financial year (Anand et al. 2025). In Madhya Pradesh, India’s second-largest state by area and home to just under 90 million people, electricity revenue collection is a major challenge. The government supplies power through three large utilities overseen by a state-run holding company. In 2020–2021, aggregate technical and commercial electricity losses were approximately 41.5 percent: for every dollar’s worth of power consumed, the government recovered less than 60 cents. Since technical losses typically amount to about 5 percent, most of this shortfall was due to illegal connections and unpaid bills. These losses were over and above tariff subsidies which already lowered the price consumers paid to below the cost of supply.

Consumer debt and utility losses are periodically socialized. Madhya Pradesh ran targeted arrear waiver schemes in 2013, 2018, and 2020. In 2025, the government announced a one-time settlement (Samadhan Yojana) with near-universal eligibility at a cost of over 320 million dollars. Rampant non-payment thus leads to a substantial financial burden on

the state.³ This non-payment problem persists despite the fact that the utility can harshly punish non-paying consumers, including fines, disconnections, termination of supply, and criminal cases and imprisonment (Appendix A.3). The last of these requires the involvement of police and the justice system, not just utility staff.

We conduct our study in partnership with Madhya Pradesh’s state-run Central and East Zone utilities, working within the Hoshangabad and Narsinghpur service regions (“circles”), roughly coincident with the districts of the same name. Appendix Figure A.1 plots the region within India. These circles are geographic administrative units defined by the utility, and, as of January 2022, contained approximately 343,000 and 194,000 domestic consumers, respectively, and were selected by the state holding company for the study because of their substantial non-payment problems, relatively good administrative data, and no exposure to domestic insurgency-related violence that would have complicated fieldwork.

2.1 Data

Administrative data Our administrative data on billing and collections, obtained from the Hoshangabad and Narsinghpur circles, includes about half a million domestic metered consumers (Table 1). For each consumer in each month, the data record a range of variables including: units of electricity consumed (kWh), components of the bill, such as energy costs and fixed charges (INR), arrears (INR), tariff type, and sanctioned load. The utilities also provided us with a separate dataset on payments (i.e., collections).⁴ Finally, the administrative data include consumer addresses and telephone numbers, which the utility used to carry out treatments and the research team used to conduct surveys.

3. There are also frequent national schemes to waive or restructure utility debt, including the Financial Restructuring Package of 2012, the Ujwal DISCOM Assurance Yojana (UDAY) of 2015, and UDAY 2.0 in 2020.

4. Appendix B.1.1 describes how we assemble these administrative billing and payment records into a monthly panel and clean the data, including assigning billing periods to calendar months, dropping consumers above the 99th percentile of baseline arrears to remove large outliers, and adjusting for a one-off utility billing error by interpolating across neighboring months.

Table 1 provides summary statistics using five months of administrative data for all domestic metered consumers and our experimental sample. Rural consumers have mean arrears (1,421 INR) about 6.6 times their monthly bill (214 INR). In urban areas, bills are higher (698 INR) but arrears are roughly the same (1,431 INR, 2.1 times the monthly bill).⁵ Only a small minority of consumers pay their bill in full; in any given month, 63 (23) percent of billed consumers in rural (urban) areas make no payments whatsoever. Average collections (289 INR) are well below the average monthly bill (407 INR), with underpayment in both rural and urban areas. Our main outcome of interest is arrears, which represents net debt owed by the consumer. We discuss the experiment sample further in Section 5, but briefly note here that we deliberately selected – on the utility’s instruction – households with high arrears (above 1,200 INR).

Table 1: Summary of the administrative data

	Full population			Experimental sample		
	Rural	Urban	Overall	Rural	Urban	Overall
Consumption (kWh)	99.97 (45.80)	128.65 (123.72)	111.42 (86.98)	108.18 (43.64)	158.74 (98.25)	123.33 (69.02)
Monthly bill (INR)	214.37 (393.44)	698.21 (1,160.04)	407.43 (828.33)	242.19 (402.20)	909.21 (976.21)	442.10 (701.61)
Arrears (INR)	1,421.11 (3,195.62)	1,431.35 (4,833.08)	1,425.20 (3,931.64)	2,129.99 (2,173.26)	3,184.99 (4,255.00)	2,446.18 (2,994.47)
Collections (INR)	98.18 (322.75)	578.14 (1,033.84)	289.69 (737.79)	78.49 (301.81)	830.89 (1,025.16)	303.99 (705.38)
% of cons who do not pay	62.97 (48.29)	22.97 (42.07)	47.01 (49.91)	86.95 (33.68)	29.81 (45.75)	69.83 (45.90)
% of cons who pay partially	32.18 (46.72)	52.22 (49.95)	40.18 (49.03)	12.96 (33.59)	69.81 (45.91)	30.00 (45.83)
% of cons who pay in full	4.85 (21.47)	24.81 (43.19)	12.81 (33.42)	0.09 (3.01)	0.38 (6.14)	0.18 (4.20)
<i>Number of bills</i>	1,409,631	966,905	2,376,536	105,038	44,950	149,988
<i>Number of consumers</i>	284,016	197,618	481,634	21,008	8,992	30,000

Notes: This table presents summary statistics for key variables from the utility administrative data, for the period August 2021 - December 2021 (prior to the start of our experiments). The experimental sample is a random sample of consumers who were on a domestic tariff, not on a below poverty line tariff, and had arrears of at least 1,200 INR at the time of sampling. We report standard deviation in parentheses.

5. For consumers who may have benefited from arrear waivers, these numbers represent debt accumulated since the most recent write-off.

Surveys We conducted baseline and endline phone surveys within the experimental sample to measure socioeconomic characteristics and beliefs about the risks of non-payment. For a utility perspective, we also conducted a survey of junior engineers in local offices within our study region. Appendix B.2 describes these surveys in detail.

3 Motivating evidence

In this section, we use administrative and survey data to provide motivating evidence that there is a substantial non-payment problem in Madhya Pradesh. Non-payment may occur for different reasons: (i) transaction costs make payment difficult; (ii) consumers cannot afford to pay their bills; or (iii) many consumers simply choose not to pay, ignoring substantial penalties enshrined in the law. The data suggest that the last of these is an important contributor to the problem.


3.1 Arrears and payments in the status quo

In addition to the summary statistics in Table 1, the temporal characteristics of our administrative data provide useful insight into utility losses. To begin, it is clear that consumer debt is increasing over time. The left panel of Figure 1 shows average arrears for all domestic metered consumers from 2021 to 2023. Arrears rose steadily in this period, growing even more rapidly after 2022.

The right panel demonstrates that non-payment is widespread: between August and December 2021, the average consumer paid only 60 percent of the total bill amount owed. Notably, consumers do pay at least some of their bill, which argues against transaction costs preventing payments writ large. Moreover, payment is straightforward in Madhya Pradesh: consumers can pay in person, electronically, or via phone. Since electricity consumption rises with income in low- and middle-income countries (Gertler et al. 2016) and in Indian households specifically (Ahmad, Pachauri, and Creutzig 2017), the fact that payment rates

are well below 100 percent across all deciles of consumption casts doubt on ability-to-pay as a primary explanation for non-payment.

Figure 1: The utility payment problem in the status quo



Notes: This figure plots arrears and payment behavior using administrative data from the Madhya Pradesh utilities. The left panel shows the trend in arrears over time in both the full domestic metered consumer base (dashed line) and our experimental sample (solid line). The grey area marks the first field experiment as described in Section 5. The right panel plots the total amount paid as a share of the total billed amount net bill from August 2021 – December 2021 for each consumer, separately by deciles of electricity consumption, over the full sample of domestic metered consumers.

3.2 New consumers

We next examine how *new* consumers behave immediately after being connected to the grid. New consumers are a useful sample, because they are informative about how exposure to the utility over time affects behavior. We use our administrative data to isolate consumers who connected between 2018 and the end of 2023 and for whom no arrear waiver schemes were announced in their first twelve months after connection.⁶ We estimate how their propensity to make *any* payment evolves with the time since first connection using the following

6. Appendix B.3.1 provides more detail on how we identify these consumers in our data. We exclude consumers joining shortly before a waiver to avoid conflating normal payment behavior with anticipation or eligibility effects.

specification:

$$\mathbf{1}[\text{Pay anything}]_{it} = \sum_{k \neq 2} \beta_k \cdot \mathbf{1}(\text{Months since connection} = k)_{it} + \alpha_i + \gamma_t + \varepsilon_{it} \quad (1)$$

where $\mathbf{1}[\text{Pay anything}]_{it}$ is an indicator equal to 1 if consumer i made a payment in month t per the collections data, α_i and γ_t are individual and trimester-of-sample fixed effects, which allow us to identify the impact of time since connection while accounting for consumer-specific traits and general time trends in the utility’s ability to collect revenue. We cluster standard errors at the consumer level. The coefficients β_k trace out the trajectory of payment behavior as a function of time since connection.


We first estimate this equation using our full sample of newly-connected consumers. The left panel of Figure 2 shows the results. Although billing is monthly, consumers in this population do not pay every month. The figure plots the probability of payment relative to the second month since being connected, during which the average payment probability is 40 percent.⁷ The probability of making a payment in a month rises slightly for the first three months – consistent with consumers learning how to pay – flattens out, and then steadily declines. At the 12 month mark, the average payment probability is only 35 percent. This pattern is consistent with consumers learning over time that low payment rates are *de facto* permissible by the utility.

Next, we split our sample of new consumers by the median amount of electricity (in kWh) used per month, and separately estimate Equation 1 for each group. The right panel of Figure 2 presents the results. The mean probability of payment in the second month since connection is approximately 40 percent for both below-median energy users (likely poorer consumers) and above-median energy users (likely richer consumers). Both groups increase their likelihood of payment over the first several months, and then level off or decline. The

7. Although the choice of reference month does not affect the shape of the trajectory, we use the second month after connecting because many consumers may not receive a bill in the month they join and thus may not make a first payment until the month after.

decline in payment probability is sharper among the above-median consumers. This suggests that the population non-payment equilibrium is not solely driven by poverty. If anything, richer consumers with larger bills, and thus more to gain from not paying, are increasingly *less* likely to pay as time goes on. These patterns are consistent with a setting in which exposure to utility staff teaches consumers that strict enforcement is unlikely.

Figure 2: New consumer payment behavior



Notes: This figure presents estimates of the relationship between months since electricity connection and the probability of making a payment for our sample of newly-connected consumers over the first 12 months of an account. This sample of consumers, described in detail in Appendix B.3.1, joined the utility between 2018 and the end of 2023 and were not eligible for arrear waiver schemes. The left panel plots pooled estimates for all newly-connected consumers; the right panel provides separate estimates for consumers with below-median (light blue) vs. above-median (dark blue) electricity consumption. Estimates come from Equation (1), which includes consumer and trimester-of-sample fixed effects. The omitted category is two months since connection, which has a baseline payment rate of 0.4. We cluster standard errors at the consumer level; error bars (light gray) show 95 percent confidence intervals.

3.3 Survey evidence


Finally, we provide descriptive evidence from surveys that we conducted with consumers and utility staff.⁸ Figure 3 presents key summary statistics from these surveys. The left panel shows consumer attitudes, which reveals two key facts. First, many consumers believe that households have a right to electricity, whether or not they pay their bill, which suggests a norms or rights-based reason for non-payment consistent with Burgess et al. (2020). Second, the utility’s credibility among consumers is low. Only 62 percent of survey respondents believe that non-payment or late payment is likely to result in a disconnection notice. 40 percent of respondents believe that linesmen are understanding and will not take action if people cannot pay; and 45 percent of respondents believe that linesmen will accept side payments in exchange for ignoring their arrears. Finally, 47 percent of respondents believe that having contacts in the utility allows consumers to avoid disconnections.

The right panel reveals that utility officials themselves view their lack of credibility as an issue. We surveyed engineers in distribution centers (who supervise linesmen), asking them about reasons for persistent non-payment. Only a small minority said that households do not pay because service quality is poor (9 percent), or because they cannot afford to (11 percent), or because of a social norm that electricity is a right (14.5 percent). A majority (60 percent) reported that consumers do not pay their bill because there is no punishment. Together, these surveys demonstrate that both consumers and utility managers view the utility’s lack of credibility as an important contributor to the non-payment problem.

Overall, it seems clear that shortfalls in revenue reflect consumers choosing not to pay in spite of the fact that non-payment in theory carries severe penalties. Both consumers and the utility management perceive these to be empty threats. The challenge for a high-

8. Appendix B.2 describes these surveys, including the procedure we used to select the sample, in detail. Notably, the consumer surveys took place over the phone, and non-response was high and non-random. Consumers who responded to our survey on average use more electricity and have higher bills than the sample as a whole. Relevant to the non-payment problem, they pay a higher share of their bill on average and are substantially more likely to pay their bills (at least partially) than the overall experimental sample, suggesting that their answers with respect to utility credibility may be biased upwards relative to the overall population.

Figure 3: Survey evidence: Consumer and utility beliefs



Notes: This figure presents attitudes collected via survey. In the left panel, we plot survey responses from control group consumers. Each bar plots the share of consumers who agree with or view as likely the following statements: (Disconnection likely) If you paid only part of your bill or paid late, how likely would it be that you would receive a disconnection notice from the distribution company? (Linesmen lenient) Utility officials and linesmen are understanding and will not take action if people have difficulty paying their bills. (Linesmen corrupt) Linesman/utility officials are willing to accept payment in exchange for leaving people alone when they have payments pending. (Utility contacts helpful) Having contacts in government or in DISCOM allows people to avoid disconnections or get reconnected quickly. The former was measured at baseline, the latter two at endline only. In the right panel, we plot the share of utility staff who reported that consumers did not pay outstanding bills because (Service quality) service quality is poor, (Affordability) they could not afford to, (Social norm) a social norm that electricity is a right, and (No punishment) there is no punishment for late payments.

level policymaker is whether threats can be made more effective even when transforming the state’s capacity to enforce is infeasible or unaffordable. These features motivate our theoretical model and field experiments, which we describe in the sections that follow.

4 A model of payment in response to threats

In this section, we develop a simple model of how consumers make payment decisions when exposed to threats. The model serves to fix ideas, to guide the design of our experiments, and to develop several testable predictions which we take to the data; we do not structurally estimate the model, however.

Intuition When the state (i.e., the utility) presents delinquent consumers with an enforcement threat, they must decide whether to pay. This choice depends on whether they think the state intends to follow through on its threat. Consumers observe the state’s choice of a messenger and use this to make inferences about state intent. If a non-credible messenger (i.e., the local linesman) is tasked with delivering the threat, consumers are less likely to believe that the state is serious, as the linesman has a history of looking the other way. If instead, the state bypasses the compromised agent (e.g., uses registered post to send threats), households are more likely to believe that the state intends to enforce. This *perception* of threat credibility raises the likelihood of payment.

Conditional on choosing to pay something, consumers must decide how much. They recognize that regardless of state intent, universal enforcement is impossible, and therefore the state must use a targeting rule such as focusing on people with arrears above some cutoff amount. Therefore they pay enough to slip under the threshold in expectation – determined by their beliefs – but not all the way down to zero.

This type of model can explain payments conditional on messenger identity and partial payments, both of which are features of our data. By invoking beliefs about state intent, the model predicts that past signals influence the effectiveness of future threats. The model is a

description of consumer perceptions and decision-making, not a description of utility policy which in the paper is set by the experiment. The remainder of this section works through the model in more detail to deliver a set of predictions that we test using the experiments described in Section 5.

Setup A representative consumer owes an arrear balance $A_0 > 0$ to the state. A_0 is known to both the state and to the consumer. The state applies a deterministic targeting function $\tau : \mathbb{R}_+ \rightarrow \{0, 1\}$. The consumer views the state as being of type $\theta \in \{R, F\}$ having the following characteristics:

- Under R (rigid), if $\tau(A) = 1$, the state enforces against a household with arrears A , imposing a cost $C_R > 0$.
- Under F (flexible), the state does not enforce regardless of τ .⁹

The consumer has prior beliefs over (i) the state's type and (ii) the targeting function. They update priors about the state's type θ based on the identity of the messenger – either registered mail ($m = M$) or the local linesman ($m = L$). They update priors about the targeting function τ based on whether they receive a written threat, given their arrears A_0 .

Beliefs about state type Consumers believe that a rigid state, intent on follow-through, is more likely to bypass a low-credibility linesman and instead use mail to deliver its threats.¹⁰

The state's choice depends on its type:

$$\alpha_R \equiv \Pr(m = M \mid \theta = R), \quad \alpha_F \equiv \Pr(m = M \mid \theta = F),$$

9. We describe F as never enforcing for ease of exposition. Our results require only that a rigid state enforces more often than a flexible one.

