From Rationed to Rational?
Improving Households' Water Usage through Education in South Africa

Report to the

WATER RESEARCH COMMISSION

by

Patrick Chiroro, (Impact Research International)

Andrea Szabo, (University of Houston)

Gergely Ujhelyi, (University of Houston)

WRC Report No. KV 332/14

ISBN No. 978-1-4312-0553-0

June 2014

Water Research Commission
Private Bag X03
Gezina, 0031

orders@wrc.org.za or download from www.wrc.org.za

DISCLAIMER

This report has been reviewed by the Water Research Commission (WRC) and approved for publication. Approval does not signify that the contents necessarily reflect the views and policies of the WRC, nor does mention of trade names or commercial products constitute endorsement or recommendation for use.

EXECUTIVE SUMMARY

BACKGROUND

The project described here implements and evaluates a water education programme among households in a group of South African township. The programme draws on education materials of the Water Research Commission to provide ordinary households with information on the water consumption process, including how to read their water meter and their bill, and how much water typical everyday activities use.

RATIONALE

In our study area, like in many parts of South Africa and other countries, the water market exhibits several anomalies, especially on the consumer side. Many households apparently waste water while accumulating large bills that they have difficulty paying. Responses by the water provider, such as installing restriction devices, are socially costly. In this study, we ask whether an information campaign can lead to improved water management and reduced non-payment by the households.

OBJECTIVES AND AIMS

The project was designed to answer the following questions:

AIM 1 Can providing information change households' water consumption (e.g., induce conservation)? How large is the change in consumption that can be achieved?

AIM 2 Does providing information affect payment behaviour (e.g., the incidence of non-payment)?

AIM 3 Are the specific education materials used effective at transmitting information to households?

METHODOLOGY

A unique feature of the project is an emphasis on a methodologically sound implementation that allows measuring the *causal effects* of the education programme. The education programme is implemented as a Randomised Controlled Trial, with a "treatment" and a "control" group of 500 households each.

RESULTS AND DISCUSSION

Aim 1

Our education campaign caused an increase in households' self-reported conservation practices. As expected, we find no change in households' average consumption, but see a decrease of about 10% for the largest consumers. These patterns are consistent with a positive effect on conservation, with some households substituting more water-intensive activities with less-water-intensive ones.

Aim 2

Our education campaign caused a substantial increase in payments. In response to our programme, households increased their payments by 25-30% over a three-month period, and the incidence of non-payment declined by 4-5%.

Aim 3

Our education campaign caused very little change in households' knowledge. As one example, even after the campaign over 88% of households could not tell their water consumption from their water bill. Thus, the large effects of our campaign on consumption and payment were not caused by improvements in household information. After ruling out several possible explanations, we conclude that the education likely elicited psychological responses from the households. Consumers may have felt like they ''should'' conserve more and pay their bills, perhaps as an expression of reciprocity for the provider's efforts in reaching out to households through the campaign.

CONCLUSIONS

Water education campaigns like the one analysed here may be an effective policy to increase conservation and reduce non-payment even if they are not effective at improving households' knowledge.

In the long-run, however, providing consumers information – and ensuring that the information they already receive on their billing statement is understandable to them – is likely to be important to improve households' water management.

RECOMMENDATIONS FOR FUTURE RESEARCH

More research is needed to understand how information may be transmitted in a more effective way. Future research should also compare the effectiveness of information provision and "social pressure" campaigns (such as comparing households' consumption to their neighbours' or giving them explicitly prescriptive messages about what is right or wrong to do). Randomised controlled trials provide a powerful way to evaluate the causal effects of such programmes.

ACKNOWLEDGEMENTS

Financial support from the Water Research Commission, the University of Houston, and Odi Water Services is gratefully acknowledged. We are grateful to Sipho Nkosi and his dedicated team for implementing the education programme and to Pieter Avenant for giving us access to the administrative dataset. We thank Inga Jacobs and the WRC for useful comments and suggestions on this report.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

DD: Difference in Differences

KL: kilolitre

MDE: Minimum detectable effect

ODI: Odi Water Services

OLS: Ordinary Least Squares

RCT: Randomised Controlled Trial

VAR: Variance

WRC: Water Research Commission

1 INTRODUCTION

The project described here implements and evaluates a water education programme among households in a group of South African townships. It is conducted in the Mabopane area, north of Pretoria, which is special in that a disadvantaged population similar to that found in many developing areas of the world is served by a state-of-the-art water infrastructure.

The programme draws on education materials of the Water Research Commission to provide ordinary households with information on the water consumption process, including how to read their water meter and their bill, and how much water typical everyday activities use.

The project is designed to answer the following questions: Can providing information change water consumption, and by how much? Does providing information affect payment behaviour (e.g., the incidence of non-payment)? Are the specific education materials used effective at transmitting information to households?

To facilitate a methodologically sound evaluation of these questions, the education programme is implemented as a Randomised Controlled Trial (RCT). Information is given to 500 randomly selected households (the “treatment group”) and their outcomes are compared to another group of 500 households (the “control group”).

While RCT’s are the gold standard in evaluating a variety of policy interventions in both developed and developing countries around the globe, to our knowledge this project is the first to use this methodology to evaluate a water education programme anywhere outside of the most developed countries. We therefore discuss the methodology in detail in Chapter 2, with the hope of providing a blueprint for future water policy interventions in South Africa and elsewhere.

The remainder of this report is organised as follows. Chapter 1 puts our project in context and relates it to the previous literature. Chapter 2 explains the methodology. Chapter 3 describes the intervention, Chapter 4 gives details on the sampling procedure and the implementation of the project, and Chapter 5 presents our hypotheses. Chapter 6 summarises the data, and Chapter 7 presents the main analysis and results. Chapter 8 explores the robustness and possible explanations of our findings. Chapter 9 studies heterogeneous treatment effects, and Chapter 10 concludes.

2 RESEARCH SETTING AND RELATED LITERATURE

2.1 Research setting and motivation

In late 2012, we implemented a water education programme among the consumers of Odi Water; a provider serving approximately 40,000 residential consumers in the Mabopane area, a group of townships located approximately an hour's drive north of Pretoria (**Error! Reference source not found.**). This area has a well-functioning water infrastructure developed in the mid-1990s as part of government efforts to develop disadvantaged areas after Apartheid. 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. 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 **Error! Reference source not found.**).

Unpaid balances accrue interest, and the provider restricts the water supply of the worst offenders (this is done by limiting the water flow to a bare minimum). Clearly, waste and non-payment are costly both to the households and to the water utility. Why do these behaviours 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 cause. Information is widely available in a format that most consumers from Western countries would consider standard (water meter on the property, detailed monthly bills, and a customer service department to answer questions). However, as we show below, households exhibit very little familiarity with the meaning of the numbers on the meter and the units in which their water consumption is being measured. There is also a lack of knowledge about the consumption process, e.g., how much water is used in various everyday activities. While some households might *choose* to consume excessive amounts of water, this is clearly not the case for most.

To study whether information can improve efficiency in this market, we designed and implemented a water education campaign. The programme drew on education materials of the Water Research Commission to provide households with information on the water consumption process, including how to read their water meter and their bill, and how much water typical everyday activities use. Implementation of the programme was randomised following the methodology of Randomised Controlled Trials (RCT) described in Chapter 3 below. A treatment group of 500 households took part in the education, while another group of 500 households serve as the control group. The programme was preceded by a baseline survey and followed by a follow-up survey. Information from these surveys was combined with billing data provided by Odi Water to measure the effects of the programme. We are interested in whether our education campaign affects consumption, payment, and households' knowledge regarding the water consumption process.

