



Regular Article

Reducing nonpayment for public utilities: Experimental evidence from South Africa[☆]

Andrea Szabó, Gergely Ujhelyi

Department of Economics, University of Houston, Houston, TX 77004, USA

ARTICLE INFO

Article history:

Received 4 October 2014

Received in revised form 2 June 2015

Accepted 11 June 2015

Available online 17 June 2015

Keywords:

Nonpayment

Public information campaigns

Water conservation

ABSTRACT

Nonpayment for public utilities is an important constraint to expanding service access in developing countries. What are the causes of nonpayment and which policies are effective at addressing them? To study these questions, we implement and evaluate a randomized water education campaign in a low income peri-urban area in South Africa. We estimate substantial short-run treatment effects: on the order of a 25% increase in payments over a three-month period after which the effect dissipates. The evidence shows that the treatment did not operate by increasing consumers' information, or by creating reminders to pay or a threat of enforcement. Instead, households may have reciprocated the provider's efforts by paying more. Our findings provide evidence that strategies other than increased enforcement can lower nonpayment.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Improving people's access to basic utilities like electricity, water, or phone service is viewed as a key challenge in many developing countries. However, consumers' ability or willingness to pay for services can be an important constraint to investment in infrastructure. For example, the difficulty to collect unpaid bills has been cited as a major obstacle to improving electricity provision in India (Ahluwalia, 2002), the former Soviet Union (Lampietti et al., 2007), and Colombia (McRae, 2015). In South Africa, nonpayment presents a major problem for local governments and prevents the efficient use of the existing infrastructure for electricity, water, and sanitation. In 2011, South African households owed municipal governments 40 billion Rand (about 4 billion USD), equivalent to 25% of these governments' annual operating expenditures (Republic of South Africa, 2011).¹

What are the causes of nonpayment and which policies are effective at addressing them? While the textbook response to nonpayment,

punishment in the form of denied service, may work well in developed countries, this is typically not the case in developing countries. First, the incentive value of denying service is limited when consumers do not have enough income left to pay the bill. Second, aggressive enforcement could go against social perceptions of fairness and erode citizens' trust in local governments, resulting in even more nonpayment or even civil unrest.² In some cases nonpayment can be caused by consumer dissatisfaction with service delivery or a lack of trust in the provider. Punishing nonpayment by denying service would be highly counterproductive in these situations.

In most developing countries, utilities' response to nonpayment requires a delicate balancing act between various costly strategies. When consumers simply refuse to pay, have no individual meters or there are widespread illegal connections, utilities may not undertake any enforcement action (e.g., World Bank, 1999). In our setting, illegal connections are virtually nonexistent, consumers have individual meters, and consumption is highly price elastic (Szabó, forthcoming). Here, the water provider has purchased and installed restriction devices that limit the flow to a bare minimum for households with large outstanding balances (about a third of the population). In many cases these households will continue not paying, and may simply leave the taps open, perhaps with a container underneath to collect water. Such limited enforcement strategies are costly to the provider, lead to waste, and often do little to incentivize payment.

We explore an alternative strategy to reduce nonpayment: providing information. Purchasing and using water is a complicated activity: quantities used by various appliances are not directly observed by the

[☆] Previously entitled "Can Information Reduce Nonpayment for Public Utilities? Experimental Evidence from South Africa." We are grateful for the cooperation of Odi Water Services, in particular to Siphon Nkosi and his dedicated team for implementing the education program, and to Pieter Avenant for giving us access to the billing data. We also thank the Water Research Commission of South Africa for supporting our project, and Patrick Chiroro and his team at Impact Research International for implementing our surveys. We thank Hunt Allcott, Paul Ferraro, Vibhuti Mendiratta, our colleagues at the University of Houston, participants of MIEDC 2014, and an anonymous referee for useful comments. All views expressed here are the authors' and do not necessarily represent those of any of the organizations involved. Detailed supplementary materials, including education brochures and survey questionnaires, are available on the authors' websites.

E-mail addresses: aszabo2@uh.edu (A. Szabó), gujhelyi@uh.edu (G. Ujhelyi).

¹ By nonpayment, we mean a failure to pay the billed amount. This is different from a policy of providing free services (e.g., to the poor), under which consumers do not have to pay.

² In South Africa, the expansion of the water and sanitation infrastructure to poor black localities took place after the fall of Apartheid, and access to these services, codified in the constitution, is viewed as a requirement for human dignity.

household, and consumption and payment occur at different times. One possible reason of nonpayment is that households are unfamiliar with the billing system or the amount of water they consume in their everyday activities and accumulate high bills which they find difficult to pay. Providing information about various aspects of water consumption could improve the households' water management and consequently lower waste and their monthly bill.

To study this possibility, we implemented a randomized water education campaign in a collection of low income peri-urban townships in South Africa. Education officers visited a treatment group of 500 households to give them accessible information about various aspects of the water consumption process, including the water meter, the bill, and the amount of water used by various everyday activities. We evaluate the program combining administrative billing data on the full population of consumers with in-depth survey information on the treatment group and a control group.

We find that our information campaign was successful in reducing nonpayment in the short run, but not the long run. Compared to a control group, we find that treated households are more likely to pay their bill and make larger payments. We estimate that our treatment reduced the fraction of consumers making no payments by 4 percentage points and increased total payments by about 25% in the three months following the treatment. While temporary, these are large effects and they provide evidence that strategies other than increased enforcement can lower nonpayment.

We use a simple model and particular features of the research design to explore the channels that could explain our findings. In the model, our treatment can lower nonpayment in four ways: by improving information on water quantities, lowering the cost of understanding the bill and making payments, creating a desire for water conservation, or increasing the cost of nonpayment.

We find that the evidence is not consistent with the treatment operating solely by increasing consumers' information. Using direct measures of information, we find modest changes in households' knowledge in response to the treatment. Compared to control households, treated households are not much more likely to understand water quantities, know how much water they consume, or be able to read their water bill. We also find little evidence that our intervention created a sufficient desire for water conservation to explain the reduction in nonpayment. In particular, we see no decrease in water use corresponding to the lower rate of nonpayment. In our model, this has an important implication. Because water use serves as an indicator for a consumer's planned payment behavior, our findings imply that nonpayment in this sample primarily reflects unplanned circumstances (such as higher than expected bills) rather than households systematically planning to avoid payment.

This leaves the possibility that the treatment increased the cost of nonpayment, and we present suggestive evidence on three channels through which this may have happened. First, households may have misinterpreted the education campaign as the provider's attempt to exert pressure for payment or signal a future increase in enforcement. We do not see much evidence to suggest that such *perceived scrutiny* could explain the payment results. In particular, the increase in payments did not come from households who were at higher risk of enforcement action. Second, the education visits may have acted as *reminders* to consumers about outstanding bills. Zwane et al. (2011), Karlan et al. (forthcoming), and Allcott and Rogers (2014) provide evidence for the relevance of reminder effects in various settings. We argue that the absence of survey effects rules out this possibility in our case. A third possibility is that increased payments are an expression of *reciprocity* for the provider's education efforts. If consumers appreciated the provider's effort in reaching out to the community through the information campaign, they may have felt guilt about not paying more. Reciprocity has been shown to affect a variety of economic transactions (see Fehr and Gächter (2000) for a survey). Alm et al. (1992, 1993), Bazart and Bonein (2014) and Luttmer and Singhal (2014) discuss how taxpayers' reciprocity towards the government can reduce tax

evasion, and Karlan et al. (2013) show how banks in the Philippines can harness feelings of reciprocity to induce borrowers to repay a loan. As we show below, reciprocity provides a consistent explanation of the patterns observed in the data.

We are not aware of any experimental study on nonpayment for public utilities or other services. Perhaps the closest to our paper is a recent economics literature studying conservation campaigns for electricity and water (Allcott, 2011; Allcott and Rogers, 2014; Ayres et al., 2013; Ferraro and Miranda, 2013; Ferraro and Price, 2013; Jessoe and Rapson, 2014; Reiss and White, 2008).³ Our work differs from these in three key ways. First, we focus on nonpayment, which is a key issue in many countries. Second, we evaluate a water education campaign in a developing country, where the lack of information is known to be a problem and where small improvements in water use could have large impacts on household welfare. Third, while most campaigns in the literature are designed to generate psychological effects (e.g., through appeals to social norms), we explicitly focus on providing information.⁴

Our paper is also related to the literature on tax evasion. Non-experimental studies on tax evasion in developing countries include, among others, Alm et al. (1991), Torgler (2005), Gordon and Li (2009), and Kumler et al. (2013). For experimental studies (mostly in richer countries), see Slemrod (2007), Hallsworth et al. (2014), and Pomeranz (forthcoming). A major issue in this literature is misreporting and the failure to declare one's income. This is conceptually distinct from the problem of nonpayment, where a consumer has already received a bill, and the two are likely to involve different calculations by the individual (for example, evasion requires weighing the probability of an audit, while nonpayment occurs in a setting where the individual's debt is common knowledge). Nonpayment is also easier to study empirically because information on the true amount owed already exists, while this typically has to be estimated to measure misreporting. A recent paper by Hallsworth et al. (2014) studies the nonpayment, as opposed to the misreporting, of taxes. Like the conservation literature, they focus on pro-social appeals in letters sent out to taxpayers in the UK and find that the resulting psychological effects achieve large reductions in tax nonpayment.