10. This also implies that if the linesman were high-credibility, then a rigid state might not bypass them.

with $\alpha_R > \alpha_F$. The consumer observes m and Bayesian updates on θ as follows:

$$\pi_M = \Pr(\theta = R \mid m_0 = M) = \frac{\pi_0 \alpha_R}{\pi_0 \alpha_R + (1 - \pi_0) \alpha_F} > \pi_0, \quad (2)$$

$$\pi_L = \Pr(\theta = R \mid m_0 = L) = \frac{\pi_0 (1 - \alpha_R)}{\pi_0 (1 - \alpha_R) + (1 - \pi_0)(1 - \alpha_F)} < \pi_0. \quad (3)$$

Receiving a threat via mail shifts the consumer's posterior about θ toward R ; receiving a threat from the linesman shifts it toward F .

Beliefs about targeting function Receiving a threat is informative about τ regardless of state type, since both R and F send threats according to the same rule. If a consumer receives a threat, they can infer that the targeting rule τ ‘turns on’ at some level of arrears that is at least as high as their own level A_0 .

Consumer decision problem After observing a threat and the identity of its messenger, the consumer chooses payment amount $P \in [0, A_0]$ to minimize expected cost. If the state is type F , there is no enforcement, so the only cost is cash paid; under R , the cost is cash plus the expected penalty:

$$\text{EC}(P) = \underbrace{P}_{\text{cash paid}} + \underbrace{\pi \cdot \mathbb{E}[\tau(A_0 - P) \mid \text{threat at } A_0] \cdot C_R}_{\text{expected enforcement cost (only under } R)}. \quad (4)$$

The first-order condition for an interior solution is

$$1 = -\pi \cdot \frac{\partial}{\partial P} \mathbb{E}[\tau(A_0 - P) \mid \text{threat at } A_0] \cdot C_R. \quad (5)$$

The marginal benefit of a rupee paid is the marginal reduction in the consumer's expected probability of being targeted if it pays ($\mathbb{E}[\tau(A_0 - P) \mid \text{threat at } A_0]$), scaled by its posterior on the state being type R (π_L or π_M depending on whether the messenger is a linesman or mail) and the penalty a rigid state imposes (C_R). The consumer pays nothing

if the marginal benefit at $P = 0$ falls below the marginal (opportunity) cost of 1 rupee, and pays its full balance if the marginal benefit remains above 1 throughout $[0, A_0]$.

Threshold targeting rules A particular targeting function is a threshold rule, $\tau(A; \bar{A}) = \mathbf{1}\{A \geq \bar{A}\}$. Consumers know that the state uses a threshold rule, but are uncertain about the cutoff $\bar{A}$, with prior CDF G and continuous density g on $[0, \infty)$.¹¹ Receipt of a threat, by linesman or mail, is equivalent to observing $\bar{A} \leq A_0$ (regardless of state type). When they receive a threat, consumers will update their CDF, truncating the prior:

$$G_1(x | A_0) = \frac{G(\min(x, A_0))}{G(A_0)} \quad (\text{posterior CDF of } \bar{A} \text{ after threat at } A_0).$$

Characterization of optimal payment Conditional on $\theta = R$ (rigid state), the consumer's expected probability of facing enforcement actions after making a payment P is:

$$\mathbb{E}[\mathbf{1}\{A_0 - P \geq \bar{A}\} | \text{threat at } A_0] = \Pr(\bar{A} \leq A_0 - P | \bar{A} \leq A_0) = \frac{G(A_0 - P)}{G(A_0)}.$$

Substituting into (4), the consumer solves

$$\min_{P \in [0, A_0]} \left\{ P + \pi \cdot \frac{G(A_0 - P)}{G(A_0)} \cdot C_R \right\}. \quad (6)$$

The first-order condition for an interior optimum is

$$\frac{g(A_0 - P^*)}{G(A_0)} = \frac{1}{\pi C_R}. \quad (7)$$

Define $\kappa \equiv G(A_0)/(\pi_{m_0} C_R)$. The FOC reads $g(A_0 - P^*) = \kappa$. Note that κ increases as the consumer's perceived probability of the state being rigid, π , goes to zero. The consumer will pay nothing if $g(A_0) \leq \kappa$. This corresponds to a case where even the first rupee does

11. In Appendix C, we show that thresholds indeed predict who the utility targets for disconnections in the data, which is corroborated by surveys with utility officials.

not lower risk enough to be worth paying, given the consumer's priors on the threshold and priors on the probability that the state is rigid. Otherwise, a consumer will make some positive payment P^* , up to a maximum of A_0 , with P^* decreasing in κ . Intuitively, even if consumers believe the state is rigid, they will make *partial payments*, in order to reduce their post-payment arrears just below their best guess at the targeting threshold.¹²

Effect of messenger on payment Consider two consumers with identical primitives who receive the same threat through different messengers. Define $\kappa_M = G(A_0)/(\pi_M C_R)$ and $\kappa_L = G(A_0)/(\pi_L C_R)$. Since $\pi_M > \pi_L$, we have $\kappa_M < \kappa_L$. The consumer receiving the threat via mail will pay more. Indeed, if $g(A_0) < \kappa_L$ but $g(A_0) > \kappa_M$, then $P_L^* = 0$ while $P_M^* > 0$. The same threat produces zero average payment when delivered via linesman and a positive average payment when delivered via mail.

Dynamics and repeated threats As with an initial threat, a subsequent threat changes beliefs about both the state type and the targeting threshold. For example, suppose a second threat arrives by mail but the first threat was delivered by the linesman. The updating rule remains Equation 2, so for a consumer receiving a second threat, simply substitute π_L for π_0 . Each new threat causes the consumer to further update their priors. A threat delivered by linesman (mail) pushes π down (up), reducing (increasing) payment probability and quantity.

On the targeting threshold, there are two cases. *Case i*: For consumers who *paid* in the previous period, the new threat further truncates the probability distribution on $\bar{A}$, because it implies that $A_0 - P^*$, the arrears after the last payment, still lies above the threshold. This may result in an additional payment. *Case ii*: For households who *did not pay* in response to the last threat, the new letter carries no additional information about the

12. Some distributions G , such as the uniform distribution, have no interior optima and consumers will pay everything or nothing. Others have an interior risk maxima. We gloss over these possibilities for expositional simplicity because they are ultimately inconsistent with the fact that 40 percent of consumers make partial payments (see Table 1).

threshold. Non-payers may still make a payment in response to a new threat, but only if π rises sufficiently.¹³

Testable predictions

- (P1) **Messenger matters:** A mailed notice generates larger arrears reductions than an identical notice delivered by linesman: $P_M^* \geq P_L^*$, with the inequality possibly binding at the corner ($P_L^* = 0$ while $P_M^* > 0$).
- (P2) **Repeat responses to mail:** If a first mailed notice produces a partial reduction in arrears, then a second letter will produce a further response, as consumers who believe both threats to be credible update on the targeting rule.
- (P3) **History matters:** Consumers that initially (at $t = 0$) received a threat via the linesman will be less responsive to a subsequent mailed notice at $t = 1$ (sequence $L \rightarrow M$) than households receiving either two mailed threats ($M \rightarrow M$) or a first threat alone (M). This follows because after the initial threat, $\pi_L < \pi_M$. Thus, after the second threat, $\pi_{LM} < \pi_{MM}$. If π_{LM} remains below the payment threshold, the second mailed notice elicits no change in arrears from those who have seen linesman delivery first, even though the same threat to a previously-untreated household would have generated a reduction in arrears.
- (P4) **Credible history matters:** The more credible the linesman is, the more likely the consumer is to infer that the state is rigid if it receives a threat from a linesman. This is implicit in our assumption that $\alpha_R > \alpha_F$. If the linesman were highly credible, a rigid state would not need to bypass them, and $\alpha_R - \alpha_F \rightarrow 0$. Therefore, consider two consumers who both received an initial threat from a linesman and then face a

13. In a richer fully dynamic specification, the consumer would solve a stochastic control problem and might make small ongoing payments. We assume payment decisions are made at threat-receipt events. This captures the empirically observed pattern of payment clustering around notice receipt, and would also be reasonable if the utility can be expected to threaten shortly before enforcing. This model also does not endogenize the state's choice of (α_R, α_F) .

subsequent threat via mail. The arrears response to the second threat should be larger for the consumer whose initial threat was delivered by the more credible linesman.

The first prediction is the main policy-relevant outcome of interest. The last three involve dynamic relationships between past experiences and contemporaneous responses.

5 Experimental design and analysis

This section describes two linked field experiments, both carried out under a partnership with the Madhya Pradesh utilities. Appendix Figure A.2 presents the full experimental timeline: Experiment 1 treatments began in January 2022, and Experiment 2 treatments began in March 2023. In Experiment 1, we randomly assigned consumers to a series of enforcement interactions, including identical threats from different messengers, or to a control group, allowing us to test the main prediction (P1). In Experiment 2, we randomly assigned consumers who had faced enforcement actions in Experiment 1 to receive a new mailed threat, in order to test the dynamic implications (P2)–(P4).¹⁴

5.1 Experiment 1: Reminders, threats, and disconnections


5.1.1 Sample selection

We drew the experimental sample from the Hoshangabad and Narsinghpur utility service regions, described in Section 2.1 above. For the experiments, we restricted the sample to consumers on domestic (i.e., residential) tariffs.¹⁵ We excluded households on the “below poverty line” tariff in order to avoid targeting enforcement actions against an economically

14. Experiment 1 was preregistered but we did not write down the full model ex-ante and did not specify the sign of the messenger effect. Thus it would be appropriate to describe the model as *rationalizing* the result in Experiment 1, with additional implications that can be tested using Experiment 2, including a replication of the response to mailed threats from Experiment 1.

15. For logistical reasons, we excluded (i) consumers with broken or missing meters, for whom the utility could not generate accurate bills; (ii) consumers with missing contact details, who could not be contacted for treatments; and (iii) consumers who were inactive (zero consumption) in August 2021, the most recent month for which data were available at the time of sampling.

Figure 4: Experimental design diagram



Notes: Figure shows the design of Experiments 1 and 2. In Experiment 2, we take all treated consumers in Experiment 1 arms (B+C), D, E, and F and filter for arrears above 1,200 INR. We randomly assign 500 consumers from each of these four samples to receive a second mailed notice. The remainder serve as within-group controls.

vulnerable population. These low-income consumers are more likely than the average consumer to lack the ability – as opposed to willingness – to pay their power bills. At the direction of our government partners, we required that households have arrears at or above 1,200 INR – approximately the 75th percentile – at the time of sampling.

Our resulting sample thus comprises non-BPL domestic consumers in the top quartile of the arrears distribution. These households account for approximately 79 percent of total billed-but-uncollected domestic revenue and represent the margin where enforcement is most policy-relevant. The experimental sample columns of Table 1 show that, relative to the overall population mean, the experimental sample has higher consumption and arrears, but a substantially lower propensity to pay.

5.1.2 Treatment arms and randomization

From this eligible population, we drew a random sample of 30,000 domestic consumers (15,000 per circle), of whom we retained 2,500 as a ‘Business as Usual’ group to check for experimentation effects (we find no such evidence; see Appendix Figure F.1). We randomly assigned the remaining 27,500 households to one of six arms, stratifying on utility distribution center (DC), terciles of outstanding arrears at the time of sampling, and above vs. below-median consumption at the time of sampling.¹⁶

Figure 4 depicts the design of the first experiment, with its six arms (sample sizes indicate treatment *assignment*):

A. Control (N = 7,500): Consumers in the control group received no additional communication beyond the information on the electricity bill. Local utility staff were asked not to carry out disconnections on these households.

16. A DC is the smallest administrative unit in the utility and the local office at which billing, collections, and disconnections are carried out. There are 57 DCs in our sample. Arrears and consumption thresholds were defined as of August 2021.

B. SMS reminder (N = 5,000): The utility sent consumers three rounds of SMS messages in January 2022, which reminded consumers of outstanding payments. The text read (translated from Hindi): *“Dear Consumer, please pay your electricity bill on time. There is [X] percentage of your bill remaining on your electricity connection number [ID XXX]. If payment is pending, please pay immediately to avoid disconnection.”*

C. SMS with social comparison (N = 5,000): The utility sent consumers SMS messages containing the same information and following the same timing as Arm B, but with an added peer-group comparison:

“Dear Consumer, please pay your electricity bill on time. There is [X] percentage of your bill remaining on your electricity connection number [ID XXX]. Other similar consumers have [Y] percentage of their bill pending. If payment is pending, please pay immediately to avoid disconnection.”

The reference group consisted of other households in the same randomization stratum as the recipient, though this was not made explicit in the message. The SMS treatments constitute a low-cost, light-touch approach to enforcement that does not involve physical visits or disconnections, inspired by evidence on the impact of peer comparisons on energy consumption (Allcott and Rogers 2014) and the effect of simple reminders on tax payments (Neve et al. 2021). The introduction of peer comparisons in Arm C was a new addition that consumers would not have seen before. However, SMS or WhatsApp messages were used by the utility to send bills to consumers who had registered their mobile number. They are also used for reminders or alerts by third-party mobile wallets in widespread use in India.

D. Mailed legal notice (N = 5,000): Consumers in this arm were sent a formal notice by registered post in January 2022, personally signed by the Executive Engineer of the corresponding circle (a senior post). The notice warned consumers that non-payment of electricity bills constitutes a criminal offense, and that accumulated arrears render them liable for disconnection. It also provided the amount due and a payment deadline following

which further action would be taken. Appendix Section A.3 provides an image of one such notice (in Hindi), along with a translation in English.

A key novelty of this arm is the mode of delivery: mailed notices from the central office bypass ground-level utility intermediaries. In the status quo, local staff—linesmen or office employees serving several villages—are the primary channel through which the utility communicates with households, carries out disconnections, and occasionally collects payments. To our knowledge, the utility had not previously used registered post to deliver legal warnings of enforcement to households.¹⁷ Importantly, which households received mailed notices was not known to local officials, only to the research team.

In-person legal notice (N = 5,000): The utility delivered a formal signed notice *identical to that in Arm D*. Unlike in Arm D, however, the utility delivered these notices in person via a local linesman. Just like in Arm D, the notices were delivered to consumers in envelopes. We further randomized this group into two arms:

E. In-person legal notice only (N = 2,500): In Arm E (2,500 households), the utility staff were told to deliver the notice, but received no further instructions. Arm E thus delivers exactly the same message as Arm D, while varying only the messenger.

F. In-person legal notice with subsequent disconnection (N = 2,500): In Arm F, the utility followed its initial hand-delivered notice with an in-person visit to disconnect the household’s electricity connection if the consumer did not pay its arrears in full within 30–45 days of notice receipt.¹⁸ The split between Arms E and F was only communicated to utility staff *after* notice delivery. Therefore, at the point of notice distribution, these two arms were indistinguishable to both households and utility staff.

17. Indeed, the utility leadership was skeptical that mailed notices would affect payment behavior and was unwilling to fund the postage costs for this experiment. This skepticism is itself informative: the utility’s prior that mailed threats would be ineffective may have contributed to the status quo in which enforcement is routed exclusively through local agents.

18. The in-person delivery of all notices began in January 2022 and continued over several months. Disconnection visits began in March and continued through May 2022.

Because disconnection visits are otherwise rare, Arm F also creates a significant exogenous increase in linesman effort relative to the status quo.¹⁹ We interpret this as increasing linesman credibility, because – unlike in the status quo – in this treatment arm, the threat is followed by action.

Note that we do *not* randomize any of our experimental sample to the most aggressive forms of enforcement that the state has at its disposal: bringing criminal proceedings against delinquent consumers. This approach requires the involvement of the police and local courts.

The specter of criminal prosecution, even though not explicitly a part of our study, is useful in two ways. First, because the state can pursue this line of enforcement while bypassing the linesman entirely, it allows the state to differ from linesmen in credibility – a combination of intent to punish and ability to punish. Second, it provides an additional reason why households might respond to our treatments beyond the possibility of disconnection itself.

5.1.3 Testing model predictions

These six arms enable us to measure the impacts of enforcement actions ranging from low-cost and light-touch (SMS messages) to high-cost and high-intensity (in-person notice with physical disconnection). The model from Section 4 predicts that consumers should respond (weakly) more strongly to a mailed legal notice than to an identical threat delivered in-person by a linesman (P1). We evaluate this prediction by comparing the effect of treatment arms D and E, where identical threats are delivered by different messengers: one bypassing existing field agents, the other using them.

19. Prior to starting Experiment 1, we attempted to collect data on how frequently households were disconnected. Based on records of households who had paid a reconnection charge, we estimate that less than 4 percent of consumers with positive arrears were targeted over the course of an entire year.

5.2 Experiment 2: Repeat mailed notices

Predictions (P2), (P3), and (P4) in Section 4 all relate to the effect of *repeated* threats. To take these predictions to the data, approximately 14 months after Experiment 1, we conducted Experiment 2.