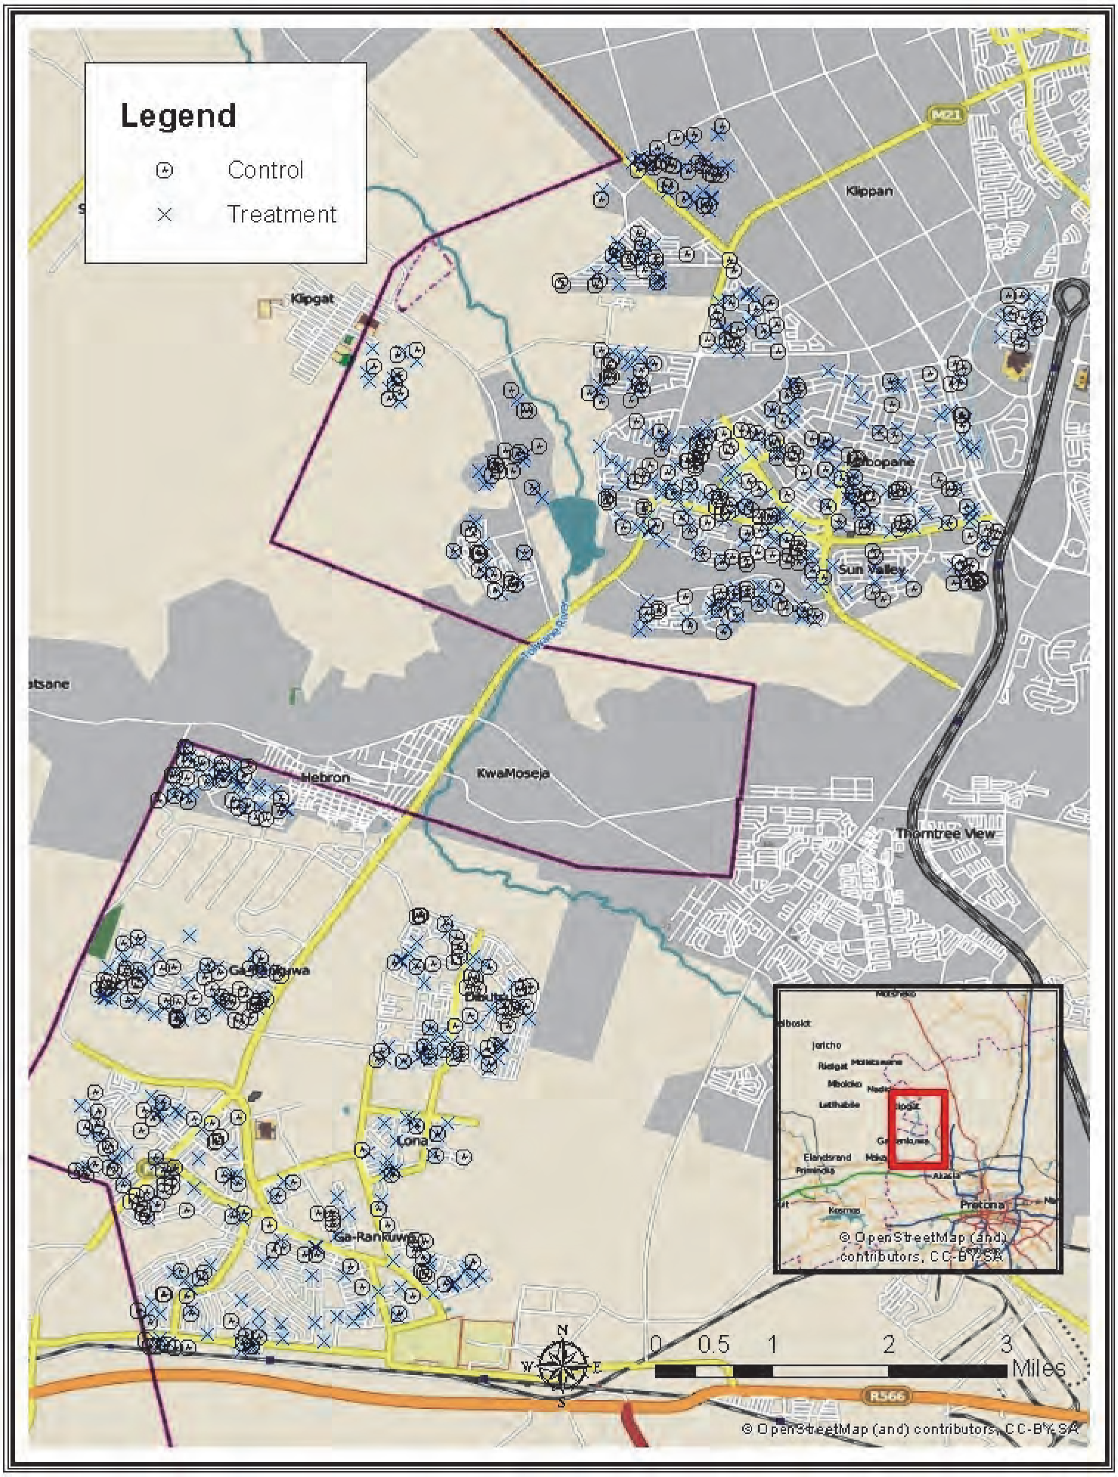

Figure 1 Study area and participating households

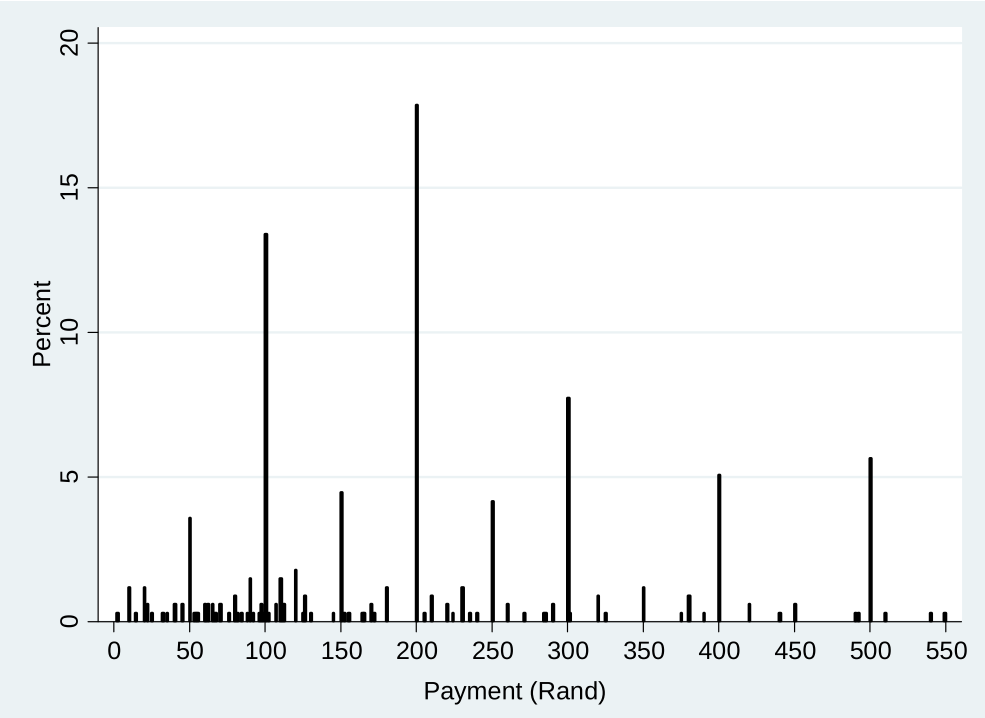

Figure 2 Distribution of payments, August 2012

2.2 Relation to previous literature

Previous analyses of conservation campaigns come almost exclusively from the US. In economics, this includes Allcott (2011), Ayres et al. (2013), Ferraro and Price (2013), and Jessoe and Rapson (2013). A related literature in marketing and psychology is reviewed by Abrahamse et al. (2005).[1] Most of these studies focus on electricity.[2]

While some campaigns emphasise psychological incentives (e.g., by providing social comparisons about a household's energy use relative to its neighbours, or including pictograms such as smiling or frowning faces on the bill), others focus on providing information (e.g., in the form of useful tips for water conservation, or real-time feedback on quantity consumed). Our work complements these studies by making progress towards understanding the mechanisms behind these programmes. In particular, we highlight the fact that even pure information campaigns can give rise to psychological effects.

The developed country context in which the studies cited above were conducted means that non-payment was not a major issue. Thus, these papers focus exclusively on consumption. In developing countries, while improving people's access to basic utilities like

[1] Many of the studies in the latter group do not offer credible identification – i.e., the campaign is not randomised, and the analysis is restricted to simple cross-sectional or time-series comparisons.
[2] One exception is a recent study on water from Australia by Fielding et al. (2013).

electricity or water is recognised as a key challenge, 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, 2013). In South Africa, non-payment presents a major problem for local governments and prevents the efficient use of the existing infrastructure for electricity, water, and sanitation (Hollingworth et al., 2011; Republic of South Africa, 2011). To our knowledge, our research is the first to study whether an information campaign can lower non-payment.

More generally, our focus on education complements studies of prices as a policy tool to achieve socially desirable changes in consumption – see, e.g., Reiss and White (2008) on the US, and Szabó (2013) or Hosking et al. (2011) for studies in the South African context.

3 RANDOMISED CONTROLLED TRIALS

In designing policy interventions, decision makers face a fundamental challenge: we would like to know the effect of the intervention before it is actually carried out. Will the policy change people's behaviour in the desired direction? How big a change can we expect? Are the benefits of the intervention going to be worth the cost of implementing it? Answering these questions is difficult before the intervention is actually implemented. While impact studies based on similar interventions elsewhere might be useful, these may not be directly applicable to the location, time period, population, or intervention under consideration. For example, an intervention could work in one city but not in another one.

The typical solution to this problem is a *pilot study* implementing the intervention on a small scale in the target area. This chapter describes how such a pilot study should be organized and evaluated in order to be informative for the policy being considered. It describes the methodology of Randomised Controlled Trials (RCT), the best-practice method of conducting and evaluating pilot studies in virtually any area of public policy.[3]

The basic principle of RCT's is simple. Randomly select 2 groups in the target population, for example, two groups of 500 water consumers each. Carry out the pilot intervention in one of these groups (the "treatment" group), but not the other (the "control" group). For example, give one of the groups tips on how to save water. Finally, measure the outcomes of interest in both groups and compare them to each other. For example, compare

[3] For a nontechnical discussion of RCT's and a long list of examples of policy interventions where the methodology has been applied, see Banerjee and Duflo (2012). For a more technical overview, see Duflo et al. (2008).

the amount of water consumed in the two groups. The difference between the two groups' outcomes is the effect of the intervention.

3.1 Why Randomise?

Imagine that, instead of the above procedure, we gave the water saving tips to the 500 highest consumers. This seems to make sense, since these are the households where saving water would be especially important. But how will we know if the water saving tips are effective? Perhaps we can compare the consumption of these 500 households before and after the intervention. If we find that consumption decreased by 1 kilolitre on average, is this the effect of the intervention?

The answer is 'no.' Consumption before and after the intervention might be different for reasons that have nothing to do with the intervention. The most obvious reason is that consumption varies over time. If we measure consumption in December, do the intervention in January, and measure consumption again in February, consumption in February is likely to be different from consumption in December not only because of the water saving tips, but simply because it is February. For example, people may spend more time at home during the holidays in December, and consumption may be lower in February simply because everyone is at work. In this case, there is no way to isolate the effect of the intervention from the passing of time.

If comparing the 500 highest consumers before and after the intervention doesn't work, perhaps we can compare them to another group after the intervention. Is consumption among the 500 highest consumers in February, after the intervention, lower than consumption among other households?

Clearly, this will not give us the effect of the intervention either. Since we are comparing the highest consumers to others, their consumption is likely to be higher regardless of the intervention. Even if the intervention was very successful and consumption among the 500 is now close to that of everyone else's, we won't know how big the change is that is due to the intervention.

The solution then is to go one step further, and instead of doing the intervention for the 500 highest consumers, select the households randomly. Because we are randomising, the 500 households receiving the water saving tips will be similar on average to the rest of the population. If we did not do the intervention, and compared their consumption to another randomly selected 500 households, the average difference would be close to 0. Thus, if by doing the intervention we find differences in the consumption of these two groups, this difference must be due to the intervention. If comparing consumption in the two random groups after one of them received the water saving tips shows a difference of 1 kilolitre, then this is exactly the effect of our policy intervention.

3.2 Advantages of RCT Compared to Other Methods

There are other ways to try and measure the effect of an intervention. However, there is typically little reason to consider them unless randomisation is not possible for some reason (for example, if the intervention to be evaluated has already happened in the past without being randomised).

Consider the above example where, instead of randomising the water saving tips, the local water provider already gave them to the 500 highest consumers. How can we measure the effect of this intervention?

One option is if we have information on the outcomes of interest in the population before and after the intervention. For example, we may know how much water households consume. Then we can perform the following comparison, called the Difference-in-Differences (DD) method. Take the difference of the consumption among the 500 treated households before and after the intervention. Do the same for another group of households selected (randomly or non-randomly) from those who did not receive the water saving tips. Compare the two differences – in other words, ask whether consumption among the 500 selected households is more or less similar to the rest of the population after the intervention than it was before. For example, we may find that while households who did not receive the tips increased their consumption by 3 kl, the 500 households who received them increased consumption by only 2 kl. In this case, under certain assumptions, the difference of 1 kl between these differences measures the effect of the intervention. Households reduced their consumption by 1 kl in response to the intervention.

It is important to note that this conclusion is only valid under specific assumptions. In particular, it is only true if in the absence of the intervention the two groups (the treatment and the comparison group) would have had similar changes in consumption. For example, suppose that the increase of 3 kl among the control group is due to the fact that it was winter before the intervention but it is summer after it. Then the question is: would the 500 highest consumers also increase their consumption by 3 kl between winter and summer if we did not give them the water saving tips? If the answer is 'yes,' the DD method is useful in measuring the impact of the intervention. But if, for example, high consumers are households who are always wasting water, then it is possible that they waste almost as much during winter as they do during summer. In this case, their water consumption may not change as much between the two seasons as that of other households, who manage their consumption more carefully. Then, the DD method simply measures this smaller difference due to waste, rather than the effect of the intervention.

3.3 Inference

When an RCT is carried out correctly, estimating the impact of the intervention is a simple matter of comparing averages. Imagine that we find that the 500 randomly selected consumers who received the water saving tips have an average consumption of 12 kl, while the control group has a consumption of 13 kl. Then our estimate of the effect of the intervention is 13 - 12 = 1 kl.

Of course, if we were to repeat this intervention with a different group of 500 customers, we may not get an effect of exactly 1 kl. To draw conclusions about the effect of the intervention on the population of consumers requires some basic statistical considerations. Suppose that, in the above example, the variance of consumption in the treatment group is 16 kl, while the variance in the control group is 36 kl. Then we can compute the *95% confidence interval* for the effect of this intervention: it is [0.37,1.63].[4] Roughly speaking, there is a 95% probability that repeating the intervention on a different group of 500 households in this population will give us an effect between 0.37 and 1.63 kl.

It is possible that the 95% confidence interval will include 0. For example, if the difference between the average consumption in the treatment group is not 12 kl but 12.5 kl, with a variance of 16 kl, then the confidence interval for the difference is [-0.13, 1.13]. Because this includes 0, we cannot rule out at the 95% confidence level that if we repeat the intervention in a different group, its effect will be 0 (or even negative). This would suggest that something might have to be changed in the intervention to make it effective.

In this manner, the results from the pilot study can be used to obtain an estimate of the likely effect of doing the intervention on a large scale. If the results suggest that the policy may be ineffective, carrying out the intervention may not be worthwhile without modifications.

[4] This can be computed using the formula $D \pm t \times \sqrt{(V / N)}$, where D is the measured difference in the outcome (D = 1 in this case), V is the sum of the two variances (V = 16 + 36), N is the size of each group (N = 500), and t is the relevant value from a Normal distribution (for a 95% confidence level, t = 1.96).

3.4 Practical Considerations

Randomisation

In practice, the best way to select random groups is by using a random number generator available in common computer software, such as MS Excel.[5] This will assign random numbers to a list of households, and one can pick the 500 households who are assigned the lowest numbers to obtain a random group of 500 households.

Other methods of randomising can be used in some situations, but one must consider carefully whether a specific method will yield a truly random group. For example, imagine using consumers' water account numbers. One could order consumers, and pick the first 500, for example. However, we must consider how account numbers were assigned to consumers in the first place. For example, if assignment was based on location, then the first 500 numbers may all correspond to a particular street or block, and we will not have a random group. Or lower account numbers can mean that consumers have lived on the property for a longer period, while newer consumers are the ones with higher numbers. Again, picking the first 500 would mean a non-random group: these consumers may be older or have bigger families, and may thus use water differently than others.

How big should the groups be?

Choosing the size of the treatment and control groups involves some basic statistical considerations (besides of course the budget available for the intervention).

Suppose the intervention is actually ineffective (i.e., has no effect on consumers). How sure do we want to be that our pilot study does not lead us to conclude that it is in fact effective? This is called *confidence level*, and is typically set at 90 or 95%. If at the given level of confidence we can reject the hypothesis that the intervention is ineffective, we say that the measured effect of the intervention is *significantly* different from 0, or, for short, that it is *significant*.

Conversely, imagine that the intervention is in fact effective at changing consumers' behaviour. How sure do we want to be that the pilot study indeed leads us to conclude that it is effective? In other words, how sure do we want to be that the effect of the pilot will be significantly different from 0? This is called *power*, and is typically set at 80 or 90%.

[5] In Excel, this can be accomplished using the "RAND()" function.

The question then is how large of an effect do we need to make the actual large-scale intervention meaningful. For example, suppose that we knew for certain that *if* we carried out the intervention being planned, the actual effect would be 1 kl for each household. Would we then decide to implement it? How about if we knew for certain that the effect would be 0.5 kl? What if it was 0.2 kl? This is important, because we want to make sure that an effect that we would consider meaningful or important in this sense is actually detected by the pilot study. In other words, if a true effect of 1 kl would be considered a huge success, we need to make sure that if the true effect is in fact 1 kl; our pilot study does not yield an insignificant result because our sample size was too low.