More generally, our paper also relates to recent studies of information provision as a policy tool in various contexts ranging from providing water quality information (Benneer and Olmstead, 2008; Jalan and Somanathan, 2008; Madajewicz et al., 2007), through mitigating misleading advertising (Glaeser and Ujhelyi, 2010), to improving households' financial decisions (Chetty and Saez, 2013; Cole et al., 2011; Duflo and Saez, 2003). Our findings provide direct evidence that public information campaigns can affect behavior through channels other than increased information.

2. Research setting and design

2.1. Research setting

We conducted our research in cooperation with Odi Water, a small public water provider serving approximately 60,000 households in a group of "townships" (low income suburbs / villages) located an hour's drive north of Pretoria. The area has a well-functioning water infrastructure developed in the mid 1990s as part of government efforts to develop black neighborhoods after Apartheid. The provider is owned

³ A related literature in marketing and psychology is reviewed in Abrahamse et al. (2005).

⁴ In the context of previous studies, information on basic ways to save water or electricity is thought to be widely available in the population. Therefore even interventions that include tips for conservation (e.g., take showers instead of baths) are viewed as pro-social appeals rather than giving consumers new information (Allcott, 2011; Ferraro and Price, 2013). One study focused on information provision is Jessoe and Rapson (2014) who analyze the provision of real-time feedback on household electricity use.

and managed by the local government which also reviews and sets the price schedule annually (in July).⁵

On the supply side, the water market operates much as it does in developed countries. All households have modern individual water meters on their property; the meter is read every month and the household receives a bill in the mail (showing amount used, current charges, as well as any previous balance); payment options available include paying at one of the many supermarkets, paying at the provider's office, paying at the bank, or paying on-line. On the demand side, however, the market exhibits several anomalies. Many consumers apparently waste water - for example, it is not uncommon to see garden taps left open, with or without an overflowing bucket underneath. Households also appear to use water on some luxuries, such as washing their cars at home, or irrigating a flowerbed or lawn in the dry season. As a result, households often accumulate large bills that they have difficulty paying. In our data, the average household's monthly water bill is around 7% of its income, and its overdue balance is 9 times as large.⁶ Most consumers pay their bills infrequently. In the 3 months preceding our treatment, about a quarter of the households in our sample did not pay their bill, and only 15% paid every month. Payments that do occur are often in round figures, unrelated to the consumer's last bill or outstanding balance. Total payments over the same 3-month period were in multiples of 100 Rand for half of the households that made any payments (see Fig. 1). Since consumers typically pay at large supermarkets, banks, or the provider's office, round figures cannot be explained by a lack of small change but likely reflect households' attempt to budget for water in the face of large outstanding balances.

Unpaid balances accrue interest, and the provider restricts the water supply of the worst offenders. This is done by installing a flow limiter that reduces water flow to a bare minimum. Restricted households are charged an additional fee for this device.

Clearly, waste and nonpayment are costly both to the households and to the water utility. Why do these behaviors arise? Based on Odi Water's experience, as well as our own visits in the field, households' lack of understanding regarding water consumption is a major candidate. The issue is not the availability of information: indeed, information is widely available in a format that most consumers from Western countries would consider standard (water meter on the property, detailed monthly bills, a customer service department to answer questions). In our baseline survey described below, 99.5% of households knew where their water meter was located,⁷ 97.8% understood the basic operation of the meter (that numbers on the dial would increase when water was being used), and 95.7% stated that they regularly receive water bills. The issue is also not that consumers simply do not care about water. In our sample, close to 40% of respondents stated recently talking to neighbors or friends about water use. Instead, the primary issue appears to be that consumers have trouble understanding the information that is presented to them. For example, over 80% of consumers were unable to tell their consumption from their water bill. In general, households exhibited very little familiarity with the meaning of the numbers on the meter and the units in which their water consumption was being measured. When asked to guess how much water their household used, only 8 households (1%) stated their consumption in kiloliters, the units of measurement used by the provider (1 kl = 1000 liters \approx 264 gallons). While in principle household could have multiplied their consumption by 1000 and responded accurately in liters, it is clear that this did not happen. Among those answering in liters, 98% gave numbers lower than the median consumption of 12,000 liters, and 61% gave numbers less than or equal to 1000 liters. There is also a




Fig. 1. Distribution of payments made (August 2012).

lack of knowledge about the consumption process, e.g., how much water is used in various everyday activities. We asked households to compare pairs of activities in terms of their water usage. In each pair, one activity used at least twice as much water as the other. Only 14% of respondents ranked each pair correctly, and 45% ranked less than half of them correctly.

It is plausible that lack of information could hinder households' ability to budget their water consumption and make sure they can afford what they use. In fact, some consumers have started to voluntarily request that the provider install restriction devices on their service to help them better manage their consumption.

2.2. Description of the intervention

In an attempt to improve households' information, we designed and implemented an in-depth water education program in cooperation with Odi Water officials. The program consisted of household visits by Odi Water education officers trained by us specifically for this project. Officers introduced themselves with the following script: "Hello, my name is XXX, I come from Odi Water. Odi Water is experimenting with a new project to provide households with information on using water. Your household was randomly selected to receive this information. I am not here to read your meter or check if you pay your bills. I would simply like to give you some information on using water which could help your household save money on water." Education officers are members of the local community employed by Odi's Marketing Department. They are fieldworkers whose regular job is to talk to people in the townships about their needs and concerns regarding the service, answer questions, and report back problems such as leaking pipes. Education officers' socio-economic background is similar to those of the households they visit. Seeing them walking around the neighborhood knocking on doors is not unusual and most households tend to welcome this type of attention. The fact that only 0.8% of households refused the education visit (including those that could not be reached) shows that respondents understood that the officers were not there to harass them or force them to pay their bills.

Visits were conducted in November and December 2012, and each visit lasted between 30 minutes to 1 hour. During the visit, the officers gave the households 5 brochures containing information on specific aspects of water usage: reading the water meter, understanding the bill, detecting and fixing leaks, tips on conserving water indoors, tips on conserving water outdoors. They explained the contents of each brochure to the household, highlighting specific points agreed upon during our training session. We designed and wrote the brochures ourselves, with feedback from Odi Water's marketing department, drawing from water information campaigns developed for primary school students in South Africa, as well as public information campaigns in the US.

⁵ Szabó (forthcoming) provides further details on the setting as well as on the administrative data used here.

⁶ By comparison, the average US household spends less than 0.5% of its income on water (American Water Works Association, 1999).

⁷ By comparison, in one North-American study, 11.3% of respondents did not know whether they had one or two water meters (American Water Works Association, 1999). This exemplifies the higher salience of water related issues in Southern Africa.

(Copies of the brochures are available on the authors' websites.) All information in the brochures was presented in an accessible and reader-friendly manner (colors, pictures, examples). For example, one section of the brochure on indoor water conservation explained how to save water with every toilet flush. "Step 1: Use a large soft drink bottle (or several small ones). Fill it partially with pebbles. Fill the rest of the way with water. Step 2: Close the lid tightly and place it in the tank. If it floats or moves around, go back and add more pebbles. Make sure that the bottle doesn't obstruct the flushing mechanism." These instructions were accompanied by a picture of someone filling a plastic bottle, and another one showing the bottle sitting in the toilet tank. Another brochure showed a picture of a water bill, highlighting and explaining the most important pieces of information shown on the bill (last month's usage, amount due, outstanding balance, etc.). Brochures were available both in English and the main local language (Setswana), and the education officers conducted the visits in the households' preferred language. Feedback from the experiment suggested that households were delighted with the information campaign.

Compared to related interventions analyzed in the literature, our treatment had two distinguishing features. First, our treatment was explicitly focused on information provision, and we deliberately tried to minimize the social pressure component as much as possible. Our education materials used descriptive rather than prescriptive language. For example, they described the various ways available for households to pay their water bill but did not say "you should pay your bills." The education officers were also trained to provide information only and not tell households what they should or should not do. Second, we conducted our campaign in a developing country setting where the lack of information is known to be a serious issue.

2.3. Sampling, implementation, and data

2.3.1. Administrative data

We were granted access to Odi Water's billing data for the full population of residential consumers (excluding commercial users). This contains consumers' names and addresses, monthly consumption in kiloliters, payment, whether they were restricted (yes/no), and whether they had a registered "indigent" status providing discounted water pricing (yes/no).⁸ The data is in the form of computer files that have to be downloaded on site from Odi's servers. We obtained this data for the period ending 6 months after our intervention.

2.3.2. Sampling

In February 2012, we randomly selected 500 treatment and 500 control households from the population to participate in the project. We excluded consumers using more than 300 kl (or 25 times the average). These accounts, comprising 0.3% of the population, are likely associated with unreported commercial activities or major leaks. We also excluded consumers whose account was less than a year old to ensure that participating households would all be experienced in using the local water infrastructure, paying the provider's bills, etc. Participating households were selected via stratified random sampling, with stratification based on the administrative information available at the time (quartiles of water consumption, registered indigent status, restricted status, and whether the consumer had made a payment on his water bill during the previous year), resulting in 32 strata.