5.2.1 Sample selection

For Experiment 2, we are interested in understanding how past exposure to threats impacts responses to future threats. We therefore restrict the sample to consumers that were both successfully treated in Experiment 1 and whose arrears were at least 1,200 INR at the time of Experiment 2 sampling, leaving us with 10,803 consumers drawn from arms B, C, D, E, and F of Experiment 1.²⁰

5.2.2 Treatment arms and randomization

The bottom half of Figure 4 shows the Experiment 2 design: among the sample described above, for each Experiment 1 treatment group (pooling the two SMS groups), we randomly assign 500 consumers to receive a mailed legal notice, stratified as in Experiment 1. The remaining consumers in each arm serve as the control group; the utility received no instructions about these consumers. The utility mailed the 2,000 treated consumers a legal notice identical to that described in Experiment 1, Arm D, between March and April 2023.²¹ Experiment 2 thus generates four main treatment groups, which we refer to as SMS $\rightarrow$ Mail, Mail $\rightarrow$ Mail, In-person $\rightarrow$ Mail, and Disconnection $\rightarrow$ Mail.

20. The compliance condition corresponding to Arms B and C was a successful SMS delivery, for Arm D a confirmed mail notice delivery, and for Arms E and F a confirmed in-person notice delivery.

21. As in Experiment 1, notices were signed by Executive Engineers and sent via registered post. We used arrears data from January 2023.

5.2.3 Testing model predictions

This design allows us to test Predictions 2, 3, and 4 from the model in Section 4, and provides an opportunity to replicate the mailed notice treatment effect, as follows.

(P2): The model predicts that if a consumer responds to the first mailed threat by reducing arrears (but not all the way to zero), they should also respond to a second mailed threat. We evaluate this by comparing treated consumers (i.e., those who received a mailed notice in Experiment 1 *and* in Experiment 2) in the Mail $\rightarrow$ Mail arm to control consumers (i.e., those who received a mailed notice in Experiment 1 only) in the Mail $\rightarrow$ Mail arm.

(P3): The model predicts that if a consumer *first* receives an in-person threat, they will be less likely to respond to a subsequent mailed legal notice than either (i) a consumer who initially received neither an in-person visit nor a mailed legal notice, or (ii) a consumer who initially received a mailed legal notice. We evaluate this by comparing the treatment effect of an Experiment 2 mailed legal notice in the In-person $\rightarrow$ Mail arm to the treatment effect of an Experiment 2 mailed legal notice in the SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail arms.

(P4): The model predicts that if the initial in-person linesman were credible, then consumers could not draw strong conclusions about the state’s intent from their choice of messenger. This implies that a second mailed threat would be more effective following an in-person delivery from a credible agent than from a compromised one.

To evaluate this, we compare the treatment effects in the Disconnection $\rightarrow$ Mail arm, where the initial messenger was rendered more credible by experimentally assigning them to follow up on their notice, to those in the In-person $\rightarrow$ Mail arm, where the initial messenger confirmed the low-credibility prior because there was no follow-up.

(Replication): Finally, we evaluate whether the Experiment 1 mailed notice effect replicates by measuring treatment effects in the SMS $\rightarrow$ Mail group. Since households in this arm received only an extremely light-touch reminder (SMS message a year ago), if any effect of mailed legal notices in Experiment 1 is robust, we should see a similar reduction of arrears in this group.

5.3 Experimental integrity

Prespecification We preregistered Experiment 1 through the AEA RCT registry as AEARCTR-0008742. We broadly followed the PAP for Experiment 1. We did not preregister Experiment 2, which we added after observing outcomes from Experiment 1 to test some of the dynamic predictions implied by our model. We list the minor deviations from the PAP in Appendix Section G.

Balance Appendix Figure D.1 tests for balance in Experiment 1, using baseline administrative data. All groups are balanced on pre-treatment electricity consumption, monthly bill, arrears, and collections: differences between groups are close to zero in magnitude, and we cannot statistically reject that each treatment group is the same as the control.

Appendix Figure D.2 tests for balance in Experiment 2, using administrative data for the period after Experiment 1 but before the start of Experiment 2. Within the four groups (SMS $\rightarrow$ Mail, Mail $\rightarrow$ Mail, In-person $\rightarrow$ Mail, and Disconnection $\rightarrow$ Mail), the consumers randomized to receive an Experiment 2 mail notice are statistically and economically indistinguishable on bills, arrears, and collections from those randomized not to receive an Experiment 2 mail notice. In the Mail $\rightarrow$ Mail and SMS $\rightarrow$ Mail arms, we estimate statistical differences in electricity consumption between the treatment and control groups, but these differences are small in magnitude.²²

Compliance Compliance was imperfect in all Experiment 1 treatment arms, as shown in Appendix Figure D.3. In the SMS arms (B and C), non-compliance occurred when SMS messages bounced, likely due to incorrect or disabled phone numbers; the overall SMS compliance rate was 62 percent. In the mailed notice arm (Arm D), 61 percent of registered letters were successfully delivered. Failures were largely due to imperfect addresses in the administrative data (not uncommon, as the utility often delivered bills by hand and did not

22. The means are 111 kWh for treated and 107 kWh for control Mail $\rightarrow$ Mail consumers; 108 kWh for treated and 114 kWh for control SMS $\rightarrow$ Mail consumers.

require perfect address data for operations) or an inability to find someone to sign for letters. In the in-person arms (E and F), linesmen confirmed notice delivery via signed receipts. This occurred in 82 percent of treatment households in the two groups, with 62 percent of Arm F treatment households also receiving a follow up disconnection visit.²³

Compliance in Experiment 2 was significantly higher. The utility recorded approximately 79 percent of Experiment 2 mailed notices as successfully delivered. Delivery rates were similar across the four groups (79.4 percent for Disconnection $\rightarrow$ Mail, 77.8 percent for In-person $\rightarrow$ Mail, 80.6 percent for Mail $\rightarrow$ Mail, 79.6 percent for SMS $\rightarrow$ Mail).

5.4 Estimating equations

For both experiments, our primary objects of interest are intent-to-treat estimates which capture the policy-relevant effect of *assigning* a consumer to a treatment. Equation 8 presents our main specification, which we run separately for each treatment group (in Experiment 1, $\{B, C, D, E, F\}$; in Experiment 2, $\{\text{SMS} \rightarrow \text{Mail}, \text{Mail} \rightarrow \text{Mail}, \text{In-person} \rightarrow \text{Mail}, \text{Disconnection} \rightarrow \text{Mail}\}$):

$$Y_{it} = \beta_0 + \beta(T_i \times post_{it}) + \gamma_i + \delta_d \times \phi_t + \varepsilon_{it} \quad (8)$$

where Y_{it} is an outcome for consumer i in month-of-sample t , T_i is an indicator for assignment to treatment, $post_{it}$ is a dummy that takes a value of 1 for observations in months after treatment, γ_i are consumer fixed effects, $\delta_d \times \phi_t$ are DC-by-month-of-sample fixed effects, and ε_{it} is an error term, clustered at the consumer level. All regressions include the control group (in Experiment 1, Arm A; in Experiment 2, the control households for the relevant group). Our two primary outcomes are arrears (i.e., the cumulative outstanding balance) and collections (i.e., payments). For completeness, we also present results for energy con-

23. Since in-person delivery and disconnections have significant time costs, we tracked compliance for 2 months for notice delivery and 3 months for disconnections and cut off after this time, by which time compliance rates had flattened out; to the extent any visits occurred after this our estimates are a lower bound.

sumption.²⁴ β is the intent-to-treat effect of interest, which reflects the effect of assignment to treatment *averaged over* the six months following our interventions.

Given the imperfect compliance in our treatments, we also estimate local average treatment effects (LATEs) using the interaction of treatment receipt and the post dummy ($D_i \times post_{it}$) as the endogenous variable, instrumented by the interaction of treatment assignment and the post dummy ($T_i \times post_{it}$). We estimate a 2SLS specification that is otherwise identical to Equation (8):

$$D_i \times post_{it} = \pi_0 + \pi_1(T_i \times post_{it}) + \gamma_i + \delta_d \times \phi_t + \nu_{it} \quad (9)$$

$$Y_{it} = \beta_0 + \beta(\widehat{D_i \times post_{it}}) + \gamma_i + \delta_d \times \phi_t + \varepsilon_{it} \quad (10)$$

Finally, in order to trace out the time path of intent-to-treat effects, we estimate event studies for each treatment arm:

$$Y_{it} = \beta_0 + \sum_{\tau \neq -1} \beta_\tau T_i \times 1\{\text{Months to treatment} = \tau\}_{it} + \gamma_i + \delta_d \times \phi_t + \varepsilon_{it} \quad (11)$$

where τ indexes months relative to treatment, and all other terms are identical to Equation (8). The β_τ s are the intent-to-treat effects of interest, which reflect the average effect of assignment to treatment *in each of* the six months following our interventions. Therefore, the event study may reveal transient effects that do not have a substantial impact on the long-run mean estimated by Equation (8).

24. We use four months prior to treatment as the pre-treatment period and six months after treatment as the post-treatment period for the arrears outcome. For the collections and consumption outcomes, we use five months prior to treatment and five months after treatment. We treat these variables differently because consumption and collections appear on the contemporaneous bill, while arrears only appears on the bill with a lag. The $t + 1$ arrear value thus corresponds to the $t + 0$ collections and consumption values. Additionally, in Appendix Table F.2 we present treatment effects using an alternative payment variable and report similar results. This check is included because of concerns the collections data may miss cash payments or incorporate them with delays.

6 Experiment 1: Enforcement treatment effects

Table 2 provides ITTs and LATEs and Figure 5 presents event studies for Experiment 1.

SMS reminders (Arms B and C): The SMS messages produce precise null effects. Neither standard SMS messages nor social comparison SMS messages meaningfully reduce arrears or increase collections. We can reject ITT effects on arrears larger than 3.9 percent and 4.4 percent of the control group mean for reminder and nudge messages, respectively. Moreover, neither SMS impacts consumption. The event studies also rule out a transient response on arrears or collections. These results contrast with evidence on the effectiveness of reminders and nudges in other contexts (Allcott and Rogers 2014; Neve et al. 2021; Hallsworth et al. 2017): simple messages do not change behavior in this setting.

In our setting these results may not be surprising, especially for Arm B, because there are other sources of SMS reminders and alerts (utility sending bills to registered users, payment alerts from third-party apps etc). A marginal message may have little impact. In addition, SMS messages do not qualify as threats in the language of our model since they are automated, mass communications that carry no information about the targeting function and need not be read by a human being before being sent.

Mailed legal notice (Arm D): Mailed notices produce a sustained and statistically significant reduction in arrears. The ITT estimate is -177.3 INR ($p < 0.001$), corresponding to a 7 percent reduction in arrears relative to the control mean. The LATE for letter recipients, accounting for the delivery rate of 61 percent, is -291.8 INR ($p < 0.001$), an 11.4 percent reduction. On payments, the ITT (13.8 INR, $p = 0.173$) and LATE (22.6 INR, $p = 0.172$) are positive, but are insignificant. We do not observe an economically meaningful or statistically significant impact of mailed notices on electricity consumption.²⁵


25. Despite the effectiveness of mailed notices at reducing arrears, using our survey, we test for impacts of our treatments on financial well-being. We find no effects of our treatments on a series of economic outcomes (Appendix Table F.3), further implying that ability to pay is not the binding constraint on consumers' responses to threats.

Table 2: ITT and LATE estimates on electricity consumption, arrears, and collections (Experiment 1)

	ITT			LATE			N consumers
	Consumption	Arrears	Collections	Consumption	Arrears	Collections	
SMS $\times$ Post	-0.17 (0.80)	-26.27 (38.01)	-9.46 (10.08)	-0.28 (1.31)	-42.94 (62.12)	-15.46 (16.47)	12,370
SMS social $\times$ Post	0.73 (0.89)	-36.64 (39.27)	-14.38 (10.72)	1.15 (1.40)	-57.66 (61.81)	-22.64 (16.87)	12,369
Mailed notice $\times$ Post	-0.43 (0.85)	-177.28 (39.18)	13.75 (10.08)	-0.70 (1.39)	-291.84 (64.62)	22.63 (16.57)	12,379
In-person notice $\times$ Post	1.51 (1.12)	-46.93 (51.20)	-1.57 (12.61)	1.85 (1.37)	-57.37 (62.66)	-1.92 (15.42)	9,900
Disconnection $\times$ Post	-0.77 (1.29)	-0.16 (41.77)	-9.37 (12.20)	-1.25 (2.09)	-0.26 (67.51)	-15.15 (19.74)	9,895
Dep var mean (control)	119.93	2,559.87	231.73	119.93	2,559.87	231.73	

Notes: This table reports Experiment 1 ITT and LATE effects estimated using Equations (8) and (10), respectively. Each cell presents a separate regression, where the outcomes are electricity consumption in kWh, arrears in INR, and collections in INR. The estimation sample for consumption and payments is restricted to the event-time window $-5 \leq$ months to treatment ≤ 5 , whereas this window is transposed forward by one month for arrears (INR) and covers the $[-4,6]$ window given the lag with which arrears appear on the bill. Note that for arrears, the post dummy instead turns on if months to treatment ≥ 1 . The last column reports the number of consumers included in the estimation sample. All regressions include consumer fixed effects and $DC \times$ month-of-sample fixed effects, with standard errors clustered at the consumer level.

Figure 5: Experiment 1 event study



Notes: This figure plots event study estimates of treatment effects on arrears (left) and collections (right) for the mailed notice arm (top), in-person notice arm (middle), and disconnection arm (bottom), all relative to the control group. Hollow markers indicate the pre-treatment period, whereas filled markers indicate the post-treatment period. Note that the time periods we present differ between arrears and collections because of the lag with which arrears appear on the bill. We estimate effects using Equation (11). Vertical lines show 95 percent confidence intervals, computed from standard errors clustered at the consumer level.

The event studies provide further insight. First, the pre-treatment coefficients are close to zero and statistically insignificant, as expected given the randomization. Second, the effect of mailed notices on arrears emerges in the first month after notice delivery, grows larger over the next three months, and persists over the full six-month post-treatment window, with no evidence of attenuation. In contrast, collections do rise in the immediate aftermath of the mailed notice, but then return to the same level as the control. The six-month ITT effect on collections indicates that equilibrium payment behavior does not change much, consistent with a model where consumers pay off some debt once. Mailed threats are highly cost-effective. Appendix Table F.4 shows that letters have a return on investment of 241 percent, as they are cheap to deliver but generate meaningful reductions in arrears.

In-person legal notice (Arm E): The in-person notice arm, which delivers an identical legal notice via the local linesman rather than registered post, produces a near-zero effect on arrears (ITT: -46.9 INR, $p = 0.36$; LATE: -57.4 INR, $p = 0.36$). The point estimate is an order of magnitude smaller than the mailed notice effect and is not statistically distinguishable from zero. The effect on collections is similarly null (ITT: -1.57 INR, $p = 0.901$; LATE: -1.92 INR, $p = 0.901$). The event study confirms that in-person notices generate no discernible effect on arrears or collections in any post-treatment month. Linesman delivery is misleadingly cheap. From a budgetary perspective it seems nearly free since no new staff are involved and no postage is paid. Accounting for time costs, it is expensive (Appendix Table F.4).

Testing (P1): Messenger matters We check the first prediction of our model by statistically testing for differences in treatment effects on arrears between Arms D and E. To do this, we use all consumers in arms A (control), D (mailed notice), and E (in-person notice) and estimate a version of Equations (8) and (10) which includes treatment indicators for both arms. First, using these jointly-estimated coefficients, we conduct an F test for equal-

ity between the two ITTs and the two LATEs.²⁶ Second, again pooling Arms A, D, and E, we run a joint version of Equation (11), and conduct a test for joint equality of all six post-treatment event study coefficients. These two tests allow us to measure differences in both the average ITT and LATEs between the two groups and to test whether the treatments yield different trajectories (e.g., a short-run spike in payments, or divergences in arrears that build over time).

Panel A of Table 3 presents the results. The treatment effect of the mailed notice on arrears over the post-treatment period is substantially larger than the effect of the in-person notice (ITT difference: 132 INR, LATE difference: 238 INR), and we strongly reject equality (ITT: $p = 0.013$, LATE: $p = 0.001$). We also strongly reject that the trajectories of arrears are equal (ITT: $p = 0.015$, LATE: $p = 0.002$). This is the central finding of the paper: identical threats produce sharply different responses depending on who delivers them. When the notice arrives by registered post, it results in a sustained reduction in arrears. In contrast, when delivered by the linesman, it has no impact. This difference demonstrates the importance of messenger identity and confirms (P1) from our theory.

In-person notice with disconnection (Arm F): Arm F adds enforcement action to the threat: households who still had arrears after a linesman delivered a legal notice in person were visited a second time by the same linesman — this time to disconnect their electricity supply. Being assigned to this treatment did not lead to economically meaningful or statistically significant changes on average over the post-treatment period for arrears (-0.16 INR), collections (-9.37 INR), or electricity consumption (-0.77 kWh). The event study demonstrates that disconnections did not cause a transient change in arrears or payments.