Let *MDE* (or *minimum detectable effect*) stand for the lowest true effect that the pilot study will allow us to detect (i.e., categorize as being significant). Let *VAR* stand for the variance of the outcome being considered (for example, the variance of consumption among all consumers). If we use an equal sized treatment and control group, a confidence level of 95%, and a power of 80%, the required sample size *N* is given by the following formula:

$$N = VAR/MDE^2 \times 31.36$$

For different combinations of confidence level and power, the formula is

$$N = VAR/MDE^2 \times k$$

where *k* is given in Table 1.

Table 1 Values of the constant *k* for power calcuations

		confidence level		
		90%	*95%*	*99%*
power	*80%*	24.60	31.36	46.51
	85%	28.73	36.00	52.13
	90%	34.11	41.99	59.29

Surveys

Follow-up surveys, conducted after the intervention, are usually an integral part of the pilot study. While in some cases the outcome of interest might be easily measured (for example, a water provider will know its consumers' consumption), this is often not the case. For example, was an information campaign successful at giving people information? This question can be answered by surveying the treatment and control groups to find out what they know.

Even when the main outcome of interest is observed, the mechanism through which it was affected might be interesting to know. For example, although we observe changes in consumption, it may be interesting to know *how* people achieved it (e.g., did they wash their cars less, did they install water-efficient fixtures or irrigation systems, etc.). This information will not be available unless it is measured through a survey.

If feasible, it is often useful to conduct a *baseline* survey before the intervention takes place. Depending on the context, this may be conducted on the treatment and control groups, or it can be conducted on a random set of individuals before the control and treatment groups are selected. Such a survey can serve several purposes. First, if there is no information on the outcome of interest before the intervention, the procedure described above to determine the correct size of the treatment and control groups cannot be used. In this case, a baseline survey can be used to estimate the variance of the outcome for these calculations. Second, results from the baseline survey can be used to guide the intervention. For example, when designing an information campaign, a baseline survey might reveal where the biggest gaps in knowledge are in the population. Third, the baseline survey can provide a useful check on the success of the randomisation. When the treatment and control groups are small, it is possible that random selection still leads to differences between them. For example it is possible, even if highly unlikely, that the 500 consumers selected in the treatment group all happen to own a car, while none of the 500 consumers in the control group do. Whether the randomisation was successful in creating similar groups can be checked by comparing the two groups in the baseline survey. If the randomisation turns out to have produced dissimilar groups, this can be controlled for using regression analysis (similar to the DD method described above).

A final use of the baseline survey might be to allow one to use a refinement of random sampling, *stratification*.

Randomisation with stratification

As mentioned above a potential issue with small groups is that randomisation might, by pure luck, still produce very different treatment and control groups. By chance, the 500 households in the treatment group could happen to be exactly the 500 highest consumers. We can eliminate this possibility if we have relevant information about the population before the intervention. For example, if consumption is known, the following procedure is available. First, sort the entire population into consumption “bins,” such as very low, low, medium, moderate, high. These could correspond, for example, to the lowest 20% of the population, the next 20%, and so on. Next, take the random samples for the treatment and control groups separately from each bin, in proportion to their share in the population. For example, when choosing 500 consumers for the treatment group, one would choose 100 consumers from each of these five bins. This ensures that the resulting groups are *representative* of the overall population: although they are chosen randomly, there are exactly as many high consumers in the sample as there are in the population, relatively speaking. This is called *stratification*.

Stratification can be done based on several characteristics. For example, the area where consumers live might be relevant. This can be combined with information on consumption to create the bins from which random groups are chosen. Of course, these characteristics must be chosen carefully, or the number of bins will quickly become very large. For example, combining the 5 consumption categories above with 4 residential areas and 3 income categories will yield 4 × 5 × 3 = 60 bins.

In general one may use stratification based on any characteristic that can have an effect on the outcome of interest. In the case of consumption, this may be, for example, consumption in the past, area of supply, income, or whether a person receives restricted service due to non-payment.

Missing data

All empirical studies have to deal with missing data. In RCT's, missing data can arise in two places: at baseline or at follow-up. Missing baseline data means that pre-treatment information is missing for some participating subjects. This is usually not very problematic, since the inclusion of baseline variables typically only serves to increase the precision of estimated treatment effects. The estimated treatment effects obtained from comparing the control and treatment groups would be valid (statistically consistent) even in the absence of any baseline information. If one does want to include baseline information, however, missing cases should not simply be dropped, since this would make the analysed sample potentially biased. For example, if 200 out of 1000 observations are missing a relevant variable, there is nothing to guarantee that dropping the 200 observations would yield a sample with the same characteristics as the 1000 observations. Valid ways of dealing with missing baseline values are reviewed in White and Thompson (2005). One best-practice method is to impute missing data in the following way: for categorical variables, create an additional "missing" category, while for continuous variables, replace missing values with their sample means, and create an additional indicator that takes a value of 1 for these missing observations, and 0 otherwise. This will result in complete baseline data for all observations that can be used in the analysis. While this procedure might lead to biased coefficients for the control variables, it will not bias the estimated treatment effects.

Missing follow-up data, when some participating subjects have missing information on potential outcome variables, is a more serious problem. Outcome data should never be imputed, since it can lead to biased treatment effect estimates. For example, suppose that missing subjects happen to be those for whom the treatment was ineffective (did not change the outcome), while nonmissing values all correspond to effective treatments. Imputing mean sample values could then lead us to find statistically significant treatment effects because we would simply be replicating the nonmissing observations. The safest way to deal with missing follow-up data is usually dropping missing observations, combined with carefully controlling for any resulting imbalance in the sample, and recognising the fact that inference will only be valid for the population represented by the smaller sample.

3.5 Randomising Large-Scale Interventions

The above discussion has focused on the use of RCT in pilot studies conducted before a large-scale policy intervention is carried out. However, in some cases the actual large-scale intervention can also be randomised. For example, the intervention considered may be a large-scale information campaign for 10,000 households. After a pilot study is conducted with a randomly treatment and control group of 500 each, it is decided that it is worth implementing the large-scale intervention. Which 10,000 households should receive the information? If there is flexibility in the decision making environment, it may be possible to select these 10,000 households randomly. Even if randomly choosing individual households is not feasible, randomisation may be possible at the level of supply areas. For example, one could randomly choose among city blocks, and provide the information to every household living in the selected blocks.

Randomising the large-scale policy intervention makes it possible to evaluate the effect of the intervention after the fact, in much the same way as was described above for the pilot study. This may provide useful information for cost-benefit calculations after the fact, and perhaps lessons for future policy interventions.

4 DESCRIPTION OF THE INTERVENTION

In an attempt to improve households' information, we designed and implemented an in-depth water education programme. Education officers of Odi Water visited the 500 households in the treatment group in November and December 2012 (see Chapter 5 below on the selection of these households). Each household was visited once, and visits 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 meter, understanding the bill, detecting and fixing leaks, conserving water indoors, conserving water outdoors. They explained the contents of the each brochure to the household, highlighting specific points in each brochure.

To make the learning experience more effective and to involve all members of the household, the officers also gave each visited household an “Activity Book.” This contained a set of simple exercises for the household to fill out. The main exercise involved completing a water diary over a period of 3 days, recording how much water was being used in and around the house. Upon completion of the activity books, these could be returned at any Odi Water office for a cash reward.

All education materials (brochures and activity books) were designed by us specifically for this project using previous WRC materials.[6] They were designed and printed to look professional and be informative to a wide range of consumers. The information was relayed in a reader-friendly manner using pictures and examples. All materials were translated into Setswana by IRI, and the translations double-checked for accuracy by an Odi Water employee.

We provide details on these materials as well as the various aspects of implementing this intervention below.

4.1 Information brochures

The brochures were designed to be two-sided A4 pages folded in 3 (see the Appendix for an example).

1. **Meter**. This brochure described the operation of the water meter, explaining what the numbers on the meter dial mean. Education officers were also instructed to show each

[6] See Appendix I for a list of materials used.

household where their meter was located on the property, and explain to households how much water a “kilolitre” represented.

2. **Bill**. This brochure presented a picture of a water bill, and explained what the various numbers meant (for example, previous and current reading, total consumption, outstanding balance and interest, total amount due). It also reminded households of the various ways they can pay their bill. Finally, it contained a table listing the current tariff, with an example of how the price of consumption could be computed.

3. **Leaks**. This brochure explained the functioning of the stop-cork (with which households can turn off all water on their property). It also described how to check for and fix simple leaks (taps, toilets).

4. **Water saving indoors**. This brochure contained tips on how households could save water indoors. For example, they could turn off the tap while washing teeth, or they could put a bottle filled with sand in the toilet tank. We focused on giving households information and advice that they could follow if they wished, rather than instructions that they must obey.

5. **Water saving outdoors**. This brochure contained advice and information on saving water outside the house. For example, it emphasised the large amount of water that a garden hose would discharge, and explained the amount of water that could be saved by washing a car with a bucket.

4.2 The activity book

The activity book contained a set of simple exercises for the households. The goal was to involve all members of the household, and make them active participants in the learning process. To give households an incentive to complete the books, they could return the filled-out activity books to Odi Water in exchange for a cash reward of R100.[7]

The first set of exercises asked participants to measure (with the help of, e.g., a large soda bottle cut in half) how much water various everyday activities used. These included activities such as washing hands or filling up a kettle. We also listed the typical amount of water used by larger-scale activities that would be hard to measure, such as taking a shower or a bath.

Next, we invited households to fill out a Water Diary over a 3-day period, recording each activity that uses water inside or outside the house. We asked them to involve all members

[7] To preserve the anonymity of the households with regard to the activity books (as well as to prevent duplication), these were identified by a serial number. Households were instructed not to write their name on the books.

of the household, including children. In a separate column, households could then use the measurements they had taken and the information we had given them to record how much water each of these activities used.

Finally, we showed them how to calculate their (approximate) monthly water bill, by multiplying their 3-day consumption by 10 and using the current tariff schedule to compute the price. We also invited them to record their meter reading before and after the 3-day period, and compare their metered consumption to what they had recorded.

Out of the 500 treated households, only 51 returned the activity books before the follow-up survey, and a further 19 returned them after the follow-up. This low take-up indicates that the activity book component of our treatment was not important for the effects we find below. Whether a similar treatment that achieved a higher take-up of this component could have different results (notably a larger impact on households' knowledge) is an open question.[8]

4.3 Details of the implementation

The intervention was implemented by 9 education officers of Odi Water trained by us specifically for this project. After a day-long training session on October 15, 2012 the officers did a series of pilot household visits on October 16 and 17. We then held another training session on October 19 to collect feedback and answer questions before rolling out the full-scale treatment.

Odi's education officers are employed by its Marketing Department and regularly visit the community to educate households on water policy and give them advice on water usage. The officers all know the area well, are used to making household visits, and they all speak English, Setswana, as well as various other local languages.

Officers were each given a list of specific consumers to visit. They were instructed to look for the person whose name is on the water bill or an adult member of his household. If the right person to talk to was not available, they would visit the household at a later time.

During our training sessions, we made sure that the education visits would be as similar as possible to each other. We gave each officer a script containing a list of guidelines, including the order in which the brochures should be presented, and the specific points that should be highlighted from each brochure. We also emphasised the important differences of this project relative to the routine visits that the officers are used to making. Most importantly,

[8] Since the households who returned the activity books may be different from the ones that did not (i.e., this is not random), our research design is not suited to measure the causal effect of the activity books alone.

we made sure that visited households would not feel pressured. All our education materials deliberately 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. The officers' task was to provide information that households might view as helpful, and let each consumer decide if and how he wanted to use it.

Based on the officers' feedback, the households appreciated the visits, and were especially interested in information pertaining to the bill as well as some specific water saving techniques. They were also surprised and delighted by the activity book and the prospect of the reward.

Compared to other interventions analysed in the literature, our treatment had three distinguishing features. First, it was an in-depth education campaign. Most programmes in the literature consist of simple interventions like a letter sent out to households or a flyer included with their monthly bill. Our officers personally visited each household and provided them with extensive information on various features of the consumption process. Second, our treatment was explicitly focused on information provision, and we deliberately tried to minimize the social pressure component as much as possible. Third, we conducted our campaign in a developing country setting where the lack of information is known to be a serious issue.

5 SAMPLING OF PARTICIPATING HOUSEHOLDS AND IMPLEMENTATION OF THE PROJECT

5.1 Sampling

The objective of the sampling design was to yield a representative sample of the residential consumers of Odi Water, based on information that was available prior to the beginning of the project (as of January 2012). We used the entire population of residential water consumers of Odi Water, excluding commercial users and consumers using more than 300 kl (0.3% of the population).[9] We used stratified random sampling (see Part I) to select 500 treatment and 500 control households. Stratification was based on monthly water

[9] Accounts with over 300 kl consumption (or 25 times the average) are likely associated with unreported commercial activities or major leaks.

consumption,[10] indigent status, whether the consumer was restricted, and whether the consumer had made a payment on his water bill during the previous year.[11] These variables defined 32 groups in the population.