2.3.3. Surveys and intervention

A baseline survey was administered to participating households in March - April 2012. This baseline survey collected household characteristics, as well as detailed information regarding households' knowledge about water consumption. Throughout the paper by "information" we

will mean knowledge of facts as measured in our surveys. This includes understanding of water quantities, one's bill and consumption, the price of water, and which everyday activities use the most water. Our findings regarding changes of information refer to the 14 measures we use to capture these (see Section 4.1 below for details).

Our survey used an independent local survey company with extensive field experience in the area, and the surveyors were young people living in neighboring communities. They introduced themselves as working for researchers at the University of Houston who were interested in gathering information about water needs and water usage in the area. Most locals are quite social, and the typical interview would feel like a conversation about the topics in the survey rather than a formal Q&A session. When training our surveyors, we particularly encouraged this approach for questions designed to measure households' knowledge. We wanted to make sure that respondents would feel at ease telling us what they knew and did not know, rather than feel that they were taking a test. The goal was to present these questions as if they were part of a "fun guessing game."

The education program took place in November - December 2012 when Odi Water education officers visited the 500 treatment households. Finally, a follow-up survey was administered to all participating households in February 2013. Note that the water price schedule, reviewed by the local government every July, is fixed throughout the intervention and the followup survey. In our regressions below, all payment and consumption data corresponds to the same price schedule.

Throughout the project our unit of analysis is the household. This makes sense because water is consumed, and paid for, jointly by all members of the household, and both consumption and payment is measured at the household level. It was also logistically infeasible to target our treatment to specific individuals within the household.⁹

2.3.4. Missing data and attrition

Due to logistical difficulties and funding issues, we only managed to gather baseline survey data for 803 households. (Note that administrative baseline data, including data on consumption, payment, restriction and indigent status is available for the entire sample.) For regressions where we control for baseline survey characteristics, we deal with missing baseline information by imputation (see, e.g., White and Thompson (2005)). Specifically, for categorical variables we create an additional "missing" category, while for continuous variables we replace missing values with their means, and create an additional indicator that takes a value of 1 for these observations and 0 otherwise.

During the education visits and the follow-up survey, 8 of our participating 1000 households could not be reached after multiple attempts or refused. Furthermore, an examination of Odi Water records revealed a name change on the account of 26 households during our study period. We exclude these households from the analysis, and restrict our attention to the remaining 966 households, implying a low attrition rate of 3.4%.¹⁰ In the month of May (5 months after our intervention) the payment data includes two major outliers in the treatment group: a payment of 9400 Rand and a payment of 170,000 Rand. By comparison, average monthly payment is around 100 Rand, and the next highest payment in our entire study period is 3200 Rand. We exclude these two outliers in the payment regressions for the month of May.

⁹ For both the surveys and the treatment, households were identified based on their billing information, which included the name (last name and first initial) and address the account was under. Surveyors and education officers were instructed to look for the person whose name was on the bill. If that person was not home, they were to talk to an adult member of his or her household (and revisit if such a person was not available either). Targeting specific individuals would have required collecting personal information to identify those individuals. This would have raised human subjects concerns and would have made respondents less willing to participate.

¹⁰ Out of these 966, we have baseline data for 776 households (80%). We impute missing baseline data as described above. Of course, missing follow-up data is never imputed, so the number of observations in some regressions is less than 966 due to missing variables.

⁸ Subject to an income threshold, households can register with the municipality as "indigent" to receive discounted pricing on various public services. In the case of water, indigent households receive the first 6 kl free of charge.

2.3.5. Sample characteristics

Table 1 shows means and standard deviations of various observables in our treatment and control groups. Not surprisingly, given the fine level at which we were able to stratify, the two groups are fully balanced on observables. As shown at the bottom of the table, the F-statistic for the hypothesis that the treatment and control groups are balanced on all observables jointly has a p-value 0.78. We also test for any differential composition of attrition in the treatment and control groups and find no evidence of it.

Table 1
Testing the balance of observables across groups.

	Control	Treatment	Difference
<i>Panel A: Administrative data</i>			
Consumption (kl)	14.965 (0.620)	16.834 (1.315)	1.869 (1.454)
Payment (Rand)	278.450 (18.337)	242.509 (15.941)	−35.941 (24.297)
Restricted	0.294 (0.021)	0.292 (0.021)	−0.003 (0.029)
Indigent	0.286 (0.021)	0.298 (0.021)	0.012 (0.029)
Payment (yes/no)	0.566 (0.023)	0.515 (0.023)	−0.050 (0.032)
<i>Panel B: Survey data</i>			
Baseline survey	0.812 (0.018)	0.795 (0.018)	−0.017 (0.026)
Informal shacks	0.123 (0.017)	0.129 (0.017)	0.006 (0.024)
Employed hh members	1.048 (0.032)	0.996 (0.030)	−0.052 (0.044)
HH size	4.338 (0.078)	4.481 (0.094)	0.143 (0.122)
No formal schooling	0.010 (0.005)	0.005 (0.004)	−0.005 (0.006)
Some primary school	0.010 (0.005)	0.010 (0.005)	−0.000 (0.007)
Primary school	0.065 (0.013)	0.088 (0.014)	0.023 (0.019)
Some high school	0.217 (0.021)	0.202 (0.020)	−0.016 (0.029)
High school	0.434 (0.025)	0.432 (0.025)	−0.003 (0.036)
Some higher educ.	0.165 (0.019)	0.152 (0.018)	−0.013 (0.026)
Higher education	0.098 (0.015)	0.111 (0.016)	0.013 (0.022)
Hot water	0.691 (0.023)	0.641 (0.024)	−0.050 (0.034)
Owns car	0.369 (0.025)	0.364 (0.025)	−0.005 (0.035)
Owns refrigerator	0.977 (0.008)	0.982 (0.007)	0.005 (0.010)
Income (Rand)	7,056.548 (236.554)	6,736.557 (226.010)	−319.990 (327.167)
N. sampled neighbors	0.125 (0.016)	0.140 (0.017)	0.014 (0.023)
Has treated neighbor	0.058 (0.011)	0.074 (0.012)	0.015 (0.016)
Joint test for balance (p-value)		0.74	(0.783)
Joint test for no differential attrition (p-value)		0.53	(0.994)

Notes: The table presents the means of various observables in the treatment and control groups as well as their difference, with standard errors in parentheses. ‘Consumption’ is average consumption in the 3 months prior to the treatment. ‘Payment (Rand)’ is the household’s total payment during this time, and ‘Payment (yes/no)’ is 1 if the household has made a payment. ‘Baseline survey’ is 1 if we have baseline survey information on the household. ‘Informal shacks’ is 1 if there are informal shacks on the property. ‘Hot water’ is 1 if the household has hot running water. ‘N. sampled neighbors’ is the number of households included in the sample in a 30 meter radius, and ‘Has treated neighbor’ is 1 if one of these households is in the treatment group. To test for differential attrition we take the F-test of a regression of attrition status on the treatment indicator, the variables in the table, and their interactions with treatment status. In the third column ***, **, * denotes statistical significance at the 1, 5, and 10 percent level, respectively.

2.4. Conceptual framework

Our intervention was designed to decrease nonpayment. While it is natural to think that it may reduce nonpayment by increasing information and improving the management of household water consumption, the previous literature suggests several other channels through which nonpayment could fall. We now formalize these various channels.

Consider an individual who consumes two water using activities, w_1 and w_2 , and a third good x . Activity w_1 is a necessity, such as drinking or bathing, and the consumer’s demand for it is inelastic at \bar{w}_1 . Activity w_2 is non-essential, such as washing the car or watering the garden. Utility from w_2 and x is given by $\gamma u(w_2) + x$, where u is concave and reaches a maximum at w_2^{max} , which is the consumer’s satiation point for this activity. The parameter γ is positive if the consumer attaches value to this activity, and 0 otherwise (for example, if the consumer does not own a car or does not like washing it at home). Choose the unit of the water using activities such that one unit of either w_1 or w_2 uses one unit (kiloliter) of water. Suppose, however, that the consumer mistakenly believes that activity w_2 uses $\alpha < 1$ units of water. Let the average price of water be p , so that the consumer expects his water bill to be $p(\bar{w}_1 + \alpha w_2)$ (while his actual bill is $p(\bar{w}_1 + w_2)$). We normalize the price of x to 1 and let y denote the consumer’s income.

The consumer is billed for water after consumption has taken place and can decide whether to pay his bill. We capture this by considering a two-stage decision problem. In stage 1, the consumer chooses w_2 and $w_1 = \bar{w}_1$ and makes a plan for stage 2 regarding the consumption of good x , and whether he will pay his bill (for simplicity we focus here on the decision to pay rather than the amount paid, so we do not allow partial payments). In stage 2, he receives his bill, and decides whether or not to pay it, spending any income he has left on good x . If the consumer chooses not to pay, he incurs a cost of m . This could represent either a utility cost, such as a ‘moral cost’ of cheating, or a monetary cost, such as the discounted present value of any future penalties associated with nonpayment. The consumer may also incur costs if he chooses to pay: he has to find the bill, figure out how much he owes, and then travel to the supermarket or the provider’s office to make a payment. We denote these costs with f , and assume that $m - f < y$.