Why do disconnections – in theory the most aggressive form of enforcement action we test – fail? First, not all consumers assigned to receive a disconnection visit were in fact visited: there was non-compliance for nearly 40% of households, in keeping with the notion

26. See Appendix E for more details. To recover our main ITTs and LATEs, we estimate effects separately for each treatment group for consistency with our estimation in Experiment 2. Jointly estimating treatment effects for all groups in Experiment 1 yields quantitatively similar results (Appendix Table F.1).

Table 3: Cross-group tests for equality for treatment effects on arrears

	ITT	LATE
<i>Panel A: Experiment 1: Mail vs In-person</i>		
Static (6-month average)	-132.22	-237.62
	[p = 0.013]	[p = 0.001]
Joint (per-period equality)	[p = 0.015]	[p = 0.002]
<i>Panel B: Experiment 2: Comparisons to In-person → Mail</i>		
<i>(i) Mail → Mail vs In-person → Mail</i>		
Static (6-month average)	-201.43	-252.41
	[p = 0.099]	[p = 0.103]
Joint (per-period equality)	[p = 0.056]	[p = 0.056]
<i>(ii) SMS → Mail vs In-person → Mail</i>		
Static (6-month average)	-345.82	-436.26
	[p = 0.018]	[p = 0.018]
Joint (per-period equality)	[p = 0.131]	[p = 0.132]
<i>Panel C: Experiment 2: Disconnection → Mail vs In-person → Mail</i>		
<i>(i) All Disconnection → Mail consumers</i>		
Static (6-month average)	-192.74	-245.37
	[p = 0.141]	[p = 0.142]
Joint (per-period equality)	[p = 0.074]	[p = 0.074]
<i>(ii) Corroborated Disconnection → Mail only</i>		
Static (6-month average)	-941.37	-1,134.10
	[p = 0.004]	[p = 0.004]
Joint (per-period equality)	[p = 0.001]	[p = 0.001]

Notes: This table reports tests of equality between treatment effects across experimental groups for monthly arrears. Each panel compares two groups; difference point estimates are the effect in the first group minus the effect in the second group. Appendix E describes our estimation procedure in detail. p -values correspond to tests of equality between the two treatment-arm effects shown in each row. The static rows compare effects averaged over a six-month post period; the dynamic rows report a joint test of equality of six post-treatment event study coefficients for each group. "Corroborated disconnections only" in Panel C(ii) restricts the Disconnection → Mail group to consumers whose Phase 1 disconnection was independently corroborated by the billing data (see Appendix B.3.2). All regressions include individual and DC-by-month-of-sample fixed effects. Experiment 2 comparisons additionally include group-by-month-of-sample fixed effects. Standard errors are clustered at the consumer level.

that disconnection visits were costly from the point of view of utility staff.²⁷ That said, our LATEs – which account for non-compliance – are small in magnitude and statistically indistinguishable from zero for arrears (-0.26 INR), collections (-15.15 INR), and electricity consumption (-1.25 kWh). Thus, non-compliance alone cannot explain the null result. Furthermore, because the experiment imposed quite stringent reporting and signed receipts to back up a claimed visit, we believe reported compliance rates are close to reality in the limited sense that linesmen who reported visiting consumers likely did so.

When linesmen report visiting they report either collecting a partial payment and not disconnecting, or doing the disconnection. We carried out a post-treatment debriefing with junior engineers and linesmen at distribution centers to rationalize these claims with the observation of no change in collections or consumption. We were able to do this in 13 DCs. We were given two explanations. First, households do make partial payments when visited but offset these by not making payments later in the month. Money is produced to avert the threat but payments are not additional. Second, households reconnect themselves.²⁸

We do not know if these explanations – which do not directly implicate the linesman as being compromised – are the whole story. Some linesmen may have visited consumers but neither collected money nor disconnected them. This could either reflect the electricity-as-a-right norm described in Burgess et al. (2020) which survey data suggests is common in our setting; or be the result of side-payments and bribes.

Overall, these results confirm that disconnections *as implemented in practice* do not generate meaningful returns for the distribution utility. Since disconnections are costly, they

27. The head office explicitly instructed linesmen to disconnect these households, and reminded them of this task over the course of two months. Nevertheless, many visits did not take place. Linesmen themselves reported having too many other things to do as a key constraint, and 91 percent of the junior engineers we surveyed said that disconnections would delay other tasks which include (as mentioned by respondents): maintenance, field work, revenue collection, meter replacement, transformer fitting, bill distribution, fault restoration, consumer complaints.

28. According to official utility rules, consumers must pay in full to avoid disconnection. In practice, in our utility survey, more than 70 percent of junior engineers reported that consumers would likely keep their electricity connection if they made only a partial payment. Out of 13 DCs, 11 junior engineers reported that households could reconnect themselves. Conversely, only one linesman said this, possibly because it would be illegal activity “on their watch.”

do not clear a cost-benefit test (Appendix Table F.4), in which case low enforcement rates can be rational. Norms or corruption may influence the linesman, but need not extend to senior management to explain why disconnections are rare.²⁹

7 Experiment 2: Dynamic impacts of enforcement

The results from Experiment 1 demonstrate that (P1) from Section 4 holds. To test (P2), (P3), and (P4), we turn to Experiment 2. Before doing so, we use the SMS $\rightarrow$ Mail group to demonstrate that the effect of the mailed notice from Experiment 1 – to reduce arrears – replicates. We present Experiment 2 ITTs and LATEs on arrears in Table 4, event studies in Figure 6, and tests for equality between treatment effects in Table 3 (above). Appendix Figure F.2 presents event study results on collections; these align with our arrear results across all groups.


29. Automated disconnection technology such as smart meters, which lower costs and remove linesman discretion, could in principle make this penalty effective. However, past work suggests installing smart meters that allow remote disconnections does not automatically turn into enforcement (Burgess et al. 2022).

Table 4: ITT and LATE estimates on electricity consumption, arrears, and collections (Experiment 2)

	ITT			LATE			N consumers
	Consumption	Arrears	Collections	Consumption	Arrears	Collections	
SMS $\rightarrow$ Mail $\times$ Post	-0.81 (1.54)	-294.16 (115.41)	53.81 (27.91)	-1.03 (1.93)	-370.57 (145.66)	67.79 (35.23)	4,680
Mail $\rightarrow$ Mail $\times$ Post	0.08 (1.41)	-142.21 (86.87)	50.46 (20.92)	0.11 (1.75)	-176.63 (108.08)	62.67 (25.97)	2,540
In-person $\rightarrow$ Mail $\times$ Post	2.24 (1.76)	48.80 (91.30)	-31.65 (25.39)	2.89 (2.26)	62.87 (117.63)	-40.78 (32.71)	1,698
Disconnection $\rightarrow$ Mail $\times$ Post	0.41 (1.95)	-136.09 (94.44)	22.74 (24.64)	0.52 (2.48)	-173.31 (120.25)	28.97 (31.38)	1,747

Notes: This table reports Experiment 2 ITT and LATE effects on administrative variables, estimated using Equations (8) and (10), respectively. Each cell presents a separate regression, where the outcomes are electricity consumption in kWh, arrears in INR, and collections in INR. The estimation sample for consumption and payments is restricted to the event-time window $-5 \leq$ months to treatment ≤ 5 , whereas this window is transposed forward by one month for arrears (INR) and covers the $[-4,6]$ window given the lag with which arrears appear on the bill. Note that for arrears, the post dummy instead turns on if months to treatment ≥ 1 . The last column reports the number of consumers included in the estimation sample. All regressions include consumer fixed effects, DC-by-month-of-sample fixed effects, and group-by-month-of-sample fixed effects, with standard errors clustered at the consumer level.

Figure 6: Experiment 2 event study



Notes: This figure presents event study estimates of treatment effects of mailed notices on arrears in Experiment 2. Each panel corresponds to one Experiment 2 group. For each panel, the estimation sample consists of only the relevant Experiment 2 group, including consumers assigned to treatment and control. For all event studies, we use month -1 as a reference category. Hollow markers indicate the pre-treatment period, whereas filled markers indicate the post-treatment period. We estimate effects using Equation (11). Vertical lines show 95 percent confidence intervals, computed from standard errors clustered at the consumer level.

Mailed notice replication (SMS $\rightarrow$ Mail): We first test whether the mailed notice treatment effects from Experiment 1 replicate by estimating the treatment effect of a mailed notice in the SMS $\rightarrow$ Mail group. Consumers that had previously received only SMS messages responded to the second-round mailed legal notice with an ITT effect on arrears of -294 INR ($p = 0.011$), and a LATE of -370.57 INR ($p = 0.011$). The event study shows a reduction in arrears that is sustained through the post-treatment period. If anything, these treatment effects are larger than what we found in Experiment 1, confirming that mailed legal notices do reduce arrears.

(P2): Repeat responses to mail (Mail $\rightarrow$ Mail): The model in Section 4 predicts that households who responded to a mailed notice the first time should reduce arrears again when they receive a second mailed notice, because these households believe that the state is rigid (hence their initial payment), and infer from the second letter that the arrear threshold is more stringent than they previously believed.

Treatment effects within the Mail $\rightarrow$ Mail arm confirm this. A second, repeated, mailed notice reduces arrears. Over six months, the average effects on arrears are -142.2 INR, $p = 0.085$ (ITT) and -176.6 INR, $p = 0.105$ (LATE). Comparing these point estimates to those from the SMS $\rightarrow$ Mail group suggests that the treatment effect of a second notice is smaller than that of first notice in this setting. The event study shows that the reduction in arrears persists over the post-treatment period.

Moreover, while we cannot identify household-specific treatment effects, in Appendix F.2.2 we show that the treatment effects of the second letter are likely driven by those who responded to the first notice.

(P3): History matters (In-person $\rightarrow$ Mail) In contrast to the effects of a first mailed notice (SMS $\rightarrow$ Mail) or a second mailed notice (Mail $\rightarrow$ Mail), choosing to use a low-credibility messenger for the first threat affects future interactions too. The model predicts that the treatment effect in the In-person $\rightarrow$ Mail group will be smaller than both the effect

in the SMS $\rightarrow$ Mail group, where there has been effectively no prior contact, and the Mail $\rightarrow$ Mail group, where prior contact has come from a credible source.

Consistent with this, in the In-person $\rightarrow$ Mail arm, where households were first given a notice by a linesman who took no further action (Arm E of Experiment 1), the second mailed threat does not reduce arrears over the post-treatment period (ITT: 49 INR, $p = 0.59$; LATE: 63 INR, $p = 0.59$).³⁰ In the event study, we observe no transient impact on arrears, in a pattern that is strikingly different from both SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail.

We formally compare treatment effects in the In-person $\rightarrow$ Mail group to treatment effects in the SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail groups. As with Experiment 1, we test for equality between ITTs, LATEs, and event study effects across groups. To do so, we construct our estimation sample by pooling all households (treatment and control) in the relevant Experiment 2 groups (in this case, SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail), and estimating a version of Equations (8), (10), and (11) that includes treatment indicators for both groups.³¹

In this cross-arm comparison of treatment effects, selection is a concern. While Experiment 2 randomizes *within* groups, the composition of consumers (treatment and control) in each group need not be identical. This is because the compliance filter applied to each in Experiment 1 precedes randomization into treatment and control for each arm in Experiment 2. This compliance filter differs by the first interaction: For Arms B and C (SMS interaction) it is a successful SMS; for Arm D (mailed notice) it is a successful postal delivery; for Arms E and F it is a successful in-person delivery.

In practice this does not seem to be a major concern. Appendix Table D.1 compares the population of consumers *across* groups in Experiment 2 over October 2022 - February 2023, a five month baseline prior to the second mailed notice treatment. The four groups are substantively very similar across variables in the administrative data though given the large sample size we can reject equality. For the key outcome of arrears, three of the four arms are

30. This treatment also has no impact on collections (ITT: -32 INR, $p = 0.21$, LATE: -41 INR, $p = 0.21$).

31. Because the control groups are not constructed identically across arms in Experiment 2, we also include group-by-month-of-sample fixed effects that flexibly control for any differences in levels and trends. See Appendix E for more details.

nearly identical, with the Mail $\rightarrow$ Mail somewhat lower, consistent with the Experiment 1 treatment effect persisting a year later.³² To mitigate against selection issues, when testing for equality of treatment effects across groups, we include individual fixed effects, group-by-month-of-sample fixed effects, and DC-by-month-of-sample fixed effects which are intended to address remaining concerns about static or time-varying cross-group differences.

Panel B of Table 3 presents the results. ITT (LATE) point estimates for Mail $\rightarrow$ Mail and SMS $\rightarrow$ Mail groups are respectively 201 (252) INR and 346 (436) INR more negative (i.e., larger in absolute value) than In-person $\rightarrow$ Mail. We reject equality on the static tests comparing the ITT estimates in these two cases ($p = 0.099$ and $p = 0.018$). We can also reject joint equality of the monthly coefficients comparing Mail $\rightarrow$ Mail to In-person $\rightarrow$ Mail ($p = 0.056$), though the joint test comparing monthly coefficients of SMS $\rightarrow$ Mail to In-person $\rightarrow$ Mail is less precise ($p = 0.131$). These results are thus consistent with (P3).

(P4): Credible history matters (Disconnection $\rightarrow$ Mail): Our final model prediction implies that if a linesman delivers the first threat then the more credible they appear, the more effective a subsequent mailed notice will be. The data suggest this is true, albeit noisily.

The point estimates on arrears from a mailed notice following a disconnection visit (Disconnection $\rightarrow$ Mail) are negative, though the estimates are noisy (ITT: -136 INR, $p = 0.15$; LATE: -173 INR, $p = 0.15$). The event study also reveals a pattern suggesting a reduction in arrears from this arm, though no individual point estimate is statistically distinguishable from zero.

Since the prediction involves a comparison across two states of the world (more vs. less credible linesman), we exploit the fact that the disconnection treatment (Arm F) in Experiment 1 serves as an exogenous shock to credibility relative to the in-person-notice-only treatment (Arm E). We therefore compare treatment effects in the Disconnection $\rightarrow$

32. For consumption, the groups range from a low of 108 kWh for the Mail $\rightarrow$ Mail group to 113 kWh in the SMS $\rightarrow$ Mail group; for monthly bill, the means range from 385 INR in the Disconnection $\rightarrow$ Mail group to 406 INR in the SMS $\rightarrow$ Mail group; for arrears the range goes from 5,361 INR for the Mail $\rightarrow$ Mail group to 5,895 INR in the In-person $\rightarrow$ Mail

Mail group to those in the In-person $\rightarrow$ Mail group. Panel C(i) of Table 3 reports the corresponding formal statistical tests. The point estimates imply that the Disconnection $\rightarrow$ Mail group has a much larger treatment effect than the In-person $\rightarrow$ Mail group over the post-treatment period, though the differences are noisy (ITT: -193 INR, $p = 0.14$; LATE: -245 INR, $p = 0.14$). We weakly reject equality of the trajectories of treatment effects ($p = 0.074$).

Since the Arm F treatment does not lead to reduced arrears in itself, the rationale for treating it as an exogenous shock to credibility rests on the assumption that some of the linesmen experimentally assigned to this group did collect spot payments and did disconnect consumers (i.e. behaved as credible enforcers) – even if these actions were subverted by consumers who reduced future payments or reconnected themselves.

As a robustness check however, we consider a subgroup of consumers where we observe more direct evidence of disconnections actually taking place. To do this, we restrict the sample to consumers where the billing data (not only the linesman’s own report) corroborate a disconnection having taken place, in the form of a spike consistent with a standard disconnect/reconnection fee (see Appendix B.3.2 for more details). We compare treatment effects from this subset of consumers in the Disconnection $\rightarrow$ Mail group to those in the In-person $\rightarrow$ Mail group in Panel C(ii) of Table 3. Using these more credible cases, we find a substantial difference between the two groups: the ITT (LATE) effect on arrears is 941 (1,134) INR larger in the Disconnection $\rightarrow$ Mail group, and the difference is highly statistically significant ($p = 0.004$ for both ITT and LATE). We also strongly reject equality of the trajectories ($p = 0.001$).

7.1 Alternative explanations

We propose that messenger credibility matters for the effectiveness of threats. The combination of results from Experiment 1 and Experiment 2 allows us to rule out a series of alternatives.

Salience and implementation failures One possibility is that mailed notices are simply more salient to consumers than in-person notices. If this were the chief explanation for the difference between the in-person and mailed notice effects, we would expect history to be irrelevant: the effect in the In-person $\rightarrow$ Mail group should be similar to that in both the SMS $\rightarrow$ Mail and the Mail $\rightarrow$ Mail groups. In practice, we find that both SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail deliver larger treatment effects than In-person $\rightarrow$ Mail: exposure to the non-credible linesman dampens the response to a future notice in a way that is not explained by salience alone.

Implementation failures also fall into this class of explanations. It is possible that linesmen did not actually deliver notices in the in-person arm, and instead falsified delivery receipts, leading to null treatment effects in Experiment 1. However, the fact that the treatment effect in In-person $\rightarrow$ Mail is smaller than that in SMS $\rightarrow$ Mail and Mail $\rightarrow$ Mail suggests that this cannot be the full explanation.