To obtain a sample of 1000 households taking into account the possibility of attrition (e.g., refusal to participate), we proceeded as follows. First, we took 3 samples of 1000 households using stratified random sampling. The survey team was instructed to interview each of the 1000 households on the first list. If a household could not be reached after two attempts or if the respondent declined to answer, it was replaced by a random household from the second list *belonging to the same group as the original household.* If that household could not be reached or declined to answer, it would be replaced by a third household from the third list. This procedure ensured that the resulting sample was representative of the surveyed population. In practice, 31 households were replaced once, and the remaining 969 household were reached from the first sample list during the first two attempts.

5.2 Implementation

The baseline and follow-up survey was carried out by surveyors employed by Impact Research International, Inc. (IRI), specifically trained for this project by the authors of this report. These surveyors visited each household at their home address and asked and recorded the questions to the questionnaire. The questionnaires were bilingual (English and Setswana), and interviews could be performed in either of these languages (in practice all but 12 interviews were conducted in Setswana). The completed questionnaires went through a quality check by the staff at IRI and by the authors, and were entered into a computer database. Baseline data collection took place in March - April 2012. The education programme took place in November - December 2012. Finally, a follow-up survey was administered to all participating households in February 2013.

The surveys were conducted in accordance with the standards of the Institutional Review Board of the University of Houston regarding the ethical treatment of human subjects. Participation in the survey was voluntary and respondents could stop participating at any time. Each questionnaire took about 30 minutes to complete. Only adults between the ages 18-65 were asked to participate. The separation of individual consumption information (including account numbers) and survey responses was maintained throughout the survey.

[10] We formed 4 groups corresponding to quartiles in the population (1-6 kl, 7-10 kl, 11-16 kl, and above 16 kl).

[11] “Restricted” status refers to households whose water supply is restricted by the water provider using a flow limiter or other device, due to unpaid water bills. “Indigent” status refers to households registered with the municipality as indigent to receive various low income subsidies.

The final database identifies respondents by a code generated by the authors, and does not contain the billing account number, service address or name of the customers.

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.[12]

5.3 Missing data

Administrative data, including data on consumption, payment, restriction and indigent status is available for the entire sample. Due to logistical difficulties, we only managed to gather baseline survey data for 803 households. As described in Chapter 2, in regressions where we control for baseline survey characteristics, simply dropping observations with missing baseline data would result in potentially biased estimates as we would be analysing a potentially imbalanced sample. Instead, we deal with missing baseline information by imputation 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. 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 relatively low attrition rate of 3.4%. Out of these 966, we have baseline data for 776 households (80%) and impute missing baseline data for the rest 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.

[12] 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.

6 HYPOTHESES

If our treatment is effective at increasing information and improving efficiency, we expect to see an improvement in household's ability to manage their consumption. We should see an increase in households' knowledge (as measured in our follow-up survey), an increase in conservation practices, and a higher propensity to make payments. With regards to quantity of water consumed, the prediction is ambiguous. Increased information might lead to less waste, which will tend to reduce consumption, but this in turn could lead to increased consumption in other activities. For example, upon learning how much water baths use compared to showers, a household could substitute taking baths with taking showers. If the substitution effect is large enough, the amount of water consumed could increase in response to our treatment.

H1 (Information effect). *The information campaign should result in increased consumer knowledge, increased use of conservation practices, and more payments. The effect on consumption is ambiguous.*

Given the emphasis on psychological motives behind conservation in the existing (mostly US) literature, another possibility presents itself. Our education programme could exert an influence on payment and consumption through a "nudging" channel. For example, the education visit may remind a consumer of his outstanding bill, or make his water choices more salient.[13] He could also feel compelled to pay his bill or reduce his water usage because the visit might suggest to him that this is what he "should" do.[14] This may be so even though, 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. Yet another possibility is that the consumer might appreciate the provider's efforts in reaching out to the households, and might make more payments to reciprocate this.[15] All these possibilities, which we will refer to as "psychological effects," can lead to more conservation and payments without changing households' information.

---

[13] For example, Jessoe and Rapson (2013) find that increasing the salience of electricity usage induces US households to conserve more energy.
[14] Alcott (2011) and Ferraro and Price (2013) show that social pressure is an effective tool to induce conservation in the US.
[15] In the tax avoidance literature, emotions like reciprocity and trust towards the government are thought to be relevant determinants of payments (see Selmrod (2007) for a survey).

H2 (Psychological effects). *The information campaign should not change consumer knowledge, but should increase the use of conservation practices and lead to more payments. The effect on consumption is ambiguous.*

7 DATA

The data analysed in this report comes from two sources. Data on household demographics as well as measures of households' information regarding the consumption process come from our two surveys. Data on consumption, payment, and households' indigent or restricted status comes from the administrative records of Odi Water.

7.1 General sample characteristics and balance across treatment and control groups

Table 2 presents various characteristics of our sample, separately for the treatment and control groups. In the 3 months prior to our intervention, average water consumption in the sample was 16 kl, and average monthly payments made by households R260. Little more than half of the households made at least one payment on their water bill in the 3 months preceding the treatment. Turning to general household characteristics, the average household has 4 members, one of whom is employed. Approximately 70% of the respondents completed high school, 66% have hot running water, and 36% own a car. Average reported household income is R6895 per month.[16]

Importantly, Table 2 shows that there are no significant differences between our treatment and control groups on any observable characteristic. In other words, the randomisation was successful.

[16] As a comparison, Statistics South Africa estimated an average income of R5803 for Black households across South Africa for 2010/11 (Statistics South Africa, 2012).

Table 2 Means of various observables in the control and treatment groups

	Control	Treatment	Difference
Consumption (kl)	15.001	16.969	1.967
	(0.628)	(1.333)	(1.473)
Payment (Rand)	278.450	242.509	-35.941
	(18.337)	(15.941)	(24.297)
Payment (yes/no)	0.566	0.515	-0.050
	(0.023)	(0.023)	(0.032)
Baseline survey	0.812	0.795	-0.017
	(0.018)	(0.018)	(0.026)
Informal shacks	0.123	0.129	0.006
	(0.017)	(0.017)	(0.024)
Employed HH members	1.048	0.996	-0.052
	(0.032)	(0.030)	(0.044)
HH size	4.338	4.481	0.143
	(0.078)	(0.094)	(0.122)
No formal schooling	0.010	0.005	-0.005
	(0.005)	(0.004)	(0.006)
Some primary school	0.010	0.010	-0.000
	(0.005)	(0.005)	(0.007)
Primary school	0.065	0.088	0.023
	(0.013)	(0.014)	(0.019)
Some high school	0.217	0.202	-0.016
	(0.021)	(0.020)	(0.029)
High school	0.434	0.432	-0.003
	(0.025)	(0.025)	(0.036)
Some higher educ.	0.165	0.152	-0.013
	(0.019)	(0.018)	(0.026)
Higher education	0.098	0.111	0.013
	(0.015)	(0.016)	(0.022)
Hot water	0.691	0.641	-0.050
	(0.023)	(0.024)	(0.034)
Owns car	0.369	0.364	-0.005
	(0.025)	(0.025)	(0.035)
Owns fridge	0.977	0.982	0.005
	(0.008)	(0.007)	(0.010)
Income (Rand)	7,056.548	6,736.557	-319.990
	(236.554)	(226.010)	(327.167)
N. sampled neighbours	1.134	1.251	0.117
	(0.050)	(0.054)	(0.074)
Has treated neighbour	0.466	0.444	-0.022
	(0.023)	(0.023)	(0.032)

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 neighbours' is the number of households included in the sample in a 100 meter radius, and 'Has treated neighbour' is 1 if one of these households is in the treatment group. In the third column, ***, **, * denote statistically significant differences at 1, 5, and 10%, respectively.

7.2 Measuring households' knowledge on water consumption

Our information campaign focused on three key areas of the water consumption process: (1) Understanding the meter; (2) Understanding the bill; (3) Understanding water quantities used in everyday activities. Separate sections in our surveys were designed to measure each of these areas.

Understanding the water meter

Each property is equipped with a water meter that is easily accessible to the households and shows the water consumption in litres and kilolitres. Monitoring one's water use is the first step towards planned consumption and conservation. We asked several questions to measure whether households know where to find and how to read the water meter.

In the surveyed population, almost everybody (96.1%) could show to our surveyors where the water meter was located on their property. An additional 3.4% asked someone from the family and they were able to show the water meter. We encountered only 4 households who did not know where their water meter was located. Households are also familiar with the basic functioning of the meter. To measure this, we asked *"Suppose you have a leaking tap in your house. What will happen to the numbers on the meter?"* Only 4.5% of respondents were confused, stating either that they did not know or that the numbers on the meter would decrease. 93.2% stated that the numbers would increase, and 2.3% that they would stay the same.

Although almost everybody knows where their meter is and they are familiar with the basic functioning of the meter, 67.5% of the respondents never check it. Only 18% of all households monitor the water meter at least once a month. One reason could be that the numbers are not meaningful to them. We asked respondents to guess how much water their households used in a month. Households were free to answer this question in any measurement unit (the bill that households receive lists consumption in kilolitres, 1 kl = 1000 l). The first thing to note is that households use litres, rather than kilolitres when thinking about water quantities. Only 8 households (1%) responded to this question in kilolitres. The second thing to note is that households do not know how much water they use. Households gave a wide range of quantities, from 6 litres to 140,000 litres. The average response among these households is 2277 litres with half the sample giving answers of 800 litres or below, and 98% giving answers below the true median of 12,000 litres. Overall, households significantly underestimate their consumption.

These findings from the baseline survey prompted us to include in the education programme an emphasis on explaining the concept of litres and kilolitres, as well as how much water specific activities use.

Understanding the water bill

Households receive regular monthly bills from the provider. We asked households several questions about their water bill to measure people's understanding of the billing process. We also suspected that one reason some households might run up large unpaid balances might be that they don't know how much they need to pay, or they do not understand how the utility computed the final charge and therefore are reluctant to pay it.

A total of 95.7% of households indicated that they received regular water bills. Considering the difficulties with mail delivery in some of these areas (e.g., lack of mailboxes), the latter is remarkably high. Again, there does not seem to be any issue with the availability of information.

We asked if households felt that the water bill was "very easy to understand / not so easy to understand / almost impossible to understand." Approximately 59% of households stated that the water bill was easily understandable, and only 7.6% viewed it as almost impossible to understand. Note that this is despite households not having a clear understanding of the concept of kilolitre, as discussed above. We also asked if households thought the bill was accurate. More than half of all households (54.8%) believe that the water bill is accurate.

Finally, we asked households to tell us from their water bill what their water consumption was that month. With the bill in their hand, 43.7% of respondents said they could not tell their consumption from the bill. However, a further 38.5% stated an incorrect number (a number that, when compared to the billing data, did not correspond to any bill the household received in the 6 months prior to the survey). Overall, less than 18% of respondents were able to tell their consumption from their bill.

Understanding the tariff schedule

As is typically the case in the water sector, the consumers of Odi Water face a block tariff, which means that they need to pay a different kilolitre price depending on their total consumption. Higher consumption typically entails a higher unit price, which tends to discourage high water consumption. We wanted to know whether the households had an understanding of this basic feature of the tariff schedule.

We asked the following question designed to ascertain whether households understood that consuming more water would yield a higher per-kilolitre price. "Imagine you flush the toilet 1000 times. Do you think the first flushing will cost more, less, or the same as the 1000th?" A single flush uses 8 litres of water, which yields 8 kl of water used during 1000 flushes. Since the first tariff block applies between 0-6 kl, flushing the toilet 1000 times would always push consumption into a higher block with a higher associated unit price. Thus, the correct answer to the question is that the 1000th flush would cost more than the first flush.

We observe large variation in households' answers. 42.9% of the households think that there is no price change. A large share of the households (23.3%) believes that the unit price will be lower if they use more water, and only 33.8% correctly guessed that the price would increase.

We also asked households how much they thought they were paying for a kilolitre of water. We indicated to them that a kilolitre was equal to 500 two-litre bottles of soda filled with water. At the time of the survey, one kilolitre of water cost between 10 and 21 Rand depending on the households' total consumption. Households' answers ranged from R0.5 to R1000, with a mean of R119. Half of the households think that they pay more than 100 Rand per kilolitre which is about five times higher than the actual cost. Only 13% of the households gave a unit cost below the largest actual value of 21 Rand. Thus, households overestimate the unit cost of water by a factor of 5-10. At the same time, recall that on average they *underestimate* their water consumption by about the same factor.

We asked households if they currently received any free water from the utility. At the time of the survey, 30% of the households were registered as indigent and received 6 kilolitres of water for free from Odi Water.[17] Our survey shows that households are not aware of the current free water policy. Over 85% of the respondents think that they do not receive any free water, and this percentage is about the same among indigent households (who actually do) as among non-indigents (who actually do not). Over 85% of the households who receive free are not aware of this.