The consumer’s problem in stage 1 is

$$\max_{w_2, x, B \in \{0,1\}} \gamma u(w_2) + x - Bf - (1-B)m$$

$$\text{s.t. } x + Bp(\bar{w}_1 + \alpha w_2) = y,$$

where $B = 1$ if the consumer plans to pay his bill and $B = 0$ if he does not. Let w_2^* denote the consumer’s optimal choice of w_2 when he plans to pay for water. This is defined by $\gamma u'(w_2^*) = \alpha p$. If he plans not to pay, the consumer’s optimal choice is $w_2 = w_2^{max}$.

There are three possible outcomes. When

$$\gamma u(w_2^*) - p(\bar{w}_1 + \alpha w_2^*) < \gamma u(w_2^{max}) - (m - f), \quad (1)$$

the consumer plans not to pay his bill (and does not pay it). When (1) does not hold, but

$$p(\bar{w}_1 + w_2^*) > m - f, \quad (2)$$

the consumer plans to pay his bill but does not pay it because it is larger than expected. Finally, when neither (1) nor (2) holds, the consumer plans to pay and does pay.

This simple model highlights that nonpayment can be planned or unplanned and implies a straightforward way in which these can be tested apart. Consider a consumer who plans not to pay ((1) holds) and imagine that some exogenous event (e.g., an increase in m) reverses the inequality in (1) so that the consumer changes his plan and pays. This will be accompanied by a decrease in water use from w_2^{max} to w_2^* . By contrast, consider a consumer who plans to pay but does not, because (2) holds. If an exogenous event reverses (2), this consumer will

Table 2
Model mechanisms for reduced nonpayment.

Mechanism	Model	Prediction	
		Consumption	Information
Better information on water quantities (e.g., showers use less water than baths)	higher α	decrease	increase (info about water usage)
Increased desire to conserve water (e.g., washing car at home wastes water)	lower γ	decrease	no prediction
Lower cost of payment (e.g., bill becomes easier to understand)	lower f	decrease (if planned nonpayment) or unchanged (if unplanned nonpayment)	increase (info about billing and payment)
Increased cost of nonpayment (e.g., reminder effect, scrutiny, feelings of reciprocity)	higher m	decrease (if planned nonpayment) or unchanged (if unplanned nonpayment)	no prediction

also start to pay, but his consumption will remain unchanged at w_2^* . The intuition is simple: because consumption occurs before payment is due, planned payment behavior should affect consumption while unplanned payment behavior should not. Thus, studying changes in payment behavior together with changes in consumption can make it possible to distinguish whether nonpayment was planned or unplanned.

The above framework highlights four channels through which our intervention can reduce nonpayment. These are summarized in Table 2 and discussed in detail below.

Better information on water quantities. Suppose the consumer learns that the amount of water α used by activity w_2 is higher than what he thought. This lowers w_2^* , the level of this activity chosen by a consumer who plans to pay ($\frac{\partial w_2^*}{\partial \alpha} = \frac{p}{\gamma u'(w_2)} < 0$). Because this leads to a lower water bill, it makes unplanned nonpayment less likely (the left-hand-side of (2) falls). A reduction in nonpayment is always accompanied by lower water use in this case.

Increased desire to conserve water. Suppose that the treatment creates a desire for water conservation. While we consciously avoided pro-social messages like the ones used in conservation campaigns, consumers may have understood from the education visits that certain behaviors (e.g., washing the car at home) are socially wasteful and therefore “bad.” For example, Allcott (2011) and Ferraro and Price (2013) show that perceived social norms induce conservation in the US.

In the above framework, we can capture these effects by assuming that the treatment lowers γ , the consumer’s utility for the non-essential water-using activity. This again lowers w_2^* , making unplanned nonpayment less likely based on (2). From (1), a lower γ also makes it more likely that the consumer will plan to pay.¹¹ In each case, the consumer’s water use falls.

Lower cost of payment. Suppose that the treatment lowers f , the consumer’s cost of paying. For example, the treatment may inform a consumer about more convenient payment options, or it may lower the effort required to understand the bill and figure out the amount due. This would make both inequality (1) and (2) less likely to hold: nonpayment becomes less likely. For consumers whose planned payment behavior changes, water use declines from w_2^{max} to w_2^* . For consumers who become more likely to pay without a change in planned payment behavior, water use remains unchanged at w_2^* . Thus, while the cost-of-payment channel also predicts increased payments, unlike the above two channels this can be consistent with no change in water use.

Increased cost of nonpayment. Finally, suppose that the treatment’s effect is to raise m , the consumer’s cost of not paying. This could be so for several reasons. First, the education visit may simply remind a consumer of his outstanding bill. Zwane et al. (2011), Karlan et al. (forthcoming) and Stango and Zinman (2014) provide evidence on the

role of reminders in various contexts. Allcott and Rogers (2014) analyze reminder effects in inducing electricity savings in US conservation campaigns, and Jessoe and Rapson (2014) find that increasing the salience of electricity usage induces households to conserve more energy. Second, the consumer could perceive the visit as increased scrutiny by the provider and feel pressured to pay his bill. As described above, we went to great lengths to ensure that the education visits focus on transmitting neutral information rather than prescriptive messages on how consumers should behave. The household’s own bill was never discussed, and the education officers did not collect any kind of information during the visit. Nevertheless it is conceivable that households may view the visits as precursors to a future “crackdown” on nonpayment, raising m . A third possibility is that our treatment increased the cost of nonpayment m by creating feelings of reciprocity towards the provider. The consumer might appreciate the provider’s efforts in reaching out to the households through the education campaign, and might “feel bad” about not paying. Karlan et al. (2013) find that reciprocity towards bank employees can motivate consumers to repay a loan. Similarly, emotions like reciprocity and trust towards the government are thought to be relevant determinants of tax payments (see Luttmer and Singhal (2014) and the references therein). Studies have found less avoidance when individuals are exposed to pro-social messages (Hallsworth et al., 2014) or when they perceive that the government provides them with valuable public goods (Alm et al., 1992, 1993). In the South African context, where Apartheid left a legacy of suspicion towards government institutions, it seems plausible that payment behavior could be affected by similar feelings towards the municipal water provider.

In the model, an increase in the cost of nonpayment m has the same effects as a lower cost of payments f . Both inequality (1) and (2) become less likely to hold, leading to reduced nonpayment and either a reduction in water use (if consumers’ planned payment behavior changes) or no change in water use (if planned behavior is unchanged). While the effect of higher m or lower f is formally identical, we will be able to differentiate the two empirically by looking directly at consumers’ information.

To preview our empirical findings below, we will show that while there are large treatment effects on nonpayment, there are no corresponding declines in consumption. This rules out the treatment operating primarily through better information on water quantities or an increased desire to conserve water (first two rows in Table 2). We also find no increase in information that would lower the cost of payment, ruling out this channel (third row in the table). This leaves the possibility that the treatment increased the cost of nonpayment (fourth row). Because nonpayment has gone down without an average change in water use, we conclude that, in the context of the above model, nonpayment was due to consumers’ unplanned, rather than planned, behavior. Based on the discussion above, we consider three ways in which our treatment may have increased the cost of nonpayment: scrutiny, reminders, and reciprocity. While we cannot directly test between these, we present suggestive evidence inconsistent with the increased scrutiny and reminder channels but consistent with the idea of reciprocity for the provider’s efforts.

¹¹ Both sides of the inequality in (1) decrease but because $u(w_2^{max}) > u(w_2^*)$, the right-hand side decreases faster.

3. Specification and results

3.1. Specification

Given our randomized treatment, we can estimate treatment effects consistently from the following simple regression:

$$y_i = \beta_0 + \beta_1 T_i + \varepsilon_i, \quad (3)$$

where y_i is the outcome of interest for household i and T_i is an indicator equal to 1 for treated households. All payment and consumption quantities are in logs (we add 1 to every observation before taking logs to account for 0 values). To increase the precision of the estimates, we also include in (3) indicators for the strata used in sampling and the baseline value of the outcome y . In the Online Appendix, we also present estimates controlling for various demographics. We estimate all regressions in the main text using OLS and in some cases report nonlinear specifications in the Online Appendix.

We have 5 months of post-intervention data (January - May 2013) for payments and 6 months (January - June) for consumption.¹² We estimate Eq. (1) pooling the first three months of post-treatment data as well as separately for each month.

3.2. Effects on payment and consumption

We begin by studying changes in consumers' payment and consumption caused by the treatment. Our main results are in the first row of Table 3 which looks at the administrative data in the 3 months following the intervention. The first cell compares households' total payment in the treatment and control groups over this period. We estimate positive and significant treatment effects: on average, our treatment increased total payment by approximately 25% in the first quarter following the treatment. The next two cells look at the incidence of payment rather than the amount paid. In the second cell the dependent variable is an indicator for whether the consumer made a payment over the given period. The results show that the treatment increased the fraction of households making at least one payment in the three months following the intervention by about 4 percentage points relative to a mean of 54%. In the third cell, we also see a small but statistically significant increase in the total number of payments over this three-month period. With an average of 1 in the control group, the point estimate of 0.09 corresponds to 45 households making an extra payment over this period.¹³

The last cell in the first row of Table 3 shows the estimated treatment effect on households' average consumption over the three months. These effects are small and statistically insignificant. Our treatment appears to have reduced nonpayment without a corresponding effect on consumption.