Other potential differences between mailed notices and in-person notices include concerns such as perceptions that mailed notices create a stronger legal record. These should not affect how consumers respond to future threats and are ruled out as being the sole explanation because of the difference between the In-person $\rightarrow$ Mail group vs. the SMS $\rightarrow$ Mail group and Mail $\rightarrow$ Mail groups.

Novelty of mailed notices The utility did not use registered mail to deliver notices in the status quo. Therefore, it is possible that the effectiveness of this treatment in Experiment 1 owes to its novelty, while in-person notices were routine and thus did not generate a response. If this were the case, we would expect all first-time recipients of a mailed notice in Experiment 2 to respond by reducing arrears. However, despite all being first-time recipients of mailed notices, treated households in the In-person $\rightarrow$ Mail group, the Disconnection $\rightarrow$ Mail group, and the SMS $\rightarrow$ Mail group do not have identical treatment effects in Experiment 2. Moreover, the Mail $\rightarrow$ Mail group reduces arrears in response to a *second* mailed notice,

while the In-person $\rightarrow$ Mail group does not respond to a *first* mailed notice, suggesting that the novelty of the delivery mechanism cannot entirely explain the effect.

Inherent dislike of linesmen or in-person enforcement It is possible that consumers did not respond to the in-person notice in Experiment 1 because they simply distrust linesmen generally or because they dislike in-person enforcement. If this were the full explanation, we would expect treatment effects in the In-person $\rightarrow$ Mail group and the Disconnection $\rightarrow$ Mail group to be similar, as both were exposed to in-person visits from linesmen. However, we see stronger impacts in the Disconnection $\rightarrow$ Mail group (especially among the corroborated disconnections subsample), suggesting that the credibility of the linesman is an important feature of the explanation.

Learning about enforcement It is possible that the different patterns we observe across groups in Experiment 2 are not driven by messenger credibility, but rather by learning. For example, suppose consumers who receive an in-person notice from the linesman but then see no follow-up (Experiment 1, Arm E) learn that the base rate of enforcement is low. Conversely, consumers who receive an in-person notice and a subsequent disconnection visit (Experiment 1, Arm F) learn that the base rate of enforcement is high. This could generate the pattern of results we observe in Experiment 2, where In-person $\rightarrow$ Mail is ineffective but Disconnection $\rightarrow$ Mail is somewhat more effective. This mechanism involves no messenger-specific inference about the credibility of threats.

However, this explanation is difficult to rationalize with the data. If consumers only responded to threats by updating their beliefs about the targeting rule – not the credibility of the messenger – we would expect to see symmetric responses across messenger types (i.e., Arms D and E) in Experiment 1 and symmetric attenuation in the effect of a mailed notice across messenger types (i.e., the In-person $\rightarrow$ Mail group and the Mail $\rightarrow$ Mail group).³³

33. The current model nests this alternative as the restricted case $\alpha_R = \alpha_F$. Under this restriction, the messenger channel is shut down: $\pi_M = \pi_L = \pi_0$ for any prior, and any updating about enforcement probability following letter receipt occurs entirely through the targeting-beliefs channel.

Instead, we reject equality in treatment effects across messenger types in both Experiment 1 and Experiment 2.

Enforcement can be subverted A remaining puzzle is why a threat should ever be effective, regardless of how it is delivered, if some combination of the consumer and linesman can render disconnections toothless (as Arm F suggests). One explanation is that the utility can bypass the linesman not only in their role as a messenger, but also in their role as an enforcer. This is achieved by filing a criminal case against the consumer, resulting in a First Information Report in a police station, and legal action in the local court. This is not unprecedented; indeed, in late 2025, the government of Madhya Pradesh announced they would build dedicated “energy police stations” to quickly use the police to crack down on consumers who do not pay their electricity bills (Gupta 2025), demonstrating a strong appetite for this approach.

8 Conclusion

To collect revenue, many governments rely heavily on threats that are backed by very limited enforcement capacity. Yet there is surprisingly little research in economics that examines which threats are more or less effective, why, and on whom.

This paper shows how the choice of whether or not to use government staff with low enforcement credibility as *messengers* can entirely determine whether a new threat reduces unpaid debt that citizens owe to the government. Our setting – the sale of state-supplied electricity to households – is of significant economic importance in many developing countries where the government runs huge markets in goods such as electricity, fertilizer, and cooking fuels. It is a convenient way to study threats, as distinct from information about non-compliance, because unpaid bills are common knowledge between citizen and state.

Working with state-run electric utilities in Madhya Pradesh, India, we document four main findings. First, when utility linesmen – the government’s default agents – deliver a

threatening legal notice to non-paying consumers, it has no impact on arrears or payments. Second, when the utility sends an *identical* notice through registered post, bypassing the linesman, consumers reduce arrears. Third, history matters: once a first notice is delivered by the linesman, a *subsequent* notice – even if mailed – is no longer effective. Fourth, the credibility of the messenger matters: when linesmen are experimentally assigned to exert higher enforcement effort, thus increasing their credibility, future mailed notices are more effective. Taken together, we argue that these results are consistent with a model where citizens use the identities of agents of the state to decide whether a threat is credible or not. Our intervention does not fully solve the payment shortfall, but scaling up mailed threats could deliver substantial reductions in debt, at a very high return on investment, by virtue of the fact that the population of defaulters is so large and postage costs quite low.

Our findings have particular relevance for developing country governments for whom revenue collection may depend heavily on how citizens respond to relatively empty threats, simply because their ability to carry out enforcement actions is hampered by low state capacity. Prior work has documented that bureaucrat discretion and decentralization can improve government effectiveness (Duflo et al. 2018; Balán et al. 2022; BenYishay and Mobarak 2019; Mookherjee 2015). We provide a reason why the opposite might be true, even absent explicit corruption. Choosing messengers or intermediaries may matter a great deal, such that we may conclude interventions fail when success might have been achieved simply by changing who implemented them or communicated information. In rolling out new policy, these considerations may be quite important.

We suggest that more research on why citizens make payments even when risk is low would be very useful. How should a state that has built up a reputation as a paper tiger (cheaply) reset expectations? Why do some people pay their bills in full, while their neighbors do not, even though both observe similar patterns of low enforcement? When and how can governments create inflated perceptions of risk? When are civic values most important?

These questions may be first-order determinants of revenue collection and, by extension, of our most basic indicators of state capacity.

References

- Ahmad, Sohail, Shonali Pachauri, and Felix Creutzig. 2017. “Synergies and trade-offs between energy-efficient urbanization and health.” *Environmental Research Letters* 12 (11): 114017.
- Allcott, Hunt, Allan Collard-Wexler, and Stephen D. O’Connell. 2016. “How do electricity shortages affect industry? Evidence from India.” *American Economic Review* 106 (3): 587–624.
- Allcott, Hunt, and Todd Rogers. 2014. “The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation.” *American Economic Review* 104 (10): 3003–37.
- Alm, James, Kim M. Bloomquist, and Michael McKee. 2017. “When you know your neighbour pays taxes: Information, peer effects and tax compliance.” *Fiscal Studies* 38 (4): 587–613.
- Anand, Abhishek, Varun Balotia, Praveen Ravi, Navneeraj Sharma, and Arvind Subramanian. 2025. *The electricity sector: A new perspective, new facts, and the need for alarm*. MIDS Working Paper 249. Chennai: Madras Institute of Development Studies, September.
- Aparicio, Gabriela, Paul E Carrillo, and M Shahe Emran. 2011. *Taxes, prisons, and CFOs: The effects of increased punishment on corporate tax compliance in Ecuador*. Technical report. Working paper 2011-02. George Washington University Institute for International Economic Policy.
- Balán, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L. Weigel. 2022. “Local elites as state capacity: How city chiefs use local information to increase tax compliance in the Democratic Republic of the Congo.” *American Economic Review* 112 (3): 762–797.
- Banerjee, Abhijit, Raghavendra Chattopadhyay, Esther Duflo, Daniel Keniston, and Nina Singh. 2021. “Improving police performance in Rajasthan, India: Experimental evidence on incentives, managerial autonomy, and training.” *American Economic Journal: Economic Policy* 13 (1): 36–66.
- Banerjee, Abhijit, Esther Duflo, Clement Imbert, Santhosh Mathew, and Rohini Pande. 2020. “E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in India.” *American Economic Journal: Applied Economics* 12 (4): 39–72.
- BenYishay, Ariel, and A. Mushfiq Mobarak. 2019. “Social learning and incentives for experimentation and communication.” *Review of Economic Studies* 86 (3): 976–1009.
- Bergeron, Augustin, Gabriel Tourek, and Jonathan L. Weigel. 2024. “The state capacity ceiling on tax rates: Evidence from randomized tax abatements in the DRC.” *Econometrica* 92 (4): 1163–1193.
- Besant-Jones, John, Clive Harris, Gary Stuggins, and Alan Townsend. 2004. *Public and private sector roles in the supply of electricity services*. Technical report. The World Bank Group, Energy and Mining Sector Board.
- Burgess, Robin, Michael Greenstone, Nicholas Ryan, and Anant Sudarshan. 2020. “The consequences of treating electricity as a right.” *Journal of Economic Perspectives* 34 (1): 145–169.
- . 2022. *The effect of smart metering on revenue collection: Evidence from an experiment in Haryana*. Working Paper. Applied Research Programme on Energy and Economic Growth (EEG).

- Chalfin, Aaron, and Justin McCrary. 2017. "Criminal deterrence: A review of the literature." *Journal of Economic Literature* 55 (1): 5–48.
- Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2018. "The value of regulatory discretion: Estimates from environmental inspections in India." *Econometrica* 86 (6): 2123–2160.
- Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke. 2016. "Extrinsic and intrinsic motivations for tax compliance: Evidence from a field experiment in Germany." *American Economic Journal: Economic Policy* 8 (3): 203–232.
- Fellner, Gerlinde, Rupert Sausgruber, and Christian Traxler. 2013. "Testing enforcement strategies in the field: Threat, moral appeal and social information." *Journal of the European Economic Association* 11 (3): 634–660.
- Gertler, Paul J., Ori Shelef, Catherine D. Wolfram, and Alan Fuchs. 2016. "The demand for energy-using assets among the world's rising middle classes." *American Economic Review* 106 (6): 1366–1401.
- Gupta, Anand. 2025. *Madhya Pradesh plans dedicated energy police stations to tackle power theft and recover dues*. EQ International. Accessed 4 June 2026.
- Hallsworth, Michael, John A. List, Robert D. Metcalfe, and Ivo Vlaev. 2017. "The behavioralist as tax collector: Using natural field experiments to enhance tax compliance." *Journal of Public Economics* 148:14–31.
- Jack, B Kelsey, and Grant Smith. 2020. "Charging ahead: Prepaid electricity metering in South Africa." *American Economic Journal: Applied Economics* 12 (2): 134–168.
- Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken. 2016. "Tax farming redux: Experimental evidence on performance pay for tax collectors." *The Quarterly Journal of Economics* 131 (1): 219–271.
- Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark." *Econometrica* 79 (3): 651–692.
- Kondylis, Florence, Valerie Mueller, and Jessica Zhu. 2017. "Seeing is believing? Evidence from an extension network experiment." *Journal of Development Economics* 125:1–20.
- Lipsky, Michael. 1980. *Street-level bureaucracy: Dilemmas of the individual in public services*. New York: Russell Sage Foundation.
- Lotta, Gabriela, Roberto Pires, Michael Hill, and Marie Ostergaard Møller. 2022. "Recontextualizing street-level bureaucracy in the developing world." *Public Administration and Development* 42 (1): 3–10.
- McRae, Shaun. 2015. "Infrastructure quality and the subsidy trap." *American Economic Review* 105 (1): 35–66.
- Mookherjee, Dilip. 2015. "Political decentralization." *Annual Review of Economics* 7:231–249.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. "Building state capacity: Evidence from biometric smartcards in India." *American Economic Review* 106 (10): 2895–2929.
- Muralidharan, Karthik, and Venkatesh Sundararaman. 2013. *Contract teachers: Experimental evidence from India*. NBER Working Paper 19440.

- Naritomi, Joana. 2019. “Consumers as tax auditors.” *American Economic Review* 109 (9): 3031–72.
- Neve, Jan-Emmanuel De, Clement Imbert, Johannes Spinnewijn, Teodora Tsankova, and Maarten Luts. 2021. “How to improve tax compliance? Evidence from population-wide experiments in Belgium.” *Journal of Political Economy* 129 (5): 1425–1463.
- Okunogbe, Oyebola, and Victor Pouliquen. 2022. “Technology, taxation, and corruption: Evidence from the introduction of electronic tax filing.” *American Economic Journal: Economic Policy* 14 (1): 341–372.
- Olken, Benjamin A. 2007. “Monitoring corruption: Evidence from a field experiment in Indonesia.” *Journal of Political Economy* 115 (2): 200–249.
- Ortega, Daniel, and Carlos Scartascini. 2020. “Don’t blame the messenger. The Delivery method of a message matters.” *Journal of Economic Behavior & Organization* 170:286–300.
- Pomeranz, Dina. 2015. “No taxation without information: Deterrence and self-enforcement in the value added tax.” *American Economic Review* 105 (8): 2539–69.
- Sudarshan, Anant. 2017. “Nudges in the marketplace: The response of household electricity consumption to information and monetary incentives.” *Journal of Economic Behavior & Organization* 134:320–335.

APPENDIX FOR:

PAYING FOR POWER

Fiona Burlig and Anant Sudarshan

Contents

A	Additional experiment details	A3
A.1	Study location within India	A3
A.2	Study timeline	A5
A.3	Legal notices and translation	A6
B	Data: Additional details	A9
B.1	Administrative data	A9
B.1.1	Data cleaning	A9
B.1.2	Computed payments	A9
B.2	Surveys	A11
B.2.1	Consumer surveys	A11
B.2.2	Utility office survey	A13
B.3	Subsamples	A14
B.3.1	Newly-connected consumers	A14
B.3.2	Corroborated disconnections	A14
C	Empirical evidence on the true targeting rule	A15
D	Experimental integrity	A17
D.1	Balance	A17
D.1.1	Experiment 1	A17
D.1.2	Experiment 2	A17
D.2	Compliance	A20
E	Formal tests of differences across treatment arms	A22
F	Additional results	A25
F.1	Experiment 1	A25
F.1.1	Jointly-estimated treatment effects	A25
F.1.2	Control vs. business-as-usual	A27
F.1.3	Computed payments	A27
F.1.4	Effects on economic well-being	A29
F.1.5	Return on investment by treatment arm	A31
F.2	Experiment 2	A32

F.2.1	Effects on collections	A32
F.2.2	Mail → Mail: Heterogeneity by Experiment 1 response	A34
G	Deviations from our pre-analysis plan	A35

A Additional experiment details


A.1 Study location within India

Figure A.1 maps the spatial distribution of the experiment sample. Each purple dot is a distribution transformer (DT), the smallest physical unit of the electricity grid, typically serving 50–100 consumers in a contiguous neighbourhood (we do not plot individual consumers for privacy reasons).³⁴

The sample is geographically dispersed across the catchment areas of the two (i.e., Central and Eastern Zone) utilities. From the map, it is clear that our experiment takes place in both dense urban or peri-urban areas where transformers are close to each other (which have more dots due to transformer density) as well as in rural parts of the state (which have fewer dots). This dispersion matters for two reasons. First, we are able to test the effect of messenger choice across heterogeneous local environments rather than in a single town. Second, the rural–urban mix is relevant for the cost calculations in Appendix F.1.5: linesman travel and time costs for in-person delivery or disconnections scale with the distance between the head office location and consumers. These costs may be lower in very densely populated cities.

34. Our administrative data contain a transformer ID for each consumer. We obtained geographic coordinates for each consumer by querying registered consumer addresses through the Google Maps API; we plot the median coordinate for each DT.

Figure A.1: Study location



Notes: This map shows our study location. The outer map shows India, with each state drawn in gray with a white outline. The light blue highlighted state is Madhya Pradesh. Our utilities serve Hoshangabad and Narsinghpur, highlighted in purple. Each dot on the inset map is a distribution transformer (DT), which serves approximately 50–100 consumers. We plot the DTs that serve the consumers in our experiment.

A.2 Study timeline

Figure A.2 reports the timing of our study activities. We carried out the sample selection for Experiment 1 in August 2021. We conducted a baseline phone survey in November–December 2021 (before any treatment). The utility sent the first round of SMS messages on Dec 30, 2021 – practically speaking, January 2022 was the first treatment month – and a second and third round about a month later. The utility delivered in-person notices in January and February 2022, and mailed notices in January–March 2022. Disconnection visits took place between March and May of 2022, because they followed an in-person visit. We surveyed utility officials (“enforcement survey”) between August and October 2022 and conducted an endline with households between December 2022 and January 2023. We conducted Experiment 2 in March and April 2023. Finally, we ran our disconnection targeting rule exercise (see Appendix C) in March – May of 2023.