Understanding the quantities of water used

We asked four questions comparing quantities of water used for different activities. To manage water usage and conserve water it is important to understand which activities use

[17] In practice, water and sanitation charges are set separately for the same quantity of water consumed. While indigent households receive 12 kl of water for free, sanitation is only free until 6 kl. Thus, indigent households pay 0 in total for the first 6 kl. The City of Tshwane introduced this policy in July 2007 replacing the earlier system under which every household, not just those registered as indigent, received 6 kl for free.

the most water around the house. To increase the response rate, the questions were worded to resemble a guessing game rather than a test, and surveyors were instructed to ask the questions accordingly.[18]

"*Please take a guess: Do you think more water is used by the baths/showers your household takes during the month, or by washing your clothes during the month?*" An average bath or shower uses about 30 litres or more and an average washing machine uses about 80 litres. A front-loading machine uses much less, about 40 litres, and washing clothes by hand uses even less. Over a one month period, the water used for bathing / showering is considerably more than that used for washing clothes. Approximately half of the households (53.7%) answered this question correctly.

"Please take a guess: Do you think more water is used if you fill 2 two-litre bottles of soda with water, or if you flush the toilet once?" A typical toilet tank used in the Mabopane area uses 8 litres of water, consequently flushing the toilet uses more water. About two-thirds of the households (64.4%) gave a correct answer.

"Please take a guess: Do you think more water is used if you use the outside hose for 10 minutes, or if you do one load of laundry?" A typical hose discharges around 10 litres per minute, so this uses more water. 58.1% of the respondents knew the correct answer.

"Please take a guess: Do you think more water is used if you open the tap for 1 minute, or with the water a person drinks in a day?" A typical tap discharges 6 litres of water per minute, while most people drink at most 4 litres of water a day. Over three quarter of the respondents (77.1%) gave correct answers.

Overall, 17% of consumers answered all four questions correctly, and 51% gave at least 3 correct answers. When interpreting these numbers, it is important to note that if respondents were to answer by flipping a coin and answering randomly, 31.25% would get at least 3 correct answers.

[18] Note that if respondents answered randomly, in a large sample we would expect half of the answers to be correct for each question.

8 SPECIFICATION AND RESULTS

8.1 Specification

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

$$y_i = \beta_0 + \beta_1 \times Treat_i + \varepsilon_i$$

where y_i is the outcome of interest for household *i*, and $Treat_i$ is an indicator equal to 1 for treated and 0 for control households. Without additional controls, the estimate of the treatment effect β_1 is simply the difference of the mean of y_i between the treatment and control groups. To increase the precision of the estimates, we sometimes include in the above regression indicators for the 32 strata used in sampling, the baseline value of the outcome *y*, and various demographic controls. Throughout, we measure consumption and payment amounts in logs. In these regressions, coefficients can be interpreted as percentage changes in the outcome in response to the treatment.

Table 3 Treatment effect on consumption

Dependent variable	(1)	(2)	(3)
January - March	0.010	0.010	0.011
average consumption	(0.031)	(0.030)	(0.030)
January consumption	-0.049	-0.068	-0.065
	(0.058)	(0.060)	(0.060)
February consumption	0.021	-0.016	-0.014
	(0.018)	(0.024)	(0.024)
March consumption	-0.059	-0.052	-0.051
	(0.096)	(0.095)	(0.096)
Number of observations	947	947	947
Strata indicators	No	Yes	Yes
Demographic controls	No	No	Yes

Notes; Each cell presents the estimated treatment effect from a different regression. The first column gives the dependent variable: log consumption over the given period. Columns (1)-(4) correspond to different specifications. Each specification controls for average consumption during the 3 months prior to the treatment (in logs). `Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

8.2 Treatment effects on consumption

We begin in Table 3 by studying the effects of our treatment on consumption. In the first row, the dependent variable is the log of average household consumption in the three months following the treatment, and the following rows look at each of these months separately. In column (1) we only control for average consumption in the three months prior to the treatment, column (2) adds sampling strata indicators, and column (3) adds a variety of socio-economic characteristics from the baseline survey. The absence of a treatment effect on average consumption can never be rejected.

The absence of an average treatment effect on consumption is not very surprising, since households consuming different quantities of water have very different possibilities to respond. While a large consumer who is wasting water may reduce his consumption, a small consumer who is consuming too little water in an effort to save money might increase his consumption in response to an effective treatment. In Table 4, we break up the sample into consumption quartiles, and find a significant treatment effect for the largest 25% of consumers. Households consuming more than 19 kl reduced their consumption by an average of 9.5% in response to the treatment. We investigate such heterogenous treatment effects more fully in Chapter 10.

Table 4 Treatment effects on consumption by consumption quartile

	First quantile	Second quantile	Third quantile	Fourth quantile
	Less than 7 kl	7 - 11 kl	12 - 19 kl	More than 19 kl
Treatment	0.083	0.043	-0.023	-0.095*
	(0.069)	(0.063)	(0.056)	(0.052)
N	227	246	245	229

Notes: Each column regresses average consumption in the 3 months following the treatment on a treatment indicator on a different consumption quartile of the sample. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

We saw that, on average, individuals did not change their consumption significantly in response to the treatment. Did individuals change their behavior (perhaps reducing consumption in some activities increasing it in others)? To analyse this question fully would require detailed data on household water use in different activities. In the absence of this, we look at households' self-reported conservation behavior. Our surveys asked whether the household recently took actions to conserve water, listing several possibilities. Table 5 looks for treatment effects in these answers. Treated households are 10 percentage points more likely to report fixing leaks around the house and 8.5 percentage points more likely to report conserving water during laundry.[19] Treated households also report taking more actions than control households, although the fraction of households taking no action does not differ significantly between the two groups. This suggests that the treatment primarily increased conservation on the intensive margin, among households already taking steps to conserve water.[20]

Table 5 Effect of treatment on conservation

Dep. var	Mean dep. var	Treatment effects			
		(1)	(2)	(3)	(4)
Use rainwater	0.013	-0.006	-0.007	-0.005	-0.007
		(0.007)	(0.008)	(0.007)	(0.007)
Reuse water	0.291	0.013	0.014	0.015	0.012
		(0.029)	(0.029)	(0.029)	(0.029)
Repair leaks	0.415	0.104***	0.103***	0.104***	0.103***
		(0.032)	(0.031)	(0.031)	(0.032)
Conserve with laundry	0.291	0.084***	0.083***	0.083***	0.082***
		(0.029)	(0.029)	(0.029)	(0.029)
Conserve with irrigation	0.252	-0.040	-0.040	-0.040	-0.041
		(0.028)	(0.028)	(0.028)	(0.028)
Number of actions	1.361	0.167**	0.166**	0.169**	0.158**
		(0.071)	(0.070)	(0.070)	(0.071)
No action	0.224	-0.033	-0.033	-0.034	-0.030
		(0.027)	(0.026)	(0.026)	(0.027)
Strata indicators		No	Yes	Yes	Yes
Baseline dep. var.		No	No	Yes	Yes
Demographic controls		No	No	No	Yes

Notes: Each cell presents the estimated treatment effect from a different regression. The first column gives the dependent variable. Except for the last two, these are dummies for whether the respondent reported having taken the action to conserve water. 'Number of actions' is the number of actions the household reported. 'No action' is a dummy equal to 1 if the household did not report taking any action. Columns (1-4) correspond to different specifications. 'Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. Robust standard errors in parentheses. N = 965. ***, **, * denote significance at 1, 5, and 10 %, respectively.

8.3 Treatment effects on payment

We next study any changes in consumers' payment behaviour caused by the treatment. The first row of Table 6 compares households' total payment (in logs) in the three months following the treatment between the treatment and control groups. In column (1) we only control for total payment in the three months prior to the treatment and find that, following the

[19] "Use washing machine less / use fuller loads."

[20] It is conceivable that households could lie about taking conservation actions to satisfy perceived social expectations. While we cannot rule this out definitively, we find it reassuring that there are no significant differences in reports of using rainwater or reusing household water in Table 5. Neither of these practices was mentioned in our education campaign. Households report more conservation activities in those areas that were explicitly covered in the campaign.

treatment, the average treated household paid 32% more than the average control household. The magnitude of this estimate drops to 25-26% but remains robustly significant when including indicators for the sampling strata (column 2), socio-economic characteristics (column 3), and average monthly consumption in the three months before the treatment (column 4). As shown above, we do not see an increase in consumption in response to the treatment. Thus, the increased payment we find is not explained by households simply using more water.

Table 6 Treatment effect on payment amount

Dependent variable	(1)	(2)	(3)	(4)
January – March total payment	0.320**	0.251*	0.266**	0.262**
	(0.136)	(0.129)	(0.130)	(0.129)
January payment	0.247*	0.214	0.229*	0.241*
	(0.136)	(0.135)	(0.136)	(0.138)
February payment	0.282**	0.246*	0.221	0.222
	(0.137)	(0.137)	(0.138)	(0.139)
March payment	0.180	0.128	0.126	0.115
	(0.141)	(0.137)	(0.138)	(0.138)
Number of observations	966	966	966	947
Strata indicators	No	Yes	Yes	Yes
Demographic controls	No	No	Yes	Yes
Pre-treatment consumption	No	No	No	Yes

Notes: Each cell presents the estimated treatment effect from a different regression. The first column gives the dependent variable, and columns (1)-(4) correspond to different specifications. All payment and consumption measures are in logs. Each specification includes average monthly payment during the 3 months prior to the treatment. `Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. `Pre-treatment consumption' is average consumption during the 3 months prior to the treatment. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

The remaining rows of the table suggest that the increase in monthly payments is short lived: the effect remains similar in the first two months but falls by almost a half and becomes insignificant by the third month.[21] Still, the 25-30% increase in payment over a 3-month period indicates a substantial revenue increase for the provider.

[21] On average, households made 1 payment over this three month period, which explains why the estimated effect on total payments is similar to the effect on monthly payments (for the first two months).

Table 7 Payment propensity

Dependent variable	(1)	(2)	(3)	(4)
January – March payment (0/1)	0.050**	0.038*	0.041*	0.040*
	(0.024)	(0.022)	(0.022)	(0.022)
January payment (0/1)	0.045*	0.038	0.040	0.043*
	(0.026)	(0.026)	(0.026)	(0.026)
February payment (0/1)	0.055**	0.048*	0.043*	0.043
	(0.026)	(0.026)	(0.026)	(0.026)
March payment (0/1)	0.034	0.024	0.023	0.021
	(0.027)	(0.026)	(0.026)	(0.026)
Number of observations	966	966	966	947
Strata indicators	No	Yes	Yes	Yes
Demographic controls	No	No	Yes	Yes
Pre-treatment consumption	No	No	No	Yes

Notes: Each cell presents the estimated treatment effect from a different regression. The first column gives the dependent variable: 1 if the household made a payment over the given period, 0 otherwise. Columns (1)-(4) correspond to different specifications. Each specification controls for whether the household made a payment during the 3 months prior to the treatment. \`Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. \`Pre-treatment consumption' is average consumption during the 3 months prior to the treatment (in logs). Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

Did the extra payment come only from households paying more, or did households' propensity to pay increase as well? Table 7 shows that the treatment increased the fraction of households making at least one payment in the three months following the treatment by 4-5 percentage points (relative to a mean of 54.5%). This effect is again driven by the two months immediately following the treatment. Table 8 shows that treated households also made significantly more payments, with a small increase of around 0.1 extra payment relative to a mean of 1.1 over the three-month period.

Since payments are bounded below by 0, estimation methods that take into account such censoring may provide more precise results. In Table 9 we estimate treatment effects on payment amounts using Tobit regressions. These give somewhat larger marginal effects than those presented above. For example, we estimate a 26% increase on 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). Estimating the effects on the propensity to make payments with Probit instead of OLS also yields larger point estimates.[22]

[22] These treatment effects could be underestimated if the treatment induced some households with low consumption to register as indigent and receive a free water allowance. However, we find no evidence that our treatment had an effect on households' registered indigent status (regressions not shown).

Table 8 Payment frequency (3 months)

	(1)	(2)	(3)	(4)
Treatment	0.097**	0.088*	0.092**	0.092**
	(0.048)	(0.046)	(0.046)	(0.047)
Number of observations	966	966	966	947
Strata indicators	No	Yes	Yes	Yes
Demographic controls	No	No	Yes	Yes
Pre-treatment cosumption	No	No	No	Yes

Notes: Each column corresponds to a different regression. The dependent variable is the number of payments made by the household in the 3 months following the treatment. Each specification controls for the number of payments the household made during the 3 months prior to the treatment. `Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. `Pre-treatment consumption' is average consumption during the 3 months prior to the treatment (in logs). Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

Table 9 Treatment effects on payment amount and propensity: Tobit and Probit estimates

Period	Treatment effect on payment (1)	Treatment effect on payment | payment > 0 (2)	Treatment effect on payment propensity (3)
Jan-March	0.367**	0.258**	0.082**
	(0.171)	(0.120)	(0.038)
January	0.228*	0.192*	0.055*
	(0.128)	(0.108)	(0.032)
February	0.247*	0.202*	0.064**
	(0.133)	(0.109)	(0.033)
March	0.200	0.159	0.045
	(0.142)	(0.113)	(0.033)

Notes: Marginal effects of the treatment indicator from Tobit (columns 1 and 2) and Probit (column 3) regressions. The first column gives the period of the dependent variable. In columns (1) and (2), the dependent variable is log payment over the given period. Column (1) presents unconditional marginal effects and column (2) marginal effects conditional on positive payments (from the same regression). In column (3) the dependent variable is an indicator equal to 1 if the household made a payment over the given period and 0 otherwise. Columns (1) and (2) control for total payment in the 3 months before the treatment, and column (3) controls for whether a payment was made during this period (marginal effects are evaluated at the means of the controls). Robust standard errors in parentheses. N = 966. ***, **, * denote significance at 1, 5, and 10 %, respectively.