In Panel B of Table 3, we show monthly payment and consumption results extending 5 months after the treatment (6 for consumption for which we have more data). The last rows of the table show the p-values for tests of equal treatment effects across months (relative to January, the first month following the treatment).

For payment measures, we see larger effects in the first two to three months, a smaller effect in the third month, and statistically insignificant negative effects for subsequent months. For example, in column (1) of Panel B we estimate a 21–25% increase in total payment in the first two months and a 13% increase in the third month. This is followed by a statistically significant decline in the effect for 1 month. We see a

¹² January 2013 is the first post-treatment month for the full sample. Because conducting the 500 household visits took several weeks in November–December 2012, the treatment did not take place simultaneously for all households.

¹³ These treatment effects could be underestimated if the treatment induced some households to register as indigent and receive a free water allowance. In the Online Appendix, we show that there is no evidence that our treatment had an effect on households' registered indigent status.

Table 3
Treatment effects on payment and consumption.

		(1)	(2)	(3)	(4)
		Payment amount	Payment propensity	Payment frequency	Consumption
<i>Panel A: First three months after treatment</i>					
Jan–March	Treatment effect	0.251*	0.038*	0.088*	0.005
		(0.129)	(0.022)	(0.046)	(0.030)
	Control mean	3.156	0.537	1.035	2.468
<i>Panel B: By month</i>					
Jan	Treatment effect	0.214	0.038		−0.052
		(0.135)	(0.026)		(0.058)
	Control mean	1.739	0.332		2.122
Feb	Treatment effect	0.246*	0.048*		0.018
		(0.137)	(0.026)		(0.017)
	Control mean	1.760	0.340		2.625
March	Treatment effect	0.128	0.024		−0.068
		(0.137)	(0.026)		(0.094)
	Control mean	1.940	0.363		1.733
Apr	Treatment effect	−0.136	−0.026		−0.058
		(0.133)	(0.025)		(0.083)
	Control mean	1.722	0.330		2.024
May	Treatment effect	−0.024	−0.014		−0.080
		(0.134)	(0.025)		(0.069)
	Control mean	1.568	0.306		2.268
June	Treatment effect				0.022
					(0.058)
	Control mean				2.097
<i>p-value for equal treatment effects</i>					
	Feb = Jan	0.855	0.744		0.231
	March = Jan	0.625	0.661		0.880
	Apr = Jan	0.033	0.038		0.955
	May = Jan	0.150	0.095		0.753
	June = Jan				0.337

Notes: Each cell presents the estimated treatment effect from a different regression. Dependent variables are for 3 months combined in Panel A and monthly in Panel B. Payment amount is log total payment over the given period; Payment propensity is an indicator equal to 1 if the household made a payment over this period and 0 otherwise; Payment frequency is the number of payments made over this period; Consumption is log average consumption over this period. All specifications control for sampling strata indicators and the value of the dependent variable for the 3 months prior to the treatment. The p-values for equal treatment effects are from Chi2 tests on the equality of the treatment coefficients when each pair of regressions is estimated as a system. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

similar pattern for payment propensity in column (2), indicating that our treatment had short run effects only.¹⁴

The pattern for the consumption results in column (4) is markedly different from the payment results. The consumption effects are always small and insignificant, with no discernible pattern over time. Increased payment does not appear systematically related to specific changes in consumption.¹⁵

¹⁴ Note that the treatment effects on payment in months 4 and 5 have a negative sign. This is not too surprising: households need to budget over time, and paying more in the first quarter will reduce their available budget in the second quarter. With data spanning a longer horizon and a repeated treatment, studying the pattern of these “rebound” effects would be an interesting question for future research.

¹⁵ Since payment and consumption are bounded below by 0, estimation methods that take into account such corner solutions may provide more precise results. In the Online Appendix, we estimate treatment effects on payment and consumption amounts using Tobit regressions and treatment effects on payment propensity using Probit. The payment results give somewhat larger marginal effects than those presented above. For example, we estimate a 26% increase in 3-month payments due to our treatment among those who make positive payments, and a much larger unconditional effect of 37% (reflecting the fact that some households switched from 0 to positive payments). The effects on consumption remain small and insignificant.

Table 4
Treatment effects on consumption by consumption quartile.

	First quartile Less than 7 kl	Second quartile 7–12 kl	Third quartile 12–19 kl	Fourth quartile More than 19 kl
<i>Panel A: Consumption</i>				
Treatment	0.075 (0.068)	0.046 (0.062)	−0.029 (0.058)	−0.093* (0.052)
<i>Panel B: Payment</i>				
Treatment	0.423 (0.293)	0.211 (0.258)	0.516** (0.247)	0.154 (0.270)
Number of observations	237	249	248	232

Notes: Each column presents treatment effects for a different consumption quartile of the sample. In Panel A, the dependent variable is average consumption during the 3 months following the treatment. In Panel B it is total payment over this period. Each specification controls for sampling strata indicators and the value of the dependent variable for the 3 months prior to the treatment. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

In Table 4 we break up the short run effects by consumption quartile. For consumption, we now find a statistically significant effect in the fourth quartile, with a reduction of 9.3% among the highest consumers in the sample. This may indicate that our treatment had an impact on conservation among the highest consumers. However, the payment effects do not appear to be driven by this: the estimated treatment effect on payment is actually the smallest in this quartile.

3.3. Discussion

While there is some indication that the treatment affected consumption, these changes were not large enough to significantly reduce water use. Although we found a statistically significant reduction in water use among the highest consumers, this particular group does not drive the payment results. On average, we find that the treatment reduced nonpayment without decreasing consumption. In terms of our model, this rules out the first two mechanisms listed in Table 2. If our treatment primarily operated by improving information on water quantities or by creating an increased desire to conserve water, we would expect to see a decline in consumption. Based on the evidence presented so far, this leaves the “lower cost of payment” and the “increased cost of nonpayment” channels as plausible explanations of our findings. We investigate these channels in detail below.¹⁶

In our model, the reduction in nonpayment coupled with no reduction in consumption has an important implication. Because payment occurs *after* consumption, planned nonpayment should be reflected in consumption, while unplanned nonpayment should not (see Section 2.4). Thus, our findings suggest that nonpayment in our sample occurs primarily for unplanned reasons (such as a higher than expected bill). The evidence is not consistent with the majority of consumers systematically planning not to pay.¹⁷

3.4. Cost-benefit calculation

A detailed back-of-the-envelope cost calculation for the project is given in the Online Appendix. The highest cost figure reported there is

¹⁶ It is of course possible that our treatment had an effect on information or conservation but these effects were not large enough to explain the payment results. In the Online Appendix, we look at households' self-reported conservation activities and find that our treatment had a significant impact on these. We look at direct measures of information in Section 4.1 below.

¹⁷ The idea of unplanned nonpayment is more plausible for consumers who pay at least some of their bills. Interacting the treatment indicator with whether the household paid at least one bill during the 6 months before the intervention provides some evidence that the increase in payment indeed came from these consumers (see Section 6 in the Online Appendix).

57.02 Rand per treated household. This includes printing costs, the imputed wage of the education officers and their supervisor, as well as fuel and vehicle amortization costs for transportation. Using an estimated treatment effect of 0.25 based on table 3 and the actual average payment of 302.18 Rand in the treatment group, the provider's extra revenue from the campaign is approximately 66.84 Rand per treated household in the 3 months following the intervention ($302.18 - 302.18/e^{0.25} = 66.84$). Based on these values, the provider achieved a positive profit of at least 4910 Rand from the campaign and a rate of return of 17% on its investment in the program. Since the effect of the treatment was short term, our findings suggest that the intervention would have to be repeated to generate sustained benefits for the provider. Whether a repeated intervention would lead to smaller or larger effects is an open question.¹⁸

4. Mechanisms

Based on our model, the empirical results above can be explained either through a “lower cost of payment” or an “increased cost of nonpayment” channel. To investigate these further, we first look at information. This provides direct evidence on the possibility that the treatment lowered the cost of payments by making it easier for consumers to understand various features of the billing process. Then, we look at different pieces of suggestive evidence to study any increase in the cost of nonpayment through reminders, scrutiny, or feelings of reciprocity.

4.1. Information

An important element of our research design is that we collected direct measures of households' knowledge targeted by our campaign. This allows us to provide direct evidence on the possible mechanisms behind our treatment effects. Our information campaign focused on four key areas of the water consumption process: (1) Understanding the meter; (2) Understanding the bill; (3) Understanding water quantities used in everyday activities and how to save water; (4) Detecting and fixing leaks. Separate sections in our surveys were designed to measure each of these areas. As described in Section 2.3, we took steps to ensure that respondents did not feel like they were being tested and felt comfortable telling our surveyors what they knew and did not know.