Figure A.2: Study timeline

Activity	2021		2022												2023				
	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May
Baseline survey																			
Experiment 1																			
SMS																			
Social comparison																			
In-person notice																			
Mailed notice																			
Disconnections																			
Enforcement survey																			
Endline survey																			
Experiment 2																			
Utility disconnection targeting exercise																			

Notes: The figure shows the timeline of the experiment.

A.3 Legal notices and translation

Figure A.3 presents a sample legal notice, used in Arms D (mailed legal notice), E (in-person notice), and F (disconnection) of Experiment 1 and the treatment group in Experiment 2, in the original Hindi. Figure A.4 presents its English translation.

Figure A.3: Sample legal notice (Hindi)

कार्यालय कार्यपालन अभियंता म.प्र.म.क्षे.वि.वि कं.लि. संभाग Harda

क्रमांक/कार्य.अभि./.....

दिनांक/...../.....

सूचना

(अंतर्गत विद्युत अधिनियम 2003 की धारा 56)

प्रति,



सर्विस क्रमांक - [Redacted]

विषय:- ऊर्जा प्रभार का भुगतान न करने के कारण विद्युत प्रदाय बंद करने बाबत सूचना।

आपको सूचित किया जाता है कि आपके सर्विस क्रमांक [Redacted] के विरुद्ध बकाया राशि [Redacted] /- रुपये माह November, 2021 के अंत तक ऊर्जा प्रभार एवं अन्य प्रभार की राशि बकाया है। जिसका देयक आपको प्रेषित किया जा चुका है, साथ ही उक्त विद्युत कनेक्शन पर बकाया राशि भुगतान हेतु नोटिस भी जारी किया जा चुका है।

अतः आपको अवगत कराया जाता है कि यदि ऊर्जा प्रभार एवं अन्य प्रभार की बकाया राशि [Redacted] /- रु. इस सूचना के जारी होने की तिथि से 15 दिनों के अंदर भुगतान नहीं किया जाता है तो विद्युत अधिनियम 2003 की धारा 56 एवं विद्युत प्रदाय संहिता 2013 की कड़िका 9.13, 9.14, 9.15, 9.16 एवं 7.27 के अंतर्गत प्रदत्त शक्तियों का उपयोग करते हुये आपके प्रतिष्ठान/परिसर की विद्युत आपूर्ति सूचना अवधि समाप्त होने के पश्चात् विच्छेदित कर दी जायेगी।

कृपया ध्यान रखें कि विद्युत आपूर्ति विच्छेदित होने के पश्चात न्यूनतम मासिक प्रभार तब तक देय होगा जब तक सकल बकाया राशि का भुगतान नहीं कर दिया जाता, एवं विद्युत विच्छेदन प्रभार एवं संयोजन प्रभार समय-समय पर मण्डल/कम्पनी द्वारा निर्धारित दरों के अनुरूप देय होगा। विद्युत आपूर्ति विच्छेदित के बावजूद मण्डल/कम्पनी को कुल बकाया राशि वसूल करने का अधिकार सुरक्षित रहेगा, तथा विद्युत देनदारियों न चुकाने के कारण विद्युत संयोजन विच्छेदित कर देने के पश्चात् विद्युत कनेक्शन का अनाधिकृत तौर पर पुनः संयोजन कर लिये जाने की दशा में अधिनियम की धारा 138 के अंतर्गत दण्डनीय अपराध पंजीबद्ध होगा।

कार्यपालन अभियंता

म.प्र.म.क्षे.वि.वि कं.लि. संभाग Harda

Notes: This figure presents a sample legal notice, representative of those used in Arms D (mailed legal notice), E (in-person notice), and F (disconnection) of Experiment 1 and the treatment group in Experiment 2, in the original Hindi.

Figure A.4: Sample legal notice (English translation)

Office of the Executive Engineer
M.P. Madhya Kshetra Vidyut Vitaran Co. Ltd., Division: Harda

Number: / Exec. Engg. / **Date:** / /

NOTICE
(Under Section 56 of the Electricity Act, 2003)

To,

Name:

Address:

Contact: **Service Number:**

Subject: Notice regarding disconnection of electricity supply due to non-payment of energy charges.

You are hereby informed that against your **Service Number** [REDACTED] an outstanding amount of **Rs. 4890/-** is due as energy charges and other charges until the end of the month of **June, 2021**. The bill for this has already been sent to you, and a notice for the payment of the outstanding amount on the said electricity connection has also been issued previously.

Therefore, you are notified that if the outstanding energy and other charges of **Rs. 4890/-** are not paid within **15 days** from the date of issuance of this notice, then exercising the powers conferred under **Section 56 of the Electricity Act 2003** and **Clauses 9.13, 9.14, 9.15, 9.16, and 7.27 of the Electricity Supply Code 2013**, the electricity supply to your establishment/premises will be disconnected after the expiry of the notice period.

Please Note:

- Following the disconnection of the electricity supply, the minimum monthly charges will still be payable until the gross outstanding amount is fully paid.
- Disconnection and reconnection charges will be applicable as per the rates determined by the Board/Company from time to time.
- Despite the disconnection of the electricity supply, the Board/Company reserves the right to recover the total outstanding amount.
- If the electricity connection is unauthorizedly reconnected after being disconnected due to non-payment of dues, it will be registered as a punishable offense under **Section 138 of the Act**.

Executive Engineer
M.P. Madhya Kshetra Vidyut Vitaran Co. Ltd., Division: Harda

Notes: This figure presents an English translation of the legal notice we use in the experiment.

B Data: Additional details

B.1 Administrative data

B.1.1 Data cleaning

We apply minimal cleaning to the administrative billing data. We perform three notable cleaning steps.

First, we trim extreme arrears values using pre-period information. For each consumer, we compute average arrears over October 2021 to December 2021 and drop consumers whose pre-period average arrears lies above the 99th percentile of this distribution. This removes a small number of large outliers in baseline indebtedness. Note that Table 1 presents summary statistics *before* removing these consumers.

Second, we clean the billing data by setting negative bill values to missing. Since negative billed amounts are not logically valid in this setting, these observations are likely errors, so we remove them.

Finally, we correct two circle-specific billing errors. In the Hoshangabad circle, there was a widespread billing error for many consumers in January 2022. For each consumer, we impute the January 2022 bill using the average of recorded bills in the adjacent months, December 2021 and February 2022. We then replace the recorded January bill with this imputed value. Because arrears are mechanically affected by billing errors, we also compute the difference between the recorded and imputed January bill and subtract this amount from arrears in the subsequent months, February 2022 to April 2022. This yields a corrected series for both bills and arrears in the affected circle.

We apply an analogous correction in Narsinghpur for a separate billing error in March 2022. In this case, the affected month is imputed using the average of the same consumer's bills in the immediately preceding and following months. Unlike Hoshangabad, this billing error does not appear to contaminate the arrears variable, so only bills are corrected in Narsinghpur, while arrears are left unchanged. After implementing these corrections, we combine the corrected circle-specific variables into unified arrears and billing variables.

B.1.2 Computed payments

Collections data are recorded in a separate administrative database, and utility officials told us that these data are often incomplete and/or record cash payments with a delay. The collections data consequently do not always perfectly align with our data on arrears – which come from the main billing data and were thus described as more reliable by our utility partners. As a robustness check, we construct an alternative measure of payments – computed payments – using variables from the billing data alone. For each bill, the total amount due is

$$\text{Net bill}_{it} = \text{current bill}_{it} + \text{arrears}_{it} + \text{surcharges}_{it} + \text{government adjustments}_{it}.$$

We therefore construct computed payments as:

$$\text{Computed payments}_{it} = \text{Net bill}_{it} - \text{arrear}_{i,t+1}.$$

Intuitively, this is the amount the customer must have paid between two billing periods, after accounting for the new bill and the remaining unpaid balance. The computed payment measure is similar to the collections variable in the data.

Both variables have many zero or near-zero observations, consistent with the substantial non-payment problem in our empirical setting. When we focus on observations where both measures are at least INR 50, the two variables are very close: their median difference is zero, and 61 percent of observations are exactly equal. The mean of computed payments is higher than the mean of collections, but this is mainly because of a small number of large values in the upper tail. In this restricted sample, the top 1 percent of positive differences between computed payments and collections explains about 72 percent of the average gap between the two variables. Alternatively, when we restrict attention to months in which the consumer made a positive recorded payment (i.e., months with $\text{Collections}_{it} > 0$), the two measures are also very close. In this sample, the median difference between computed payments and collections is zero, while the mean difference is only about INR 35. Thus, in months where we observe a positive collection, computed payments and recorded collections are nearly identical on average.

These comparisons validate the use of the collections measure as our preferred outcome in the main text. We present robustness to using computed payments for Experiment 1 in Appendix Table F.2; results are very similar.

B.2 Surveys

B.2.1 Consumer surveys

Overview We conducted our baseline and endline surveys by phone by drawing a random sample of households from each circle and calling numbers down the list. Within our budget of maximum attempted calls (with two attempts per number), we completed 3,141 baseline surveys. At endline we first attempted to call consumers with complete baseline surveys. We were able to complete 2,122 such surveys and additionally completed a further 1,035 endline surveys for consumers who had not completed a baseline (3,157 endline surveys total).

We designed the baseline survey to collect information on electricity supply, correlates of income, and household financial indicators that include consumption expenditure, income, and credit. In the endline survey, we added a small number of questions related to beliefs about the credibility of enforcement.

Comparing survey respondents to the experiment sample The baseline and endline phone surveys were administered to random subsamples of the 30,000 experimental households, but a potential shortcoming of the survey is that the realized survey sample had a high non-response rate: 59.5 percent of numbers called were invalid, switched off, unreachable, engaged, did not pick up, or disconnected immediately. Table B.1 compares respondents to the full experiment sample on key administrative variables observed for all households.

Survey respondents consume more electricity on average than non-respondents (about 13 percent higher mean monthly consumption), carry slightly higher mean arrears, and pay a meaningfully higher share of their monthly bill. Households with active phone numbers and willingness to take a survey call may also be households with greater administrative attachment to the utility. We use the survey primarily to document attitudes about the non-payment problem (see Section 3.3). Given the composition of survey participants, the true problem is likely worse than what these households report. We use the survey to report the effect of our Experiment 1 treatments on households' financial well-being in Appendix F.1.4; the results are therefore not fully representative of the experimental population. These concerns do not apply to the administrative data, because we observe these outcomes for the universe of households in the experiment.

Table B.1: Selection into survey response on administrative covariates

	(1) Experimental sample	(2) Answered baseline survey	(3) Answered endline survey
Consumption (kWh)	123.27 (68.86)	130.68 (73.27) [p=0.000]	129.79 (75.26) [p=0.000]
Monthly bill (INR)	440.99 (699.09)	524.24 (748.71) [p=0.000]	527.02 (766.82) [p=0.000]
Arrears (INR)	2,444.93 (2,992.57)	2,491.75 (3,039.58) [p=0.361]	2,495.73 (2,822.87) [p=0.289]
Collections (INR)	303.25 (703.48)	421.25 (815.66) [p=0.000]	422.61 (781.78) [p=0.000]
Share of net bill paid	0.11 (0.21)	0.15 (0.23) [p=0.000]	0.16 (0.25) [p=0.000]
% of cons who do not pay	0.70 (0.46)	0.60 (0.49) [p=0.000]	0.60 (0.49) [p=0.000]
% of cons who pay partially	0.30 (0.46)	0.39 (0.49) [p=0.000]	0.40 (0.49) [p=0.000]
% of cons who pay in full	0.00 (0.04)	0.00 (0.06) [p=0.126]	0.00 (0.05) [p=0.565]
<i>Observations</i>	29,984	3,139	3,155

Notes: This table demonstrates how the consumers who answered our phone surveys compare to the overall experimental sample. Each cell reports the mean (with standard deviation in parentheses) of a variable from our administrative data over the pre-treatment window (August 2021 - December 2021). Column (1) presents the full experimental sample, Column (2) restricts to consumers who completed the baseline survey, and Column (3) restricts to consumers who completed the endline survey. We present p -values from a test of equality between the overall experimental sample and the relevant sub-sample in brackets.

B.2.2 Utility office survey

In addition to surveying households, we also interviewed one junior engineer from each of the 57 distribution centers. Alongside asking them many of the same belief-related questions included in the household survey, we asked a separate set of questions focused on how they make disconnection and reconnection decisions. For example, we asked why they think some consumers do not pay their electricity bills, whether they view electricity as a right, and how effective they believe disconnection notices are. We also asked about operational constraints, such as whether the utility has sufficient capacity to issue notices while carrying out its other functions on time.

B.3 Subsamples

B.3.1 Newly-connected consumers

To isolate a clean sample of genuinely new consumers for our motivating evidence in Section 3.2, we first restrict to the Narsinghpur administrative panel (as we lack utility connection dates for Hoshangabad) over the roughly 5-year period from 2018 to the end of 2023. Within this panel, we identify consumers whose first positive monthly bill appears within three months of their recorded connection date and who carry zero arrears on that first bill. This screens out accounts whose “new connection” entry in the dataset reflects an administrative migration or reconnection rather than a true start of service from a new consumer. We further keep only consumers who were on domestic metered tariffs throughout the twelve-month analysis window, and drop experimental participants so enforcement treatments do not confound the analysis. We also drop any consumer enrolled in the 2020/2022 Samadhan arrear waiver scheme, to avoid confounding the relationship between payments and time from connection with the impact of the arrear waiver. Finally, we require that each consumer’s twelve-month window ends before September 2023, when a statewide arrear waiver was applied – this rules out both the mechanical write-down and any anticipation effects in the run-up. After these restrictions we are left with 5,013 consumers, each of whom we observe for up to twelve months following their connection.

B.3.2 Corroborated disconnections

We isolate the subgroup of consumers who are the most likely to have had their electricity disconnected in Experiment 1 for our discussion of (P4) in Section 7. We rely on the fact that the utility charges a 200 INR disconnection/reconnection fee. We therefore identify consumers with “corroborated disconnections” as those for whom the linesman recorded a disconnection *and* we observe a bill at least 180 INR higher than (the consumer’s own) pre-disconnection visit average in the post-treatment period. We use 180 INR rather than 200 to account for natural fluctuations in bills over time. This bill spike serves as an independent check that a disconnection actually occurred. Note that being disconnected by the utility does not preclude consumers from reconnecting themselves, but does suggest that the linesman did their job. Overall, approximately 22 percent of consumers that the linesman reported disconnecting (237 out of 1060) are labeled as corroborated disconnections under this definition.

C Empirical evidence on the true targeting rule

In our model (Section 4), consumers believe that the utility uses a threshold rule for disconnection targeting. Here, we provide empirical evidence that corroborates this setup.

To start, although experiment subjects cannot know this directly, it is a fact that the utility enforced a threshold targeting rule on the *experiment* by requiring that the sample only include households with arrears above 1,200 INR – even for treatment arms involving only a written notice or an SMS reminder.

More generally, in order to understand how the utility field staff would target new disconnections, we restricted the universe of consumers to those who met the experiment’s inclusion criteria, and drew a new sample of 1,701 consumers at random (excluding those in the original experiment) in March 2023. We drew these consumers using the same strata as before, thus distributed over the 57 field offices (DCs) in the two circles of the experiment. Utility leadership provided each DC with the subset of this longlist that fell in their jurisdiction, and instructed them to select 25 percent of this list for disconnections.³⁵ The utility did not provide instructions on how each DC should choose who to disconnect. The DCs were asked to submit the consumer IDs of those that they intended to disconnect to the research team and the utility head office.

We compare the set of consumers that the DCs slated for disconnection against the overall longlist in order to learn about how the utility decides who to disconnect. We train a LASSO model using covariates in the administrative data to see what predicts this selection. We provided the LASSO with data on the monthly bill; arrears (in thousands); the square of arrears; a series of indicators for arrears being above round numbers in the three months before selection (i.e., 1,000, 2,000, ..., 10,000); the share of monthly bill paid over the 3 months, 6 months, 9 months, and 12 months before the selection month; changes in arrears over the last 3 months; and the number of months with zero payments in the last 3 months, 6 months, 9 months, and 12 months. We calibrate the LASSO to choose the three most important variables in explaining the utility’s disconnection decision. It selects two threshold indicators (cutoffs for arrears being above 6,000 and above 10,000), and arrears (in thousands of INR) as the most important explanatory variables. Appendix Table C.1 presents the results of a linear probability model with disconnection choice on the left-hand side and the LASSO-chosen covariates on the right-hand side. These results are consistent with the utility using arrear thresholds to target households for disconnections. An important caveat is that the selection for this exercise required choosing 25 percent of a candidate list to disconnect, which is higher than the share of the population with outstanding arrears that is actually targeted.