8.4 Treatment effects on information

We saw that our treatment raised households' propensity to pay their water bill and increased their self-reported conservation activities. What is the mechanism behind these effects? An important element of our research design is that we are able to look households' information directly, and evaluate the extent to which change in knowledge is responsible for the treatment effects we found above.

Our information campaign focused on three key areas of the water consumption process: (1) Understanding the meter; (2) Understanding the bill; (3) Understanding water quantities used in everyday activities. Separate sections in our surveys were designed to measure each of these areas.

The meter. As described in Chapter 4, although households know where their meter is located and understand what it is for, there is a lack of understanding about what the numbers mean and the units in which water is measured. Our information campaign significantly raised households' familiarity with the concept of a kilolitre, increasing the number of households responding in kilolitres by 4 percentage points (first row of Table 10). However, many households seem to have become familiar with the word "kilolitre" without learning what it means. Over 60% of those who answered in kilolitre after the treatment gave unrealistic numbers of several hundred or even thousands of kilolitres.[23]

When asked how many litres a kilolitre represented, only 3 households gave the correct answer. Others didn't know or were off by a factor of 10, with no significant difference between treatment and control. Households' learning about how water consumption is measured was superficial.

The bill. As described in Chapter 4, a large majority of households state that they understand their water bill. Our information campaign may have increased this even further, although the effect is not significant (Table 10). However, stating that the bill is understood does not mean that the respondent actually understands it. 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. There was no significant difference between treatment and control (Table 10).

Regarding the price of water, in the follow-up survey less than 5% of households gave numbers in the ballpark of the true kilolitre price. These are the households who state prices between 5 and 25 Rand (the true kilolitre price is between 10 and 21, depending on consumption).[24] There was no difference between treatment and control either in the fraction of these households, or in how far off reported prices were from realistic values (Table 10). There was no difference in knowing the fact that the price schedule is increasing, i.e., that an additional kilolitre costs more when consumption is high than when it is low (Table 10). Households' understanding of their water bill did not increase in response to the treatment.

---

[23] As in the baseline, responses in litres were too low, with 90% giving numbers less than 1000 litres.
[24] 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.

Table 10 Effect of treatment on information

Dep. var	Mean dep. var	N	Treatment effects			
			(1)	(2)	(3)	(4)
Response in kl	0.104	953	0.038*	0.038*	0.037*	0.036*
			(0.020)	(0.020)	(0.020)	(0.020)
Bill hard to understand	0.062	952	-0.024	-0.024	-0.023	-0.026
			(0.016)	(0.016)	(0.016)	(0.016)
Reads consumption from bill	0.397	731	0.036	0.042	0.043	0.046
			(0.036)	(0.037)	(0.037)	(0.037)
Consumption accurate	0.114	731	0.037	0.037	0.038	0.034
			(0.023)	(0.023)	(0.023)	(0.023)
Tariff in ballpark	0.045	820	-0.017	-0.017	-0.016	-0.017
			(0.014)	(0.014)	(0.015)	(0.015)
Tariff error	70.591	396	-21.967	-5.935	-6.141	-3.922
			(17.540)	(9.637)	(9.256)	(9.217)
Increasing tariff	0.699	964	-0.023	-0.021	-0.023	-0.020
			(0.030)	(0.030)	(0.030)	(0.030)
N. correct answers	2.485	965	0.057	0.056	0.059	0.046
			(0.064)	(0.064)	(0.064)	(0.065)
Q1 correct	0.459	966	-0.006	-0.006	-0.005	-0.012
			(0.032)	(0.032)	(0.032)	(0.032)
Q2 correct	0.726	966	0.052*	0.053*	0.050*	0.051*
			(0.029)	(0.029)	(0.028)	(0.029)
Q3 correct	0.604	965	-0.005	-0.005	-0.005	-0.013
			(0.032)	(0.032)	(0.032)	(0.032)
Q4 correct	0.697	966	0.015	0.013	0.017	0.019
			(0.030)	(0.030)	(0.030)	(0.030)
Strata indicators			No	Yes	Yes	Yes
Baseline dep. var.			No	No	Yes	Yes
Demographic controls			No	No	No	Yes

Notes: Each cell presents the estimated treatment effect from a different regression. The first column gives the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. The last 5 rows are the number of correct answers to the quiz and indicators for whether individual questions were answered correctly. 'Please take a guess: Do you think more water is used... Q1. By the baths/showers your household takes during the month OR by washing your clothes during the month. Q2. If you fill 2 two-litter bottles of soda with water OR if you flush the toilet once. Q3. If you use the outside hose for 10 minutes OR if you do one load of laundry. Q4. If you open the tap for 1 minute OR with the water a person drinks in a day.' Columns (1-4) correspond to different specifications. 'Demographic controls' are the number of children, teenagers, adults in the household, number of employed members, education of respondent, household income, and whether the household has hot running water, owns a car, or owns a fridge. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 %, respectively.

Quantities of water used. In our quiz to measure households' understanding of quantities of water used, 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 was also very similar. There were no differences between treatment and control (Table 10).

Looking at the fraction of correct answers to individual questions, we only find a significant treatment effect for one of them, regarding the amount of water used when flushing the toilet.[25] This is a relevant dimension as the toilet is typically a major source of water consumption, responsible for a quarter or more of indoor water use (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% (Table 10). Overall, households' understanding of quantities of water used did not increase much in response to the information campaign.

Taken together, our education treatment had at best a modest effect on households' knowledge. Thus, surprisingly, an increase in information does not seem to explain the reduction in nonpayment achieved by our education campaign.

9 POSSIBLE EXPLANATIONS

9.1 Spillovers

A potential concern with any information treatment is the possibility of spillovers. Treated households could talk to their neighbours about what they have learned, or they could give them the information brochures. Even if the treatment was effective at increasing households' information, such spillovers could result in no difference in information between the treatment and control groups. Could this be responsible for the lack of information effects we found above?

Note first that in most cases, we did not simply find the information of treatment and control groups to be similar, but also that they were both similarly low both before and after the treatment. While spillover effects from our treatment could potentially explain the first of these patterns, they are unlikely to account for the second. If the treatment had increased information and there were spillovers, we would expect to find increased knowledge in both the treatment and control groups.

To formally test whether spillovers were present in our intervention, we collected data to identify individuals who would be most likely to be exposed to information spillovers. First, our survey collected information on whether the respondent had talked to his neighbours or friends about water in the previous 6 months. If there were information spillovers, these would likely be present among the 39% who reported talking about water with others. Second, we collected each household's GPS coordinates and thus know their location

[25] "Do you think more water is used (a) if you fill 2 two-liter bottles of soda with water, or (b) if you flush the toilet once?" The correct answer is (b) for all toilet tanks used in this area.

relative to other households.[26] Information spillovers could occur between neighbours, and we can capture this by creating an indicator for whether a household has other treated households nearby.

Let *Exposure* represent one of the above proxies for exposure to information spillovers. We estimate

$$Y_i = \beta_0 + \beta_1 \times Treat_i + \beta_2 \times Exposure_i + \beta_3 \times Exposure_i \times Treat_i + \varepsilon_i,$$

where, Y_i is one of our measures of respondent *i*'s knowledge. If the treatment did have an effect on Y_i, but large spillovers caused us to find no effect, then we expect to find $\beta_1 > 0$ and $\beta_2 > 0$. By contrast, if the treatment was indeed ineffective at increasing knowledge, we expect $\beta_1 = \beta_2 = 0$.

In Table 11, our measure of exposure is *Talks*, which takes a value of 1 if the respondent talked to neighbours about water in the previous 6 months. Our dependent variables are the main information measures that we found to be unaffected by our treatment. The table also presents an F-test and the corresponding p-value for the hypothesis that $\beta_1 = \beta_2 = 0$ (no spillovers). For 5 out of 8 variables, the hypothesis of no spillovers is not rejected. In the remaining 3 columns (1, 7, and 8), the coefficients on *Talks* is negative: if anything, individuals who talk to others have less information. In Column (6), we find a positive and significant β_1 but the point estimate on *Talks* remains insignificant and negative. This suggests that the treatment may have been relatively more effective in improving the respondents' quiz scores among individuals who do not talk about water.[27] In none of these regressions does the evidence support the idea that the treatment raised information but was accompanied by large information spillovers ($\beta_1 > 0$ and $\beta_2 > 0$).

Table 12 presents corresponding regressions using GPS coordinates to identify a household's neighbours, and using the treatment status of a household's neighbours to capture potential exposure to information spillovers. The variable *Treated neighbours* takes a value of 1 if there is one or more treated household in a 100 meter radius around the respondent. 45% of the households in our study have such a neighbour, and the number of treated neighbours ranges between 0 and 4. The results in Table 12 also reject the idea that spillover effects could explain the lack of information effects found above. We do not find any

[26] 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 are the only way to map these households.
[27] This does not appear to be because these individuals had a lower level of knowledge to start with. In fact, individuals with $Talks_i = 0$ had a slightly higher average score at baseline (2.55 vs. 2.52, the difference is not statistically significant). Instead, individuals who talk less about water with others may have been more attentive during the education visit.

support for the hypothesis that $\beta_1 > 0$ and $\beta_2 > 0$. In some cases, having neighbours in the treatment group is associated with significantly worse information.

Table 11 Checking for spillover effects I: Talking to neighbours

	(1) Response in kl	(2) Bill hard to understand	(3) Reads consumption from bill	(4) Consumption accurate	(5) Tariff in ballpark	(6) Tariff error	(7) Increasing tariff	(8) N. correct answers
Treatment	0.040	-0.022	0.041	0.015	-0.020	-30.092	0.006	0.138*
	(0.026)	(0.021)	(0.049)	(0.031)	(0.018)	(21.937)	(0.036)	(0.082)
Talks	-0.028	-0.015	-0.042	-0.009	0.001	-7.799	-0.085*	-0.101
	(0.025)	(0.025)	(0.052)	(0.031)	(0.023)	(31.669)	(0.044)	(0.092)
Talks x Treatment	0.002	-0.006	-0.008	0.043	0.008	27.315	-0.063	-0.179
	(0.039)	(0.031)	(0.073)	(0.048)	(0.030)	(34.665)	(0.062)	(0.133)
N	947	946	726	726	815	391	958	959
F test (Treatment, Talks)	2.999	0.577	1.207	0.278	0.812	1.317	2.371	3.315
p value	0.050	0.562	0.300	0.757	0.445	0.269	0.094	0.037
Mean dep. var.	0.104	0.062	0.397	0.114	0.045	70.591	0.699	2.485

Notes: Each column lists the coefficient estimates from a different regression (apart from the constant). The column headings give the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. 'N. correct answers' is the number of correct answers in our quiz. 'Talks' is 1 if the respondent answered Yes to 'In the previous 6 months have you talked to friends or neighbours about the way you use or save water?' Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 %, respectively.

Table 12 Checking for spillover effects II: Treated neighbours

	(1) Response in kl	(2) Bill hard to understand	(3) Reads consumption from bill	(4) Consumption accurate	(5) Tariff in ballpark	(6) Tariff error	(7) Increasing tariff	(8) N. correct answers
Treatment	0.021	-0.045**	0.008	0.028	0.020	-13.413	-0.079*	0.030
	(0.027)	(0.021)	(0.049)	(0.032)	(0.020)	(9.537)	(0.041)	(0.088)
Treated neighbours	-0.023	-0.026	-0.098*	-0.020	0.037	32.376	-0.015	0.070
	(0.025)	(0.024)	(0.051)	(0.031)	(0.023)	(39.460)	(0.042)	(0.088)
Treated neighbours x Treatment	0.037	0.046	0.058	0.018	-0.082***	-20.982	0.127**	0.064
	(0.040)	(0.031)	(0.073)	(0.047)	(0.028)	(40.964)	(0.059)	(0.130)
N	953	952	731	731	820	396	964	965
F test (Treatment, Treated neighbours)	1.416	2.190	2.708	1.154	1.418	1.487	2.100	0.319
p value	0.243	0.112	0.067	0.316	0.243	0.227	0.123	0.727

Notes: Each column lists the coefficient estimates from a different regression (apart from the constant). The column headings give the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. 'N. correct answers' is the number of correct answers in our quiz. 'Treated neighbours' is 1 if the respondent has a neighbour in the treatment group (in a 100 meter radius). Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 %, respectively.