We consider a total of 14 measures of information. The Online Appendix presents detailed results, and we present a subset of these in Table 5. Table 9 in the Appendix gives the definition of the information measures used in this table. In short, we find at most a modest impact of our treatment on consumer knowledge.

In Table 5, the only effect that reaches statistical significance is an increase of 4 percentage points in the number of households who gave us their estimated water usage in kiloliters.¹⁹ However, many households seem to have become more familiar with the word “kiloliter” without learning what it means.²⁰ In the follow-up survey, with the water bill in their hands, 60% of respondents admit to not being able to tell their consumption from the bill, and another 28% read out an incorrect number from the bill. Overall, less than 12% of households are able to tell their consumption from the bill.²¹ There was no significant difference between treatment and control (Table 5, columns (2–3)).

¹⁸ If, as we argue below, increased payments are an expression of reciprocity, one would expect the effect of the first campaign to be the largest.

¹⁹ This effect loses significance if we control for the false discovery rate under multiple inference using the Benjamini and Hochberg (1995) method (see Anderson (2008)).

²⁰ Over 60% of those who answered in kiloliter after the treatment gave unrealistic numbers of several hundred or even thousands of kiloliters. When asked how many liters a kiloliter represented, only 3 respondents gave the correct answer. Others didn't know or were off by a factor of 10, with no significant difference between treatment and control.

²¹ We cannot be entirely sure about exactly which bill the respondent was looking at. 11.4% of respondents stated a number that corresponded to *any* bill the household received in the 6 months prior to the survey. This likely overestimates the fraction of households giving correct answers.

Table 5
Treatment effects on information.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Response in kl	Reads consumption from bill	Consumption accurate	Tariff in ballpark	Tariff error	Increasing tariff	Quiz score
Treatment effect	0.037* (0.020)	0.043 (0.037)	0.038 (0.023)	−0.016 (0.015)	−6.141 (9.256)	−0.023 (0.030)	0.059 (0.064)
Control mean and std. dev.	0.085 (0.279)	0.379 (0.486)	0.095 (0.294)	0.054 (0.225)	81.796 (232.596)	0.711 (0.454)	2.456 (0.963)
Multiple inference p-values	0.364	0.55	0.364	0.55	0.71	0.676	0.634
N	953	731	731	820	396	964	965

Notes: Each column corresponds to a different regression. The column headings give the dependent variable (definitions are in the Appendix). All regressions control for sampling strata indicators and the value of the dependent variable at baseline. We also present p-values that control for the false discovery rate under multiple inference using the Benjamini and Hochberg (1995) method. These are computed taking into account the full set of 14 information measures we consider (see Online Appendix for the measures not included in this table). Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

We also had several questions asking about the price of water. In the follow-up, less than 5% of households gave numbers in the ballpark of the true kiloliter price. These are the households who state prices between 5 and 25 Rand (the true kiloliter price is between 10 and 21, depending on consumption). About half of the remaining households say that they don't know the price, and the other half report prices that are much higher – the mean answer is 95 Rand. There was no difference between treatment and control either in the fraction of households whose answers were in the ballpark of the true price, or in how far off reported prices were from realistic values (Table 5, columns (4–5)). There was no difference in knowing the fact that the price schedule is increasing, i.e., that an additional kiloliter costs more when consumption is high than when it is low (column (6)).²²

These findings suggests that our treatment did not increase consumers' familiarity with their bill. In our model, this rules out the treatment operating primarily by lowering the cognitive costs associated with reading the bill and making payments (third row of Table 2).

We can also use the information measures to provide further evidence on the first mechanism in Table 2 which would imply that we should see improved understanding of water quantities. We used a "guessing game" to measure households' understanding of quantities of water used in various everyday activities. Four questions asked households to guess which of two activities used more water (e.g., using the outside hose for 10 minutes, or doing one load of laundry). In each pair, one activity typically uses at least twice as much water as the other. We based these questions on materials used in a South African primary school program. Like all other questions, these were read to the respondent by the surveyor, and we trained our surveyors to present the questions as a fun guessing exercise rather than as a test. The average number of correct answers is 2.5 for both the baseline and the follow-up survey, and the distribution of the number of correct responses is also very similar. There were no differences between treatment and control (Table 5, column (7)).²³

Overall, we do not see changes in information on a scale that would explain the large reductions in nonpayment we found.²⁴ From the mechanisms described in Table 2, the only one that can explain lower

nonpayment without corresponding changes in consumption and information is the last one, "increased cost of nonpayment." We now explore this mechanism in detail.

4.2. Increased cost of nonpayment: Reminders, scrutiny, and reciprocity

Section 2.4 describes three ways in which our treatment may have increased a household's cost of nonpayment. Based on previous literature, the visit may have (i) reminded consumers of their unpaid bills; (ii) increased perceived scrutiny and pressure towards making payments; and (iii) created feelings of reciprocity towards the provider and increased the psychological cost of nonpayment. While obtaining direct evidence on these channels is particularly hard, our research setting allows us to make some progress towards testing them apart.

4.2.1. Reminders

We can directly address the possibility that the education visit reminded consumers of their outstanding bills because if this was the main channel behind our effects, our surveys should have had a similar impact. Our surveys inquired at length about households' payment behavior. For example, we explicitly asked consumers whether they had ever missed a payment and if yes, why. We also asked them to find their water bill and read out their consumption, and respondents saw the interviewer recording all these answers. This is in stark contrast to the education visits, where the household's own bill or payment behavior was never discussed. Our education officers began their visits by telling the households that they were not there to check their bills or whether they had paid (see Section 2.2), and we explicitly trained the officers not to collect any kind of information during the visits. Thus, if the primary effect of the education visits was to act as reminders, then being surveyed should have an even larger effect on payments.²⁵

Because we have access to administrative data for the entire population, we can directly test for such survey effects. We randomly select a "new control group" of 500 households who did not participate in our study in any way (using the same stratification procedure as for participating households). In Table 6 we compare these households to our actual control group. The dependent variables in this table are for March, which is the first month after the follow-up survey (we present regressions for the 3-month period after the survey in the Online Appendix and find similar results). The variable *Control* takes the value of 1 if a household was surveyed (but not treated) in our study, and 0 if it did not participate in the study. Because of random sampling, the coefficient on *Control* consistently estimates the change in behavior caused by our two surveys only. We find no effect for either payment or consumption. Being surveyed did not affect behavior, the education visits did. This makes it unlikely that the effect of the education visits operated primarily by increasing the salience of households' unpaid bills and acting as reminders.

²⁵ See Zwane et al. (2011) for a discussion of the reminder effects of surveys in other contexts.

²² To measure this, we asked households to imagine flushing the toilet 1000 times and to guess whether the last flush would cost them more or less than the first.

²³ As we show in the Online Appendix, looking at the fraction of correct answers to individual questions we only find a significant effect for one of them, regarding the amount of water used when flushing the toilet. This is relevant because the toilet is typically the largest single source of indoor water use, responsible for over a quarter of consumption (American Water Works Association, 1999). Our treatment had a modest effect, raising the fraction of correct answers by 5%, relative to a mean of 73%.

²⁴ In the Online Appendix, we break down the information results by consumption quartile. We again do not see information effects corresponding to reduced nonpayment. We also present a detailed analysis of two potential confounds that could bias our estimated information effects downward. First, spillovers between control and treatment households could cause us to underestimate the information effects. Second, a lack of information sharing within the household could mean that we increased the knowledge of some but not all household members. We present extensive evidence showing that neither of these can account for the small information effects we find.

Table 6
Survey effects.

	(1)	(2)	(3)
	Payment amount	Payment propensity	Consumption
Control	0.076 (0.131)	0.012 (0.024)	0.029 (0.094)
N	985	985	985

Notes: The table estimates survey effects by comparing the control group in our study (Control = 1) to 500 randomly selected households who did not participate in our study. Each column corresponds to a different regression. The column headings give the dependent variable, measured in March (the first month after the follow-up survey). Every regression controls for sampling strata indicators and the pre-survey value of the dependent variable. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

4.2.2. Scrutiny

In our setting, household visits by utility employees are not unusual, and less than 1 percent of households refused to participate in the education visits. This makes it unlikely that the visits were perceived as unwanted scrutiny or as part of a crackdown on nonpayment.

The ideal experiment to formally test whether a perceived threat of enforcement could explain our findings would require the same education officers to visit the households without delivering the education program. Since the officers would still have to do *something* during the visit, it is far from obvious how such an experiment could be designed. A practical alternative is to use our treatment but study the neighbors of our treated households. These consumers saw Odi officers visiting their neighbor's house without seeing the details of the visit. Thus, they were simply exposed to a "treatment" of increased presence of the provider's employees. If our information campaign increased payments through an increase in perceived scrutiny, then this "treatment" of the neighbors should have had a similar effect.