Moreover, we find evidence of threshold rules in our survey of field offices. When we asked utility officials how consumers were selected for disconnections, the most common responses in a multiple choice question were High Arrears (79 percent), whether the last bill was paid (65 percent) and the number of months since payment (68 percent). For those who mentioned arrears as relevant, we asked them to list a cutoff above which they would disconnect a consumer. The median reported cutoff was 5,000 INR. Taken together,

35. This implied delivering a disconnection notice to selected households, followed by cutting power to those who did not clear arrears.

Table C.1: Administrative variables in utility targeting rule: Post-LASSO Estimates

Variable	= 1[Selected for disconnection]
Arrears (INR 1,000s)	0.006
= 1[Arrears>6k]	0.147
= 1[Arrears>10k]	0.327
Number of consumers	1,701
R-squared	0.180
Number of post-LASSO variables	3

Notes: This table reports OLS estimates from a linear probability model where the dependent variable equals one if the utility selected a consumer for disconnection from the longlist. We measure arrears (in thousands of INR) at the time of selection. The arrears-threshold variable equals one if arrears exceeded 6,000 (10,000) INR in the selection month or the prior three months. We present robust standard errors in parentheses.

these results suggest that the utility does use arrear thresholds to target consumers for disconnections, consistent with the framework we lay out in Section 4.

D Experimental integrity


D.1 Balance

D.1.1 Experiment 1

Appendix Figure D.1 reports balance in Experiment 1 between each treatment arm and the control group on a battery of pre-treatment administrative variables, measured during the August 2021 - December 2021 window.

For each covariate X_{ia} for consumer i in arm a , we record the mean value in each arm. Numbers on top of each bar represent the p -value arising from a t -test for equality of the corresponding treatment arm with the control. Across covariates and across arms, differences are small and the corresponding p -values are not near zero, consistent with successful random assignment.

Figure D.1: Experimental balance (Experiment 1)




Notes: The figure above depicts balance between the control and various treatment arms for the pre-treatment period from August 2021 - December 2021. The value inside the bar represents the mean within each arm, and the value above the bar is the p -value of the two-sided t -test comparing each arm to the control group.

D.1.2 Experiment 2

Within-group balance Appendix Figure D.2 reports balance in Experiment 2, where we randomize consumers to receive a mailed legal notice or to control *within* a group (e.g., SMS $\rightarrow$ Mail). We present balance on administrative data variables spanning October 2022 – February 2023, after the conclusion of Experiment 1 but prior to the start of Experiment 2. For each group, we present the control mean of each covariate in the hollow bar, and

the treatment mean in the solid bar. We present p -values on the test for equality between treatment and control above the treatment bars. We generally fail to reject equality between treatment and control, and in the cases we do, the differences in the administrative variables are small in magnitude, consistent with successful random assignment.

Figure D.2: Experimental balance (Experiment 2)



Notes: This figure presents balance within Experiment 2, measured on administrative data variables for the period October 2022 – February 2023. The values inside the solid bars represent the mean of the experimental arms assigned to treatment, whereas the values inside the white bars are the control group means. The value above the bar is the p -value of the two-sided t -test comparing the mean of the treatment group with that of the corresponding control group.

Across-group balance In Experiment 2, we construct four groups: SMS $\rightarrow$ Mail, Mail $\rightarrow$ Mail, In-person $\rightarrow$ Mail, and Disconnection $\rightarrow$ Mail. Within each group, we then randomize consumers to treatment (i.e., a mailed legal notice) or control. To construct each group, we first filter each of the Experiment 1 treatment arms to select compliers, and then restrict to those with arrears above 1200 INR. Here, we consider comparability across each group. Appendix Table D.1 presents a test for cross-group balance in each circle separately as well as pooled.³⁶

36. One oddity here is that the spread in monthly bills is much larger than consumption and arrears, especially in Narsinghpur. At first glance, this is surprising since bills are based on consumption but the divergence here arises for two reasons: some households hover around the lifeline tariff where a bill discontinuity exists, and the top 1 percent of bills is not statistically similar across arms during the baseline period.

Table D.1: Cross-group balance in Experiment 2

	SMS → Mail	Mail → Mail	In-person → Mail	Disconnection → Mail	Equality p-value
<i>Narsinghpur</i>					
Consumption (kWh)	93.97 (37.02)	91.95 (26.30)	92.24 (26.40)	92.96 (28.55)	0.243
Monthly bill (INR)	423.03 (309.59)	394.03 (187.91)	402.70 (200.44)	410.94 (225.89)	0.006
Arrears (INR)	4,479.85 (2,382.20)	4,431.02 (1,832.22)	4,587.43 (2,338.20)	4,489.79 (1,820.06)	0.401
<i>Number of consumers</i>	2,059	1,493	866	880	
<i>Hoshangabad</i>					
Consumption (kWh)	128.15 (41.73)	129.92 (37.32)	128.04 (38.78)	125.48 (37.65)	0.114
Monthly bill (INR)	392.00 (387.70)	384.70 (358.21)	377.09 (356.88)	357.97 (321.79)	0.120
Arrears (INR)	6,793.76 (4,175.43)	6,685.18 (3,923.10)	7,254.81 (4,655.80)	7,075.84 (4,301.08)	0.009
<i>Number of consumers</i>	2,621	1,049	833	869	
<i>Both Circles</i>					
Consumption (kWh)	113.11 (43.20)	107.62 (36.47)	109.79 (37.58)	109.12 (37.12)	0.336
Monthly bill (INR)	405.65 (355.75)	390.18 (271.43)	390.14 (288.16)	384.62 (278.89)	0.013
Arrears (INR)	5,775.74 (3,684.81)	5,361.24 (3,090.45)	5,895.22 (3,896.85)	5,774.68 (3,539.03)	0.003
<i>Number of consumers</i>	4,680	2,542	1,699	1,749	

Notes: This table records sample means of administrative variables across the four groups in Experiment 2. The sample for each group consists of Experiment 1 compliers (consumers who received treatment in Experiment 1) with January 2023 arrears $\geq 1,200$ INR. SMS notice and SMS-social are pooled into one arm. Cell entries are consumer-level means with standard deviations in parentheses, computed over October 2022 – February 2023, prior to the start of Experiment 2. Equality p -values are from F -tests of equality across the four groups after accounting for circle fixed effects.


Within both circles, the four groups are almost identical on consumption. We can reject joint equality of monthly bills but the differences are substantively small. In the data they appear to be driven by a few outlying high bills. Additionally in Narsinghpur many consumers have usage which hovers around the lifeline level below which bills are flat by design (marginal price zero). Arrears are nearly identical across three of the four arms (SMS $\rightarrow$ Mail, In-person $\rightarrow$ Mail, Disconnection $\rightarrow$ Mail), all of whom had insignificant treatment effects in the first experiment. Arrears for the Mail $\rightarrow$ Mail arm are lower, consistent with the treatment effect of the Experiment 1 legal notice persisting more than a year later. In the cross-group tests we report in Table 3, we control for individual fixed effects, DC-by-month-of-sample fixed effects, and group-by-month-of-sample fixed effects, which addresses any differences across groups, in levels or trends.

D.2 Compliance

Appendix Figure D.3 presents compliance for Experiment 1. By the end of the study, 62 percent of Arm B and C consumers successfully received an SMS delivery. Mailed legal notices (Arm D) were dispatched by registered post towards the end of January 2022. However, delivery occurred over the next few weeks and we received reports of successful deliveries for about 3-4 weeks thereafter. We were able to confirm that 61 percent of Arm D households successfully received a mailed notice.

In-person notices (Arms E and F) were delivered by linesmen with instructions given the same month (January 2022). However, deliveries only happened gradually – compliance kept rising until March before tapering off at 82 percent. Disconnection visits for Arm F households occurred substantially later, as they had to be preceded by an in-person notice delivery and a waiting period to allow consumers to pay down their arrears. Given the slow roll-out of notices, utility senior management only gave instructions to carry out this treatment in March 2022. As with notices, these disconnections occurred gradually, rising to about 62 percent compliance.

Figure D.3: Compliance (Experiment 1)



Notes: The figure shows the cumulative share of consumers treated by date for each intervention arm. Treatment rollout was fastest for Arms B and C (SMS and SMS social comparisons) and Arm E (in-person notices), while Arm D (mailed notices) and especially Arm F (disconnections) were implemented more gradually over time. The line for the disconnection treatment starts after the others because disconnection follows the delivery of in-person notices.

E Formal tests of differences across treatment arms

The headline ITT and LATE tables in this paper estimate the effect of each treatment group separately. However, several claims in the paper compare treatment effects on arrears across groups. In this section, we describe our procedure for formally conducting the following comparisons, whose results we present in Table 3:

- (A) **Experiment 1:** Arm D (Mailed legal notices) vs. Arm E (in-person notices)
- (B.i) **Experiment 2:** Mail $\rightarrow$ Mail vs. In-person $\rightarrow$ Mail
- (B.ii) **Experiment 2:** SMS $\rightarrow$ Mail vs. In-person $\rightarrow$ Mail
- (C.i) **Experiment 2:** Disconnection $\rightarrow$ Mail vs. In-person $\rightarrow$ Mail
- (C.ii) **Experiment 2:** Corroborated disconnection $\rightarrow$ Mail vs. In-person $\rightarrow$ Mail

For each comparison, we conduct two tests for both the ITT and LATE estimates.

1. Comparing static post-period means. For Experiment 1, this test pools the two treatment groups being compared, along with the relevant control households, in a single regression as in Equation E.1 with monthly arrears as the outcome. We restrict the sample for groups X and Y to the four months before treatment, the treatment month, and the six months after treatment, noting that these months may differ across group. The control group spans four months before the start of the earlier treatment through six months after the later treatment.

For Experiment 2, this test pools the treatment and control households from the two groups being compared in a single regression as in Equation E.1 with monthly arrears as the outcome. We restrict the sample for groups X and Y to the four months before treatment, the treatment month, and the six months after treatment, noting that these months may differ across group.

For the ITT, we estimate:

$$y_{it} = \beta_X (T_i^X \times Post_{it}^X) + \beta_Y (T_i^Y \times Post_{it}^Y) + \alpha_i + \delta_d \times \phi_t + \varepsilon_{it}, \quad (\text{E.1})$$

where y_{it} is arrears for consumer i in month-of-sample t , T_i^X and T_i^Y are treatment assignment indicators for groups X and Y , α_i are household fixed effects, and $\delta_d \times \phi_m$ are DC-by-month-of-sample fixed effects. For Experiment 2, we also include $\eta_g \times \phi_m$ – group-by-month-of-sample fixed effects – to use the within-group experimental variation and to address any (time-varying) differences in group composition, since randomization took place within but not across groups. ε_{it} is an error term, clustered at the consumer level.

For the LATE, we estimate the corresponding IV specification:

$$y_{it} = \theta_X (D_i^X \times Post_{it}^X) + \theta_Y (D_i^Y \times Post_{it}^Y) + \alpha_i + \delta_d \times \phi_m + u_{it}, \quad (\text{E.2})$$

where D_i^X and D_i^Y are indicators for actual receipt/selection into arms X and Y . The endogenous variables $D_i^X \times Post_{it}^X$ and $D_i^Y \times Post_{it}^Y$ are instrumented by $T_i^X \times Post_{it}^X$ and

$T_i^Y \times Post_{it}^Y$, respectively. All other terms are as in Equation (E.1), including the addition of $\eta_g \times \phi_m$ for Experiment 2. u_{it} is an error term, clustered at the consumer level.

After estimating these pooled specifications, we test whether the two group-specific treatment effects are equal. For the ITT, the null hypothesis is

$$H_0 : \beta_X = \beta_Y,$$

and for the LATE, the null hypothesis is

$$H_0 : \theta_X = \theta_Y.$$

Equivalently, the reported static estimate is the difference between the two effects, $\hat{\beta}_X - \hat{\beta}_Y$ for ITT and $\hat{\theta}_X - \hat{\theta}_Y$ for LATE.

2. Joint test of monthly effects. This test evaluates whether the *trajectory* of treatment effects is the same across groups. Instead of estimating a single average post-period effect for each arm, we estimate group-specific event-study coefficients and test whether the two groups have equal effects in each post-treatment month.

For the ITT, we estimate:

$$y_{it} = \sum_{k \neq -1} \beta_k^X (T_i^X \times \mathbf{1}\{mtt_{it}^X = k\}) + \sum_{k \neq -1} \beta_k^Y (T_i^Y \times \mathbf{1}\{mtt_{it}^Y = k\}) + \alpha_i + \delta_d \times \phi_t + \varepsilon_{it}, \quad (\text{E.3})$$

where mtt_{it}^X and mtt_{it}^Y denote event time relative to the relevant treatment date for arms X and Y , respectively. The omitted event-time period is $k = -1$, and the estimation window is $k = -4, \dots, 6$. All other terms are as in Equation (E.1), including the addition of $\eta_g \times \phi_t$ for Experiment 2.

For the LATE, we estimate the corresponding IV event-study specification:

$$y_{it} = \sum_{k \neq -1} \theta_k^X (D_i^X \times \mathbf{1}\{mtt_{it}^X = k\}) + \sum_{k \neq -1} \theta_k^Y (D_i^Y \times \mathbf{1}\{mtt_{it}^Y = k\}) + \alpha_i + \delta_d \times \phi_t + u_{it}, \quad (\text{E.4})$$

where each endogenous term

$$D_i^a \times \mathbf{1}\{mtt_{it}^a = k\}$$

is instrumented by the corresponding assignment term

$$T_i^a \times \mathbf{1}\{mtt_{it}^a = k\},$$

for $a \in \{X, Y\}$.

After estimating these event-study specifications, we jointly test whether the two groups have equal effects in each of the six post-treatment months. For the ITT, the null hypothesis is

$$H_0 : \beta_1^X = \beta_1^Y, \beta_2^X = \beta_2^Y, \dots, \beta_6^X = \beta_6^Y.$$

For the LATE, the corresponding null hypothesis is

$$H_0 : \theta_1^X = \theta_1^Y, \theta_2^X = \theta_2^Y, \dots, \theta_6^X = \theta_6^Y.$$

This produces an F -test with six numerator degrees of freedom. Unlike the static test, which compares average effects over the full post-period, this joint test rejects equality if the two treatment arms have different dynamic effects at any point in the six-month post-period.

F Additional results

F.1 Experiment 1

F.1.1 Jointly-estimated treatment effects

For consistency with Experiment 2, we estimate treatment effects for each arm separately in Table 2. When we test for equality across treatment groups as described in Appendix E and presented in Table 3, however, we estimate treatment effects jointly. Appendix Table F.1 presents our main treatment effects estimated in a single regression:

$$Y_{it} = \beta_0 + \sum_{k \in \{B, C, D, E, F\}} \beta^k (T_i^k \times post_{it}^k) + \gamma_i + \delta_d \times \phi_t + \varepsilon_{it} \quad (\text{F.1})$$

where Y_{it} is an outcome for consumer i in month-of-sample t , T_i^k is an indicator for assignment to treatment group $k \in \{B, C, D, E, F\}$, $post_{it}^k$ is a dummy that takes a value of 1 for observations in months after treatment k , γ_i are consumer fixed effects, $\delta_d \times \phi_t$ are DC-by-month-of-sample fixed effects, and ε_{it} is an error term, clustered at the consumer level.

Our estimated treatment effects are very similar to those in Table 2 (though not identical because we do not fully interact the model).

Table F.1: Experiment 1 ITTs and LATEs on electricity consumption, arrears, and collections (jointly estimated)


	ITT			LATE			N consumers
	Consumption	Arrears	Collections	Consumption	Arrears	Collections	
SMS $\times$ Post	-0.75 (0.76)	-9.54 (36.00)	-17.08 (9.59)	-1.23 (1.24)	-17.64 (58.91)	-27.72 (15.70)	27,225
SMS social $\times$ Post	0.16 (0.86)	-20.18 (37.19)	-21.99 (10.36)	0.25 (1.35)	-33.77 (58.61)	-34.42 (16.33)	27,225
Mailed notice $\times$ Post	-0.71 (0.79)	-160.28 (35.98)	11.30 (9.82)	-1.16 (1.29)	-264.11 (59.27)	18.91 (16.13)	27,225
In-person notice $\times$ Post	1.23 (1.07)	-29.65 (48.45)	-4.31 (12.30)	1.51 (1.31)	-35.57 (59.19)	-5.17 (15.02)	27,225
Disconnection $\times$ Post	-0.89 (1.23)	0.11 (38.88)	-7.62 (11.25)	-1.45 (1.98)	-0.08 (62.87)	-12.53 (18.21)	27,225
Dep var mean (control)	123.13	2,729.04	231.92	123.13	2,729.04	231.92	

Notes: This table reports Experiment 1 ITT and LATE effects, estimated using a version of Equation (8) and (10) that includes all treatments arms and the control group in the same regression. The estimation sample for consumption (kWh) and payments (INR) is restricted to the event-time window $-5 \leq$ months to treatment ≤ 5 , whereas this window is transposed forward by one month for arrears (INR) and covers the $[-4,6]$ window. Note that for arrears, the post dummy instead turns on if months to treatment ≥ 1 . The control group spans the earliest pre-treatment month through the latest post-treatment month. The last column reports the number of consumers included in the arrears estimation sample. All regressions include consumer fixed effects and DC-by-month-of-sample fixed effects, with standard errors clustered at the consumer level.