9.2 Lack of information sharing within households

Another possible explanation for the lack of an information effect is that information may not be shared within the household. As described in Chapter 5, it makes sense to consider the household as the unit of analysis since consumption and payment are measured at the household level. However, this raises the possibility that surveyed individuals within the household are different from treated individuals. To fix ideas, suppose that the education officers met with the wife, who is responsible for paying the water bill, and the treatment successfully increased her knowledge. Suppose this information channel explains the findings above. We may still measure no treatment effect on information if our surveyors in the follow-up survey talked to the husband and the wife failed to share her information with him.[28]

We perform two further tests to assess the possibility that information sharing within the household might be an important factor in explaining the findings above. First, based on the respondent's age and gender, we identify households where the same respondent is likely to have answered the baseline and the follow-up survey. If the same person answered both surveys, it is more likely that (s)he was also home during the education visit. Under the information story, these households should show the biggest increase in knowledge relative to the control group. We have 28 such households in the control and 25 in the treatment group. Including this indicator and its interaction with treatment status yields a significant interaction in only one case, but with the wrong sign (Table 13). Relative to the control group, treated households where the same person was home during both surveys do not have significantly more information than others.

Our second test is based on the idea that if information sharing within the household is a major factor, we would expect treatment effects to diminish as households get larger. This is both because information sharing within the household becomes harder in a larger household, and because a larger household makes it more likely that the education officers and the surveyors met with different members of the household. As before, the interaction of household size with treatment status is only significant in one regression, but with the wrong sign (Table 14). Relative to the control group, smaller treated households do not have more information than larger households.

[28] In some sense, this explanation also falls under psychological effects (H2 in Chapter 6). If the treated individual does not share her information within the household but, e.g., simply tells her husband and children to change their water consumption habits, then the households' behavior as a whole changed due to a "social pressure" exerted by the wife. If by contrast she had explained to her family what she had learned, then this would have been reflected in the follow-up survey.

Table 13 Information sharing within the household I: Same respondent

	(1) Response in kl	(2) Bill hard to understand	(3) Reads consumption from bill	(4) Consumption accurate	(5) Tariff in ballpark	(6) Tariff error	(7) Increasing tariff	(8) N. correct answers
Treatment	0.042*	-0.033*	0.057	0.049*	-0.028*	-29.914	-0.024	0.090
	(0.023)	(0.018)	(0.042)	(0.027)	(0.017)	(21.584)	(0.034)	(0.075)
Same respondent	-0.010	0.067	0.244**	0.175*	-0.021	-15.463	-0.142	0.072
	(0.051)	(0.068)	(0.106)	(0.093)	(0.042)	(28.969)	(0.097)	(0.179)
Same respondent x Treatment	0.136	0.027	-0.277*	-0.087	0.034	12.192	0.013	-0.146
	(0.104)	(0.100)	(0.160)	(0.137)	(0.062)	(37.602)	(0.141)	(0.295)
N	763	763	578	578	666	327	774	774
F test (Treatment | Same respondent)	3.116	0.004	2.037	0.082	0.008	0.331	0.007	0.038
p value	0.078	0.948	0.154	0.774	0.927	0.565	0.933	0.845

Notes: Each column lists the coefficient estimates from a different regression (apart from the constant). The column headings give the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. 'N. correct answers' is the number of correct answers in our quiz. 'Same respondent' is 1 the follow-up respondent's gender and age match the baseline respondent's gender and age, age+1, or age+2. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 %, respectively.

Table 14 Information sharing within the household II: HH size

	(1) Response in kl	(2) Bill hard to understand	(3) Reads consumption from bill	(4) Consumption accurate	(5) Tariff in ballpark	(6) Tariff error	(7) Increasing tariff	(8) N. correct answers
Treatment	0.063	0.010	0.001	0.000	-0.061*	-13.067	0.075	0.293*
	(0.040)	(0.034)	(0.090)	(0.058)	(0.033)	(40.686)	(0.071)	(0.165)
HH size	-0.004	0.015**	0.000	-0.017***	-0.003	3.970	0.002	0.060***
	(0.005)	(0.007)	(0.013)	(0.006)	(0.005)	(5.902)	(0.010)	(0.023)
HH size x Treatment	-0.006	-0.008	0.008	0.008	0.010	-1.971	-0.023	-0.054
	(0.007)	(0.008)	(0.019)	(0.011)	(0.007)	(6.783)	(0.015)	(0.036)
N	950	949	730	730	817	395	961	962
Treatment at mean HH size	0.038	-0.024	0.034	0.035	-0.016	-21.739	-0.025	0.055
	(0.020)	(0.016)	(0.036)	(0.023)	(0.015)	(17.531)	(0.030)	(0.065)

Notes: Each column lists the coefficient estimates from a different regression (apart from the constant). The column headings give the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. 'N. correct answers' is the number of correct answers in our quiz. 'HH size' is the number of individuals in the household. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 %, respectively.

Based on these measures, we do not see any evidence to suggest that our finding of no treatment effects on information is due to the lack of information sharing with households

9.3 Psychological effects

Chapter 6 describes various psychological effects that may be consistent with our findings of increased payments and no change in information. While our data does not permit us to identify all of these separately, we can test for one particular psychological channel.

One possibility is that the education visit acted as a reminder for the household about any outstanding bills. Spending 1/2-1 hour talking to the education officers is likely to have made water consumption in general more salient, and increased payments could have been a response to this. This could imply that the involvement of utility employees might not be important: perhaps simply sending reminders in the mail could have similar effects?

Table 15 Survey effects

Dep. var:	Jan - March total payment (1)	Jan - March payment (0/1) (2)	Jan - March payment frequency (3)	Jan - March avg. consumption (4)
		Panel A		
Control	0.010	0.005	0.010	-0.002
	(0.131)	(0.022)	(0.047)	(0.030)
N	985	985	985	962
		Panel B		
Treatment	0.259**	0.043*	0.098**	0.010
	(0.127)	(0.022)	(0.045)	(0.031)
N	988	988	988	970

Notes: Panel A estimates survey effects by comparing the control group (Control = 1) to 500 randomly selected households who did not participate in our study. Panel B compares our treatment group to this "new control group." Each column corresponds to a different dependent variable, and every regression controls for sampling strata indicators and the pre-treatment value of the dependent variable. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

Our data allows us to test this because in this case our surveys should also have increased the salience of water consumption and unpaid bills. Our surveys inquired at length about households' conservation and payment behavior, and even explicitly asked respondents to find their water bill and read out their consumption. If anything, this should have made payments even more salient than the education visits (where a household's own water bill was never explicitly discussed).[29] Because we have access to administrative data

[29] Similarly, Zwane et al. (2011) argue that surveying households can change their behavior.

for the entire population, we can address this question directly. 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 15, Panel A, we compare these households to our actual control group. Thus, 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. 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. By comparison, Panel B compares the new group and our treatment group. As expected, the results are numerically similar to those found earlier. Being selected to participate in our study and being surveyed did not affect behavior; the education visits did.

10 HETEROGENOUS TREATMENT EFFECTS

In this chapter we ask whether different groups of households may have been affected differently by our treatment. We chose to focus on five dimensions of heterogeneity: restricted status at baseline, indigent status at baseline, water consumption before the treatment, the respondent's education, and household income. The first three of these variables were used in our stratified sampling procedure because they are natural candidates for determinants of households' ability or willingness to respond to our treatment. For example, restricted households or those consuming low amounts of water may not be able to adjust their consumption by much, and indigent households may find it more difficult to increase their payments. We add income and education because they are obvious dimensions of heterogeneity, especially for an information campaign.[30]

Table 16 looks at payment and consumption. Each panel interacts our treatment indicator with one of the five variables mentioned above. In each case, a test of heterogenous treatment effects is equivalent to asking whether the interaction term is statistically significant. We find evidence of heterogeneity for three variables. In Panel C, higher consumers appear to have lowered their consumption in response to the treatment, while low consumers may have increased it. The effect of the treatment is significantly positive below 6 kl (log average consumption = 1.8) and negative for consumption exceeding 30 kl (log average consumption = 3.4). This is in line with our findings from Table 4 above, where we saw a reduction in consumption for the highest consumption quartile.

[30] To maximize our sample size, we use income and education measures from the follow-up survey.

Panels D and E of Table 16 show evidence that the average increase in payment amount and propensity found earlier comes from more educated and higher-income households. In Panel D, more educated households significantly increased both their payment amount and their propensity to pay, while less educated households did not change their behavior in response to the treatment.[31] Similarly, in Panel E, higher income households were more likely to pay and paid more.[32] This makes sense: our treatment increased payments among those who are more able to pay.

Table 17 adds the five interactions, one at a time, to the information regressions. Did the more educated and higher income households benefit more from our treatment (in terms of increased knowledge)? In Panel E, we do not see any heterogeneity by income, while in Panel D, if anything, it is the less educated who show evidence of increased knowledge. In column (1), the less educated are more likely to respond in kilolitres, and in column (4), they show an increased ability to tell their consumption from the bill. Thus, while our treatment shows some impact on the information of specific groups, our earlier conclusion remains: the change in information was not responsible for the increase in payments.[33]

[31] We measure education by whether the respondent completed high school. The share of such respondents is 58% in the control and 57% in the treatment group.
[32] The increase in amount paid becomes statistically significant at a household income of 6300 Rand, which just below the median of 6600. The increase in payment propensity becomes significant at 7800 Rand.
[33] In Table 17, Panel B, indigent households are also more likely to be able to tell their consumption from the bill. As we saw in Table 16, the treatment did not impact the payment or consumption behavior of these households differentially. In Panel C of Table 17, there is some improvement in the information of low consumers.

Table 16 Heterogeneous treatment effects on payment and consumption

	(1) Payment amount	(2) Payment propensity	(3) Payment frequency	(4) Consumption
	Panel A: Restricted			
Treatment	0.227	0.035	0.075	0.029
	(0.158)	(0.027)	(0.058)	(0.036)
Interaction	0.294	0.045	0.071	-0.066
	(0.308)	(0.053)	(0.101)	(0.070)
Restricted	-0.728***	-0.131***	-0.214***	0.001
	(0.226)	(0.039)	(0.076)	(0.048)
N	966	966	966	947
	Panel B: Indigent			
Treatment	0.334**	0.043	0.078	0.016
	(0.159)	(0.027)	(0.057)	(0.037)
Interaction	-0.046	0.023	0.065	-0.021
	(0.309)	(0.055)	(0.106)	(0.067)
Indigent	-0.063	-0.028	-0.046	0.041
	(0.231)	(0.041)	(0.079)	(0.045)
N	966	966	966	947
	Panel C: Pre Consumption			
Treatment	0.740	0.097	0.222	0.270*
	(0.514)	(0.090)	(0.159)	(0.138)
Interaction	-0.173	-0.020	-0.051	-0.103**
	(0.192)	(0.033)	(0.061)	(0.051)
Pre Consumption	0.426***	0.051**	0.121**	0.694***
	(0.146)	(0.024)	(0.048)	(0.031)
N	947	947	947	947
	Panel D: Education			
Treatment	-0.010	0.001	0.050	0.002
	(0.199)	(0.035)	(0.073)	(0.046)
Interaction	0.586**	0.085*	0.092	0.013
	(0.274)	(0.048)	(0.097)	(0.062)
Education	-0.243	-0.039	0.037	0.030
	(0.202)	(0.035)	(0.073)	(0.041)
N	960	960	960	941
	Panel E: Income			
Treatment	0.038	-0.007	0.056	0.043
	(0.210)	(0.037)	(0.073)	(0.048)
Interaction	0.032*	0.006*	0.005	-0.003
	(0.019)	(0.003)	(0.007)	(0.004)
Income	0.004	-0.000	0.008	0.009***
	(0.013)	(0.002)	(0.005)	(0.003)
N	857	857	857	844

Notes: Panels A-E investigate heterogeneous treatment effects by different grouping variables. 'Restricted' is 1 if the consumer was restricted at baseline. 'Indigent' is 1 if the consumer was registered as indigent at baseline. 'Pre Consumption' is average consumption in the 3 months before the treatment (in logs). 'Education' is 1 if the follow-up respondent has completed high school and 0 otherwise. Income is total household income in 1000 Rand at follow-up. 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. 'Consumption' is average consumption in the 3 months following the treatment (in logs). All regressions control for 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.