The data shows no evidence of this. In Panel A of Table 7 we restrict the sample to the control group, and regress payment on an indicator equal to 1 if a household has a treated neighbor within a 30 meter radius (measured using GPS coordinates).²⁶ The effect of this "treatment" on payment is *negative* and mostly insignificant: households whose neighbor was visited by our officers did not increase their payments. We obtain similar results when looking at a radius of 20, 30, 40, or 50 meters. Consumers who experienced an increased presence of the provider's employees without receiving the information campaign did not increase payments.

An alternative approach to checking whether a perceived threat of enforcement could explain our findings is to proxy for a household's risk of being subject to enforcement. If increased scrutiny was driving the results, the increased payment should come from households who have a higher risk of being punished. Recall that, in our context, the primary enforcement mechanism of the provider is to install a restriction device on a consumer's water service. Therefore households who are not yet restricted should perceive a higher risk of enforcement action than households who are already restricted. If increased scrutiny was driving the payment results, we would expect the treatment to operate primarily among non-restricted households.

We do not find this to be the case. In Panel B of Table 7 we interact the treatment indicator with the household's restriction status at baseline (1 if restricted). None of the interactions are statistically significant, and their coefficients are always positive. The increased payments did not come from those who are not yet restricted and therefore would have more reason to respond to increased scrutiny by the provider.

²⁶ Anticipating that neighbors' behavior could be important, we collected GPS coordinates of each household's location during the surveys. Although each property has a street address used for mail delivery, there is no official map of our study area that would contain these addresses. House numbers often follow each-other in surprising orders. Thus, GPS coordinates were the only way to map these households.

Table 7
Scrutiny effects on payment.

	(1)	(2)	(3)
	Payment amount	Payment propensity	Payment frequency
<i>Panel A</i>			
Neighbor visited	-0.185 (0.351)	-0.033 (0.060)	-0.194* (0.106)
N	479	479	479
<i>Panel B</i>			
Treatment	0.180 (0.151)	0.028 (0.026)	0.069 (0.057)
Restricted	-0.644 (0.400)	-0.122* (0.070)	-0.159 (0.110)
Interaction	0.243 (0.288)	0.036 (0.050)	0.065 (0.097)
N	966	966	966

Notes: Regressions in Panel A are on the control group only. 'Neighbor visited' is 1 if there is a treatment household within a 30 meter radius. Regressions in Panel B are on the full sample. 'Restricted' is 1 if the consumer was restricted at baseline. The columns in each panel correspond to separate regressions. The column headings give the dependent variable. 'Payment amount' is total payment in the 3 months following the treatment in logs; 'Payment propensity' is 1 if the household made a payment during this period, and 'Payment frequency' is the number of payments made. All regressions control for sampling strata indicators and the value of the dependent variable during the 3 months prior to the treatment. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

4.2.3. Reciprocity

As described in Section 2.4, a third possible channel to explain the increased payments is consumers' reciprocity towards the provider. While obtaining direct evidence on this channel seems especially difficult, we did ask treated households in the follow-up survey whether they had found the information brochures very useful / fairly useful / a little useful, or not at all useful. Given our finding that households gained at most limited information from the education visits, we might expect the majority of respondents to view the brochures as not very useful. Surprisingly, we find the opposite: 89.0% said the brochures were very useful or fairly useful, and only 2.6% that they were not at all useful. Thus, consumers were appreciative of the education campaign in spite of the modest information effects. One possible explanation is that brochures were valued for reasons that our surveys were not designed to measure. For example, one respondent told us that she thought the colorful brochures would help her convince her children that it was important to conserve water at home. Alternatively, households may have appreciated the provider's efforts in attempting to do something useful for the community. One respondent told us that although he himself already knew most of what was in the brochures, he thought this was a very valuable service to others in the township.

Under reciprocity, it is not surprising that our surveys had no effect, or that visits to a neighbor's house did not result in increased payments, as we saw above. Because the treatment was a one-time intervention, reciprocity can also explain why the payment response was short-lived, and suggests that a sustained campaign may result in longer term effects.²⁷ Thus, while our measures of consumers' feelings are necessarily limited, reciprocity provides a consistent explanation for the patterns observed in the data.

5. Heterogenous treatment effects

In Section 3 we found some heterogeneity in our treatment effects by quantity of water consumed. In this section, we explore other potential dimensions of heterogeneity. This is relevant because heterogenous treatment effects across subgroups could in principle reconcile our findings with information being the dominant mechanism behind our

²⁷ A short-lived response is difficult to square with some of the other mechanisms discussed above. For example, if consumers had changed their behavior due to new information or perceived scrutiny, we would expect more long-lasting treatment effects.

intervention. For example, suppose that our information campaign raised knowledge among the less educated but not among the highly educated (who were more knowledgeable to begin with) and that the less educated made large payments as a result. Then as long as behavior was sufficiently responsive to information in this subgroup, the lack of an average treatment effect on information can be consistent with the large effects on nonpayment that we found.

We focus on six dimensions of heterogeneity: household income, education, baseline knowledge, restricted status at baseline, indigent status at baseline, and payment behavior before the treatment. The last four of these variables were also used in our stratified sampling procedure because they are natural candidates for determinants of households' ability or willingness to respond to our treatment. We add income, education and baseline knowledge because they are obvious dimensions of heterogeneity in the context of an information campaign and payment behavior. We focus on the income, education and baseline knowledge results in the main text and present the rest of the analysis in the Online Appendix.

Table 8 presents estimates from regressions that interact the treatment indicator with one of our variables measuring heterogeneity. Columns 1–3 focus on payment and consumption and columns 4–10 study our various information measures. In Panel A the payment effects of our treatment come mainly from higher income consumers, although the effects are imprecisely estimated. This makes sense since higher income consumers can more easily adjust their payments. At the same time, we find no evidence that these consumers experienced large improvements in knowledge. Higher income consumers also did not reduce their consumption significantly in response to the treatment. Thus, our earlier findings hold up in these subgroups.

In Panels B and C we find some evidence of heterogeneous effects on information based on education levels and baseline knowledge. In Panel B, the less educated show some evidence of increased knowledge

relative to the control group. For example, these households are more likely to use the word 'kiloliter' (column 4) and more likely to be able to tell their consumption from the bill (column 6). In Panel C, we find similar patterns when we break up the treatment effects by whether the respondent had a good sense of the water tariff at baseline. Less knowledgeable households show some evidence of improved information. Can this improved information explain the payment results? The answer seems to be 'no.' Columns 1 and 2 show that it is in fact the *more* educated who account for the increased payments we find. Similarly, we estimate larger payment effects for those who had a better understanding of tariffs (although these are again estimated imprecisely). Thus, while our treatment shows a larger impact on the information of specific groups, this change in information is not responsible for the reduction in nonpayment.

6. Conclusion

We implemented and evaluated an information campaign as a potential response to nonpayment for water in South African townships. Our education visits had a substantial impact in the short run, reducing the fraction of households making no payments by 4 percentage points and increasing the amount of payments by approximately 25% over a three-month period. This provides evidence that strategies other than increased enforcement can lower nonpayment. The evidence shows some reduction in water use among the highest consumers, but we find no treatment effect on average consumption. Because consumption serves as an indicator of the household's planned behavior, our simple model suggests that on average nonpayment in this population is largely unplanned.

Using direct measures of households' knowledge, we can rule out the treatment operating through improvements in information. On average, treated households are not much more likely to understand quantities of

Table 8
Heterogenous treatment effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Payment amount	Payment propensity	Consumption	Response in kl	Reads consumption from bill	Consumption accurate	Tariff in ballpark	Tariff error	Increasing tariff	Quiz score
<i>Panel A: Income</i>										
Treatment	0.047 (0.193)	0.003 (0.034)	0.044 (0.047)	0.034 (0.030)	0.074 (0.053)	0.069* (0.039)	-0.002 (0.020)	-18.862 (20.248)	-0.046 (0.047)	-0.002 (0.098)
Interaction	0.428 (0.279)	0.067 (0.048)	-0.053 (0.065)	0.008 (0.043)	-0.060 (0.079)	-0.053 (0.050)	-0.008 (0.031)	8.719 (23.164)	0.084 (0.064)	0.085 (0.139)
N	857	857	857	846	657	657	735	351	855	856
<i>Panel B: Education</i>										
Treatment	-0.017 (0.187)	-0.000 (0.033)	-0.005 (0.044)	0.102*** (0.033)	0.104* (0.057)	0.085** (0.039)	-0.037* (0.022)	-29.196 (19.100)	-0.045 (0.044)	-0.015 (0.096)
Interaction	0.479* (0.258)	0.068 (0.045)	0.018 (0.060)	-0.117*** (0.041)	-0.115 (0.075)	-0.080 (0.050)	0.038 (0.031)	43.087 (29.205)	0.034 (0.060)	0.115 (0.131)
N	960	960	960	948	728	728	816	393	958	959
<i>Panel C: Tariff in ballpark</i>										
Treatment	0.082 (0.148)	0.012 (0.025)	-0.026 (0.035)	0.069*** (0.024)	0.052 (0.044)	0.043 (0.029)	-0.031* (0.017)	-1.169 (15.672)	-0.030 (0.036)	0.081 (0.077)
Interaction	0.288 (0.522)	0.059 (0.086)	0.204** (0.097)	-0.177*** (0.050)	-0.130 (0.132)	-0.050 (0.089)	0.031 (0.057)	-55.641 (76.368)	-0.000 (0.108)	0.026 (0.215)
N	776	776	776	764	578	578	667	328	775	775
<i>Panel D: Quiz score</i>										
Treatment	0.484 (0.436)	0.089 (0.073)	0.004 (0.101)	0.184** (0.073)	-0.122 (0.122)	0.056 (0.091)	-0.031 (0.049)	-64.095 (52.105)	-0.035 (0.100)	0.005 (0.210)
Interaction	-0.147 (0.156)	-0.028 (0.027)	-0.001 (0.038)	-0.053* (0.027)	0.064 (0.045)	-0.007 (0.034)	0.001 (0.018)	21.105 (20.878)	0.004 (0.038)	0.026 (0.078)
N	776	776	776	764	578	578	667	328	775	775