F.1.2 Control vs. business-as-usual

It is possible that the utility treated control consumers in the experiment differently than they would have in the status quo (e.g., perhaps they received less enforcement effort). To test for this, we compare consumers assigned to the control group (where the utility told staff not to conduct disconnections) to business-as-usual hold-out group (where the utility provided no instructions). Appendix Figure F.1 presents event studies, estimated using Equation (11), comparing these two groups. We use January 2022 – the month of the first experimental treatments – as the “treatment time.” We see no economically meaningful or statistically significant differences between the control and business-as-usual groups. Therefore, we interpret the effects of our treatment groups B, C, D, E, and F as relative to both the control and to the status quo.

Figure F.1: Control vs. business-as-usual



Notes: This figure presents event studies comparing the Experiment 1 control group to the business-as-usual group. Event time 0 is January 2022, when the experiment treatments began. For all event studies, we use month -1 as the reference category. Hollow markers indicate the pre-treatment period, whereas filled markers indicate the post-treatment period. We estimate effects using Equation (11). Vertical lines show 95 percent confidence intervals, computed from standard errors clustered at the customer level.

F.1.3 Computed payments

Based on our interactions with utility staff, there is a concern that the collections data are of lower quality than the billing data. Therefore we test for robustness using a new payment variable – computed payments – constructed using billing data alone (see Appendix B.1.2 above for details). Appendix Table F.2 presents ITT and LATE estimates for Experiment 1 using computed payments as the outcome. For Arms B, D, E, and F, the treatment effect estimates are within 1.96 standard errors of the effects on collections we report in Table 2. For Arm C (SMS social), we observe positive effects on computed payments and negative effects on collections, though both are extremely noisy. These results suggest that problems with the collections variable are not driving our estimates.

Table F.2: ITT and LATE estimates on computed payments

	ITT	LATE	N consumers
SMS $\times$ Post	-4.55 (15.30)	-7.44 (25.01)	12,370
SMS social $\times$ Post	11.00 (17.55)	17.30 (27.60)	12,369
Mailed notice $\times$ Post	19.23 (16.45)	31.64 (27.03)	12,379
In-person notice $\times$ Post	-5.68 (19.08)	-6.94 (23.32)	9,900
Disconnection $\times$ Post	-27.42 (19.51)	-44.26 (31.56)	9,895
Dep var mean (control)	550.51	550.51	

Notes: This table reports Experiment 1 ITT and LATE effects on computed payments estimated using Equations (8) and (10), respectively. Each cell presents a separate regression. The estimation sample is restricted to the event-time window $-5 \leq \text{months to treatment} \leq 5$. The last column reports the number of consumers included in the estimation sample. All regressions include consumer fixed effects and DC-by-month-of-sample fixed effects, with standard errors clustered at the consumer level.

F.1.4 Effects on economic well-being

A natural concern with any enforcement intervention that successfully extracts payments from indebted households is that the gains in utility revenue may come at the expense of household consumption of other goods. To test for this, we estimate impacts on a battery of self-reported endline survey outcomes covering monthly expenditure on food, non-food essentials, healthcare, and education; total monthly savings; and reported outstanding loan amounts. We use a cross-sectional ANCOVA regression:

$$Y_i^{Endline} = \sum_{k \in \{B, C, D, E, F\}} \beta_k Treat_i^k + \gamma Y_i^{Baseline} + \eta_s + \varepsilon_i$$

where $Y_i^{Endline}$ and $Y_i^{Baseline}$ are outcome variables measured at endline and baseline, respectively, for survey respondent i , $Treat_i^k$ are indicators for assignment to treatment arms B through F in Experiment 1; η_s is a strata fixed effect, and ε_i is an error term. Negative-coded sentinel responses (Refused = -1, Did not consent / not reach = -100, Don't know = -999) are treated as missing on both sides. An important caveat is that there was substantial and non-random non-response to these surveys (see Appendix B.2). Appendix Table F.3 reports the ITT estimates by treatment arm.

The point estimates imply that enforcement did not cause substantial financial harm to consumers in our sample, though the standard errors are large. We do not find evidence that mailed notices – the only arm to produce a meaningful reduction in arrears – led households to cut consumption on food, essentials, healthcare, or education in the months following treatment, nor that they reduced reported savings or increased reported borrowing. These results are in keeping with our evidence in Section 3 that non-payment is unlikely to be driven by an inability to pay.

Table F.3: ITT estimates on the economic wellbeing of the experimental sample

	(1) Expenditure (INR)	(2) Savings (INR)	(3) Highest monthly income (INR)	(4) Lowest monthly income (INR)
SMS	21.50 (446.6)	1035.2 (1814.6)	-7206.3 (5614.8)	-559.7 (922.7)
SMS social	-456.2 (323.2)	-1311.7 (1615.9)	-4528.2 (5448.5)	-171.0 (1025.3)
Mailed notice	8.014 (376.6)	1816.9 (2263.0)	322.6 (6527.9)	432.1 (1073.4)
In-person notice	175.4 (374.4)	2384.2 (2581.3)	-6297.1 (5580.4)	-1298.3 (958.2)
Disconnection	160.4 (553.6)	2623.9 (2198.0)	-2339.4 (6114.4)	98.19 (1021.1)
Baseline expenditure (INR)	0.0213 (0.00803)			
Baseline savings (INR)		0.0160 (0.0211)		
Baseline highest income (INR)			0.169 (0.136)	
Baseline lowest income (INR)				0.180 (0.0329)
Observations	1,519	1,025	727	784
Dep var mean (control)	7,267.3	7,408.2	23,799.1	7,600.0

Notes: This table presents effects of Experiment 1 treatments on measures of economic well-being, using an ANCOVA specification with endline survey outcomes. Each column uses its own complete-case sample; samples differ across columns because refusal rates vary by question. Robust standard errors in parentheses.

F.1.5 Return on investment by treatment arm

To calculate return on investment, we define the net return for each treatment arm as the reduction in arrears net of implementation costs. The arrears reduction for each arm is given by the corresponding 6-month ITT estimate on arrears (as negative ITT estimates indicate a reduction in outstanding debt). For treatment arm k , we compute:

$$\text{Net return}_k = -\hat{\beta}_k - C_k,$$

where $\hat{\beta}_k$ is the 6-month ITT estimate on arrears and C_k is the implementation cost per-consumer for treatment arm k . We then compute return on investment as:

$$ROI_k = \frac{\text{Net return}_k}{C_k}.$$

The cost computation for each arm is described below:

In-person notice. The cost of delivering an in-person notice has three components: travel, staff time, and printing. For travel, we assume that each delivery requires a 10 kilometre round trip, undertaken on a two-wheeler with mileage of 25 kilometres per litre. At a petrol price of 110 INR per litre, this implies a travel cost of 44 INR per notice. Staff time is calculated on a pro rata basis using the mean monthly salary of full-time linesmen, estimated at 51,695 INR from the utility office survey.³⁷ Assuming 26 working days per month and an eight-hour workday, a one-hour trip to deliver the notice costs 248 INR in staff time. Finally, we add 12 INR per notice for printing.

Disconnections. The cost of the disconnection arm includes the cost of the in-person notice described above, plus the additional cost of physically visiting the household to carry out the disconnection. For this additional visit, we assume that two linesmen are required. Using the travel-cost assumptions described above, the additional visit costs 44 INR. Labour costs are also calculated using the hourly wage implied by the salary assumptions above. Since two linesmen are assumed to spend one hour traveling, travel-time compensation amounts to 497 INR. The average time taken to disconnect one consumer is 27 minutes, based on responses from the utility office survey,³⁸ which implies an additional labour cost of 224 INR for the disconnection itself. Thus, over and above the in-person notice cost, the cost of carrying out a disconnection is 765 INR.

Mail notice. This arm involved postage (40 INR) and printing charges (12 INR).

SMS messages. The cost of sending one text message was 0.9 INR.

37. We asked two questions in the utility office survey that are used to compute the average monthly salary: (i) Who delivers the majority of disconnection notices, carries out disconnections, and carries out reconnections in your DC? (ii) How much are they paid every month?

38. The utility office survey asked: “How long does it take to disconnect one consumer? Enter in minutes.”

Table F.4: Return on investment (ROI) by treatment arm

Arm	Cost (INR)	Arrear reduction (INR)	ROI (%)	Scaled reduction (millions of INR)
SMS	0.9	26.27	2818.89	24.43
SMS social	0.9	36.64	3971.11	34.42
Mailed notice	52	177.28	240.92	120.66
In-person notice	304	46.93	-84.54	-247.15
Disconnection	1,075	0.16	-99.99	-1,035.29


Notes: This table presents return on investment for each treatment arm in Experiment 1. Cost and arrear reduction columns are in INR. These values represent amounts per consumer. The scaled reduction column measures expected revenue from scaling these treatments to the full eligible population, and is computed as (arrears reduction - cost) $\times$ 963,140, where 963,140 is the total number of consumers with average baseline arrears greater than 1,200 between August 2021 and December 2021.

F.2 Experiment 2

F.2.1 Effects on collections

Appendix Figure F.2 presents event studies for collections, in the same four-panel layout as Figure 6. As in Experiment 1, the payment response to a mailed notice is concentrated in the month of notice delivery and the immediately following month, with payment behaviour reverting to baseline thereafter. This is the characteristic “one-time payment” pattern: households respond to a threat from a credible messenger by making a single larger-than-usual payment, after which their month-to-month payment frequency returns to its pre-treatment level.

Figure F.2: Second mailed notice: Impacts on collections




Notes: This figure presents event study estimates of treatment effects of mailed notices on collections in Experiment 2. Each panel corresponds to one Experiment 2 group. For each panel, the estimation sample consists of only the relevant Experiment 2 group, including consumers assigned to treatment and control. For all event studies, we use month -1 as a reference category. Hollow markers indicate the pre-treatment period, whereas filled markers indicate the post-treatment period. We estimate effects using Equation (11). Vertical lines show 95 percent confidence intervals, computed from standard errors clustered at the customer level.

F.2.2 Mail → Mail: Heterogeneity by Experiment 1 response

Our model (Section 4) predicts that consumers who respond to a mailed notice in Experiment 1 by reducing arrears should respond to a subsequent mailed notice also. While we cannot estimate consumer-specific treatment effects, we test whether the effects of a mailed notice in Experiment 2 vary by the size of the change in arrears around the time of Experiment 1 for Arm D consumers.

Specifically, we first calculate the *household-specific* change in arrears following the receipt of a letter in *Experiment 1*. We then group households into quartiles based on the size of this difference (note that because our treatment effect is negative, quartile 1 contains the largest responders). Next, we estimate a variant of Equation 8 on the Mail → Mail group from Experiment 2, where we interact the $T_i \times post_{it}$ term with quartile dummies Q_i , yielding intent-to-treat effects for each quartile.

Figure F.3: Experiment 2: ITT effects by responsiveness to Experiment 1



Notes: This figure shows ITT effects for consumers in the Mail → Mail group in Experiment 2, separately by quartile of change in arrears (i.e., average post-period arrears, less average pre-period arrears) around Experiment 1. We estimate this using a modified version of Equation 8 which includes interactions between the quartile indicator and the treat $\times$ post indicator. Since the overall effect of mailed notices on arrears is negative, quartile 1 corresponds to the largest responders and quartile 4 to the smallest responders. Vertical bars denote 95 percent confidence intervals. Standard errors are clustered at the consumer level.

Figure F.3 presents the results of this exercise. The ITT reduction in arrears that we observe in this treatment group (142 INR from Table 4) is driven entirely by significant reductions from households who were in the first quartile of changes in arrears around the Experiment 1 treatment. To the extent that households who reduced arrears in response to the first treatment are more likely to be found in the first quartile than any other, the fact that Experiment 2 reductions also come mostly from this group suggests that households

who responded to the first letter might be responding again to the second, in keeping with our model prediction.

G Deviations from our pre-analysis plan

The first experiment in this paper was registered with the AEA RCT Registry (AEARCTR-0008742) prior to treatment assignment.³⁹ We filed the pre-analysis plan (henceforth, PAP) in December 2021. Though we have broadly followed the PAP for Experiment 1, here, we document deviations from the plan, including reasons for the change. We have grouped deviations into substantive additions beyond the PAP, changes to specifications, changes to terminology and treatment-arm labelling, changes to surveys and outcome measures, and items that are consistent with the PAP and noted only for completeness.

Substantive additions beyond the PAP

1. **Second experiment (Experiment 2).** The PAP describes only Experiment 1 – including the comparison of in-person and mailed threats. We designed Experiment 2 (the repeat-mailed-notice experiment) in response to the Experiment 1 results, approximately 14 months after the first experiment began to test the theoretical framework. As such, it is not part of the AEA-registered PAP.
2. **Theoretical model.** The PAP did not contain a formal theoretical model. We developed the model in this paper after we analyzed the Experiment 1 data.

Changes to specifications

3. **Main estimating equation** The PAP specifies a cross-sectional regression with strata fixed effects and pre-treatment levels as controls (PAP Equation 1):

$$Y_{is} = \beta_0 + \sum_k \beta_k T_i^k + \alpha_s + X_i + \varepsilon_{is}.$$

We instead estimate a panel difference-in-differences specification with consumer and month-by-DC fixed effects (Equation 8), run separately for each treatment arm. We use the panel regression as our primary specification because it allows us to use our monthly data more efficiently, and include consumer-level fixed effects. We use event-time specifications for each arm because in implementing the experiment treatment dates could not be kept identical across households (e.g., SMS messages deliver immediately, in-person notices take longer). As such there is not a common ‘post period’ across all arms and consumers. Appendix Table F.1 presents results from a specification where we estimate treatment effects for all arms jointly, which is perhaps closer to the original PAP; results are similar.

39. The registry entry is available from <https://www.socialscienceregistry.org/trials/8742>.

4. **Post-treatment window.** The PAP specifies a four-month averaging window for outcomes. The paper uses a six-month post-treatment window for the event studies and the LATE and ITT estimates. We gained access to more administrative data than we expected at the time of writing the PAP, so we prefer to use the longer series. This is especially valuable because the four-month specification risks being too short to capture steady state behavior. This is both because households may respond with a lag and because payments may get reflected in arrears with a one month lag depending on when they are made and credited.
5. **Dynamic specification.** The PAP specifies treatment interactions with three two-month period indicators (six months total). The paper instead reports monthly event-study coefficients (Equation 11). The monthly specification is more granular and strictly more informative. We prespecified a two-month resolution to guard against the possibility of some of our sample being on bimonthly billing cycles, but in practice, this was rare.

Changes to treatment-arm labelling and terminology

6. **Arm labelling.** The PAP uses arm labels A–G; the paper uses labels A–F. The hold-out sample that is used to test for experimentation effects is labelled in the PAP but not in the paper. This is for clarity and consistency with standard practice. We show the correspondence between PAP and paper labels in Appendix Table G.1.

Table G.1: Correspondence between PAP and paper arm labels

PAP	Paper	N
A. Pure control	A. Control	7,500
B. Business-as-usual	(hold back sample, no label)	2,500
C. SMS message	B. SMS reminder	5,000
D. SMS Social comparison	C. SMS with social comparison	5,000
E. Mailed notice	D. Mailed legal notice	5,000
F. In-person, no disconn.	E. In-person legal notice	2,500
G. In-person + disconn.	F. In-person + disconnection	2,500

7. **Notice terminology.** The PAP refers to the legal notice as a “disconnection notice”. The paper refers to it as a “legal notice”, with content identical to the PAP-described instrument. For a reader unfamiliar with the setting, “legal notice” more accurately describes the document, whose primary content is a warning of legal liability under the Electricity Act, 2003; disconnection follows only if non-payment continues.

Changes to surveys and outcomes

8. **Endline sample composition.** The PAP reports our plans to contact households who completed the baseline survey. The paper’s endline sample of 3,157 households

includes 2,122 baseline-endline matched households and 1,035 endline-only households drawn fresh from the experimental sample. The additional endline-only sample was added to increase statistical power given baseline non-response.

9. **Reconnections as an outcome.** The PAP lists reconnections among the administrative outcomes. The paper discusses reconnections qualitatively in the context of Arm F but does not report formal treatment effects on a reconnection outcome variable in the main text. This is because we were unable to obtain reliable data on re-connections.

Dropping heterogeneity analysis

10. **Heterogeneity strategy.** The PAP pre-specifies two approaches: (i) a pre-defined heterogeneity decomposition by division, binned pre-treatment arrears, and binned pre-treatment consumption, yielding 42 groups, and (ii) the Chernozhukov et al. (2020) machine learning approach. The paper does not report these because they are not relevant to our analysis and the PAP did not specify a hypothesis to be tested based on these results.