Notes: Panels A-D investigate heterogenous treatment effects by income, education, and baseline knowledge. Income is total household income at follow-up. Education is 1 if the follow-up respondent has completed high school and 0 otherwise. The information measures 'Tariff in ballpark' and 'Quiz score' in Panels C and D are measured at baseline. The columns in each panel correspond to separate regressions. The column headings give the dependent variable. 'Payment amount' is total payment in the 3 months following the treatment in logs; 'Payment propensity' is 1 if the household made a payment during this period. 'Consumption' is average consumption in the 3 months following the treatment (in logs). All regressions control for sampling strata indicators and the baseline value of the dependent variable. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

water used or their water bill than households in the control group. Although we see larger increases in information among the less educated, the reduction in nonpayment is not driven by these changes. A consistent explanation for the patterns seen in the data is provided by households reciprocating the provider's efforts by paying more. These findings show that public information campaigns can change behavior through other channels besides increased information.

Appendix A

Table 9

Information measures.

Measures	Description
Response in kl	1 if the respondent's guess about their consumption is stated in kiloliters
Reads consumption from bill	1 if the respondent reads out his consumption from the bill, 0 if he cannot, missing if he did not find a water bill
Consumption accurate	1 if this number matches any consumption in the administrative data from the prior 6 months, 0 if not
Tariff in ballpark	1 if the respondent's guess about the kiloliter price is between 5–25 Rand
Tariff error	$\max(0, \text{the respondent's guess about kiloliter price} - 25)$
Increasing tariff	1 if the respondent understands that the tariff schedule is increasing
Quiz score	Number of correct answers to 'quiz' questions on water quantities

Appendix B. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jdeveco.2015.06.002>.

References

- Abrahamse, W., Steg, L., Vlek, C., Rothengatter, T., 2005. A review of intervention studies aimed at household energy conservation. *J. Environ. Psychol.* 273–291.
- Ahluwalia, M.S., 2002. Economic Reforms in India since 1991: Has Gradualism Worked? *J. Econ. Perspect.* 16 (3), 67–88.
- Allcott, H., 2011. Social norms and energy conservation. *J. Public Econ.* 95, 1082–1095.
- Allcott, H., Rogers, T., 2014. The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation. *Am. Econ. Rev.* 104 (10), 3003–3037.
- Alm, J., Bahl, R., Murray, M.N., 1991. Tax Base Erosion in Developing Countries. *Econ. Dev. Cult. Chang.* 39 (4), 849–872.
- Alm, J., McClelland, G.H., Schulze, W.D., 1992. Why do people pay taxes? *J. Public Econ.* 48, 21–38.
- Alm, J., Jackson, B.R., McKee, M., 1993. Fiscal exchange, collective decision institutions, and tax compliance. *J. Econ. Behav. Organ.* 22, 285–303.
- American Water Works Association, 1999. Residential End Uses of Water. AWWA Research Foundation, Denver, CO.
- Anderson, M.L., 2008. Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *J. Am. Stat. Assoc.* 103 (484), 1481–1495.
- Ayres, I., Raseman, S., Shih, A., 2013. Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage. *J. Law Econ. Org.* 29 (5), 992–1022.
- Bazart, C., Bonein, A., 2014. Reciprocal relationships in tax compliance decisions. *J. Econ. Psychol.* 40, 83–102.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate. *J. R. Stat. Soc. Ser. B* 57, 289–300.
- Benneer, L.S., Olmstead, S.M., 2008. The impacts of the "right to know": Information disclosure and the violation of drinking water standards. *J. Environ. Econ. Manag.* 56, 117–130.
- Chetty, R., Saez, E., 2013. Teaching the tax code: Earnings responses to an experiment with EITC recipients. *Am. Econ. J. Appl. Econ.* 5 (1), 1–31.
- Cole, S., Sampson, T., Zia, B., 2011. Prices or Knowledge? What Drives Demand for Financial Services in Emerging Markets? *J. Financ.* 66 (6), 1933–1967.
- Duflo, E., Saez, E., 2003. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Q. J. Econ.* 118, 815–842.
- Fehr, E., Gächter, S., 2000. Fairness and retaliation: The economics of reciprocity. *J. Econ. Perspect.* 14, 159–181.
- Ferraro, P.J., Miranda, J.J., 2013. Heterogeneous treatment effects and mechanisms in information-based environmental policies: Evidence from a large-scale field experiment. *Resour. Energy Econ.* 35, 356–379.
- Ferraro, P.J., Price, M.K., 2013. Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment. *Rev. Econ. Stat.* 95 (1), 64–73.
- Glaeser, E.L., Ujhelyi, G., 2010. Regulating Misinformation. *J. Public Econ.* 94 (3–4), 247–257.
- Gordon, R., Li, W., 2009. Tax structures in developing countries: Many puzzles and a possible explanation. *J. Public Econ.* 93, 855–866.
- Hallsworth, M., List, J.A., Metcalfe, R.D., Vlaev, I., 2014. The Behavioralist As Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance. NBER Working Paper No. 20007.
- Jalan, J., Somanathan, E., 2008. The importance of being informed: experimental evidence on demand for environmental quality. *J. Dev. Econ.* 87, 14–28.
- Jesoe, K.K., Rapson, D., 2014. Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use. *Am. Econ. Rev.* 104 (4), 1417–1438.
- Karlan, D., Morten, M., Zinman, J., 2013. A personal touch: Text messaging for loan repayment. Working paper.
- Karlan, D., McConnell, M., Mullainathan, S., Zinman, J., 2014. Getting to the Top of Mind: How Reminders Increase Saving. *Manag. Sci.* (forthcoming).
- Kumler, T., Fras, J.A., Verhoogen, E., 2013. Enlisting Employees in Improving Payroll-Tax Compliance: Evidence from Mexico. NBER working paper no. 19385.
- Lampietti, J.A., Banerjee, S.G., Branczik, A., 2007. People and Power: Electricity Sector Reforms and the Poor in Europe and Central Asia. The World Bank, Washington, D.C.
- Luttmer, E.F.P., Singhal, M., 2014. Tax Morale. *J. Econ. Perspect.* 28 (4), 149–168.
- Madajewicz, M., Pfaff, A., van Geen, A., Graziano, J., Hussein, I., Momotaj, H., Sylvi, R., Ahsan, H., 2007. Can information alone change behavior? Response to arsenic contamination of groundwater in Bangladesh. *J. Dev. Econ.* 84, 731–754.
- McRae, S., 2015. Infrastructure Quality and the Subsidy Trap. *Am. Econ. Rev.* 105 (1), 35–66.
- Pomeranz, D., 2015. No taxation without information: deterrence and self-enforcement in the value added tax. *Am. Econ. Rev.* (forthcoming).
- Reiss, P., White, M.W., 2008. What Changes Energy Consumption? Prices and Public Pressures. *RAND J. Econ.* 39 (3), 636–663.
- Republic of South Africa, 2011. Intergovernment Fiscal Review 2011. National Treasury, Pretoria, South Africa.
- Slemrod, J., 2007. Cheating Ourselves: The Economics of Tax Evasion. *J. Econ. Perspect.* 21 (1), 25–48.
- Stango, V., Zinman, J., 2014. Limited and Varying Consumer Attention Evidence from Shocks to the Salience of Bank Overdraft Fees. *Rev. Financ. Stud.* 27 (4), 990–1030.
- Szabó, A., 2015. The Value of Free Water: Evaluating South Africa's Free Basic Water Policy. *Econometrica* (forthcoming).
- Torgler, B., 2005. Tax morale in Latin America. *Public Choice* 122, 133–157.
- White, I.R., Thompson, S.G., 2005. Adjusting for partially missing baseline measurements in randomized trials. *Stat. Med.* 24, 993–1007.
- World Bank, 1999. Non-Payment in the Electricity Sector in Eastern Europe and the Former Soviet Union. World Bank Technical Paper No. 423. The World Bank, Washington, DC.
- Zwane, A.P., Zinman, J., Van Dusen, E., Pariente, W., Null, C., Miguel, E., Kremer, M., Karlan, D.S., Hornbeck, R., Gine, X., Duflo, E., Devoto, F., Crepon, B., Banerjee, A., 2011. Being surveyed can change later behavior and related parameter estimates. *Proc. Natl. Acad. Sci.* 108 (5), 1821–1826.