Table 17 Heterogeneous treatment effects on information

	(1) Response in kl	(2) Bill hard to understand	(3) Reads consumption from bill	(4) Consumption accurate	(5) Tariff in ballpark	(6) Tariff error	(7) Increasing tariff	(8) N. correct answers
Panel A: Restricted								
Treatment	0.031	-0.018	0.017	0.035	-0.020	-24.676	-0.033	0.079
	(0.024)	(0.019)	(0.043)	(0.028)	(0.015)	(22.039)	(0.036)	(0.077)
Interaction	0.025	-0.019	0.066	0.005	0.011	19.581	0.035	-0.075
	(0.043)	(0.033)	(0.079)	(0.052)	(0.037)	(24.942)	(0.064)	(0.142)
Restricted	-0.019	-0.002	-0.060	-0.003	0.032	-45.110**	0.018	-0.033
	(0.027)	(0.026)	(0.055)	(0.034)	(0.027)	(22.319)	(0.045)	(0.097)
N	953	952	731	731	820	396	964	965
Panel B: Indigent								
Treatment	0.042*	-0.038**	-0.018	0.008	-0.028	-2.831	0.007	0.070
	(0.023)	(0.018)	(0.043)	(0.026)	(0.018)	(8.385)	(0.035)	(0.076)
Interaction	-0.011	0.049	0.186**	0.099*	0.037	-79.015	-0.101	-0.047
	(0.043)	(0.035)	(0.079)	(0.056)	(0.030)	(67.470)	(0.065)	(0.143)
Indigent	0.005	-0.020	-0.099*	0.010	-0.029	77.600	0.045	0.082
	(0.029)	(0.025)	(0.054)	(0.035)	(0.022)	(66.399)	(0.045)	(0.097)
N	953	952	731	731	820	396	964	965
Panel C: Pre Consumption								
Treatment	0.105*	-0.140**	0.302**	0.061	0.012	29.809	-0.019	-0.136
	(0.059)	(0.056)	(0.126)	(0.067)	(0.047)	(24.112)	(0.102)	(0.211)
Interaction	-0.025	0.045**	-0.104**	-0.009	-0.011	-21.277*	-0.000	0.075
	(0.024)	(0.019)	(0.048)	(0.024)	(0.018)	(11.038)	(0.039)	(0.079)
Pre Consumption	0.045***	-0.053***	0.098***	-0.008	0.009	16.547*	-0.024	-0.057
	(0.017)	(0.015)	(0.035)	(0.016)	(0.015)	(9.518)	(0.030)	(0.059)
N	934	933	719	719	801	382	945	946
Panel D: Education								
Treatment	0.102***	-0.028	0.087	0.085**	-0.026	-54.035	-0.051	-0.021
	(0.033)	(0.027)	(0.056)	(0.037)	(0.021)	(37.991)	(0.043)	(0.097)
Interaction	-0.115***	0.006	-0.098	-0.085*	0.016	57.549	0.045	0.122
	(0.041)	(0.033)	(0.073)	(0.048)	(0.029)	(39.549)	(0.059)	(0.130)
Education	0.006	-0.038	-0.108**	0.019	0.000	-46.624	-0.089**	-0.047
	(0.026)	(0.025)	(0.052)	(0.031)	(0.023)	(38.102)	(0.041)	(0.088)
N	948	946	728	728	816	393	958	959
Panel E: Income								
Treatment	0.052	-0.021	0.037	0.083*	0.006	-23.625	-0.040	0.036
	(0.035)	(0.022)	(0.056)	(0.044)	(0.033)	(14.859)	(0.045)	(0.095)
Interaction	-0.001	0.001	-0.000	-0.005	-0.001	0.896	0.003	0.001
	(0.003)	(0.002)	(0.005)	(0.004)	(0.004)	(0.783)	(0.003)	(0.007)
Income	0.000	-0.001	0.004	0.001	0.006*	-0.970*	0.007***	0.014***
	(0.002)	(0.001)	(0.004)	(0.003)	(0.003)	(0.567)	(0.002)	(0.005)
N	846	843	657	657	735	351	855	856

Notes: Panels A-E investigate heterogeneous treatment effects by different grouping variables. 'Restricted' is 1 if the consumer was restricted at baseline. 'Indigent' is 1 if the consumer was registered as indigent at baseline. 'Pre Consumption' is average consumption in the 3 months before the treatment (in logs). 'Education' is 1 if the follow-up respondent has completed high school and 0 otherwise. Income is total household income in 1000 Rand at follow-up. The columns in each panel correspond to separate regressions. The column headings give the dependent variable. 'Response in kl' is 1 if the respondent's guess about their consumption is stated in kilolitres. 'Reads consumption from bill' is 1 if the respondent was able to find a water bill and reads out their consumption from the bill. 'Consumption accurate' is 1 if this number matches any consumption in the administrative data from the prior 6 months. 'Tariff in ballpark' is 1 if the respondent's guess about the kilolitre price is between 5-25 Rand. 'Tariff error' is max (0, the respondent's guess about kilolitre price - 25). 'Increasing tariff' is 1 if the respondent understands that the tariff schedule is increasing. 'N. correct answers' is the number of correct answers in our quiz. Robust standard errors in parentheses. ***, **, * denote significance at 1, 5, and 10 percent, respectively.

11 CONCLUSION

We implemented and evaluated an information campaign as a potential response to non-payment for water in South African townships. Our education visits had a substantial impact, reducing the fraction of households making no payments by 4-5 percentage points and increasing the amount of payments by approximately 30% over a three-month period. We find no effect on average consumption, but treated households report an increase in conservation practices, which is consistent with households substituting more water-intensive activities with less water-intensive ones. We do find a reduction in quantities consumed for the highest consumers: those consuming above 19 kl reduced their monthly consumption by 9.5% on average.

These effects are not driven by increased knowledge. Our treatment had modest effects on information, with a small increase in households' familiarity with the word "kilolitre" and an increased sense of the amount of water used when flushing the toilet. Overall, treated households are not much more likely to understand quantities of water used or their water bill than households in the control group. This suggests a psychological explanation for the effect of our treatment on consumers' behavior: households were "nudged" by the education visits or reciprocated the providers' efforts by paying more. Consistent with this, we find that involvement of the providers' employees - as opposed to simply the visit of our independent surveyors – was crucial to achieve the increase in payments. These findings show that public information campaigns may generate unintended consequences, including psychological responses, that can improve their effectiveness.

Our findings of little information effects raise a number of interesting questions for future research. First, although we did not find large increases in information, feedback from the households clearly indicates that the education visits and the distributed materials were viewed as valuable. In particular, 89% of treated households described the information brochures as "very useful" or "fairly useful" (8.4% said that they were a little useful and only 2.6% that they were not at all useful). Such positive attitudes were likely important for the changes in behavior that we found. For example, it is unlikely that households would have reciprocated education visits they did not view as valuable by increasing their payments. This suggests that defining and measuring the "quality'' or "effectiveness" of an education programme is a subtle exercise. Whether consumers "learned" the information given to them based on conventional measures is not always relevant for evaluating the effectiveness of such programmes.

Second, the impact of any education programme potentially depends on a large set of factors. In our case, this includes among other things: the original WRC materials we adapted for the project, the adaptation itself, the translation of the materials into the local language, the training of the education officers, the delivery of the education visits by the education officers, etc. Our project was designed to measure the impact of the entire "bundle" of factors encompassed by this particular programme. Whether changing a specific factor would change the results, including the effects on information, consumption, and payment, is a relevant question for future research. In the long-run, educating consumers is crucial to improve households' water management. Our findings highlight the need for more research to carefully evaluate various information vehicles and to select the ones most likely be effective at improving consumer knowledge.

12 LIST OF REFERENCES

Abrahamse, W., L. Steg, C. Vlek, and T. Rothengatter (2005): "A review of intervention studies aimed at household energy conservation," *Journal of Environmental Psychology*, 273-291.

Ahluwalia, M.S. (2002): "Economic Reforms in India since 1991: Has Gradualism Worked?" *Journal of Economic Perspectives* 16(3), 67-88.

Allcott, H. (2011): "Social norms and energy conservation," *Journal of Public Economics* 95, 1082-1095.

American Water Works Association (1999): *Residential End Uses of Water*, Denver, CO: AWWA Research Foundation.

Ayres, I., S. Raseman, and A. Shih (2013): "Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage," *Journal of Law Economics and Organization* 29(5), 992-1022.

Banerjee, A., and E. Duflo (2012): *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, PublicAffairs.

Duflo, E., R. Glennerster, and M. Kremer (2008): “Using Randomization in Development Economics Research: A Toolkit,” *in*: T. P. Schultz and J. A. Strauss (ed.): *Handbook of Development Economics*, Elsevier.

Ferraro, P.J. and M. K. Price (2013): "Using Nonpecuniray Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment," *The Review of Economics and Statistics*, 95(1), 64-73.

Fielding, K.S., S. Anneliese, S. Russell, R. McCrea, R. Stewart, and J. Gardner (2013): “An experimental test of voluntary strategies to promote urban water demand management,” *Journal of Environmental Management* 114, 343-351.

Hollingworth, B., P. Koch, S. Chimuti, and D. Malzbender (2011): *An investigation into the water infrastructure development financial allocation pathways in municipalities*, WRC Report No. TT 476/10, Pretoria, South Africa: Water Research Commission.

Hosking, P., K. Jacoby, G. Sharp, and J. Hosking (2011): *Investigating the Mechanism and Processes Used in Setting Water Services Tariffs*, WRC Report No. 1871/11, Pretoria, South Africa: Water Research Commission.

Jessoe, K. K. and D. Rapson (2013): "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use," *American Economic Review*, forthcoming.

Lampietti, J.A., S.G. Banerjee, and A. Branczik (2007): *People and Power : Electricity Sector Reforms and the Poor in Europe and Central Asia*, Washington, D.C.: The World Bank.

McRae, S. (2013): "Infrastructure Quality and the Subsidy Trap," *American Economic Review*, forthcoming.

Reiss, P. and M. W. White (2008): "What Changes Energy Consumption? Prices and Public Pressures," *The RAND Journal of Economics* 39(3), 636-663.

Republic of South Africa (2011): *Intergovernment Fiscal Review 2011*, Pretoria, South Africa: National Treasury.

Selmrod, J. (2007): "Cheating Ourselves: The Economics of Tax Evasion," *Journal of Economic Perspectives* 21(1), 25-48.

Statistics South Africa (2012): *Income and Expenditure Survey (IES) 2010/2011 Press Statement,* 6 November 2012, Pretoria, South Africa.

Szabó, A. (2013): "The Value of Free Water: Evaluating South Africa's Free Basic Water Policy," working paper, University of Houston.

White, I.R., and S.G. Thompson (2005): "Adjusting for partially missing baseline measurements in randomized trials," *Statistics in Medicine* 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, Washington, DC: The World Bank.

Zwane, A.P., J. Zinman, E. Van Dusen, W. Pariente, C. Null, E. Miguel, M. Kremer, D.S. Karlan, R. Hornbeck, X. Gine, E. Duflo, F. Devoto, B. Crepon, and A. Banerjee (2011): "Being surveyed can change later behavior and related parameter estimates," *Proceedings of the National Academy of Sciences* 108(5), 1821-1826.

APPENDIX

1. WRC RESOURCES USED FOR DEVELOPING THE EDUCATION MATERIALS

We drew from a series of lesson plans from Grade R to 10 developed by the WRC In support of learning and teaching about water and water-related issues. These are available at http://www.wrc.org.za/Pages/Learning_School_lessonplans.aspx.

In addition, we also used specific content from the following WRC reports:

- Peddie, C., D. Hibbert and C. Conway-Physick (2008): *Learning and Teaching About Water In Our Classrooms: A Series Of Lesson Plans For Grades R – 7*, WRC Report No TT 345/08, Pretoria, South Africa: Water Research Commission.
- Peddie, C., D. Hibbert and C. Conway-Physick (2008): *A Series of Lesson Plans for Grades 8 – 10*, WRC Report No TT 346/08, Pretoria, South Africa: Water Research Commission.
- Sarah Slabbert Associates (2010): *Guidelines On Domestic Water Accounts – Towards A Consistent Approach in the RSA*, WRC Report No. TT 457/10, Pretoria, South Africa: Water Research Commission.
- Sarah Slabbert Associates (2010): *Towards Standards for Municipal Invoices in South Africa*, WRC Report No. TT 458/10, Pretoria, South Africa: Water Research Commission.

2. PHOTOS FROM THE PROJECT

Photos from the surveys

Fieldworkers attended a full day workshop in March 2012 where they familiarised themselves with the purpose of the project and the questionnaire. We went through all questions in the questionnaire and made sure everyone understood the information to be collected and how it should be recorded. The workshop also included role-playing a household visit to fine-tune the interview process. We wanted to make sure that household visits would be as uniform as possible. Interviewers were young people in their 20s. Most of them live in the survey area, and all of them have extensive experience conducting household visits in this type of setting.

GPS coordinates were collected for the location of each household. Households' location relative to each other was used to assess any spillover effects from the education intervention.

Typical properties in the area are relatively small; the average household size is 4.3 persons. Many housing units are uniform single-family buildings provided by the government with limited modifications made by the residents. Thus, living conditions are fairly similar within our sample.

Each participating household has water-using sanitation and tap(s) inside the house or outside on the property. Consumption is metered individually by a meter located on the property and easily accessible to the household.

Photos from the education programme

Education officers delivering the education programme were trained by us specifically for this project. They are employed by Odi Water and regularly make household visits in our area of study.

The training included role-playing a household visit.

During the education programme, education officers were supervised by Odi Water's marketing department.

The education officers visited the households to give them the education brochures and the activity book. They explained the contents of each, emphasizing a list of specific items (for example, how the dial on the water meter should be read).

Examples of education materials

Selected pages from the Activity Books

Water Audit

Auditing your family's water consumption is a good way to find out how you can avoid wasting water and how you can reduce your bill.
This booklet will help you make a Water Diary and compute how much you spend on water through your everyday activities.
To make the Water Diary, you will need to measure your water consumption over a 3 day period. This is a simple and fun activity that you and your family can do together.

Once your audit is complete, we invite you to return the booklet in exchange for a cash reward.
We simply ask that you complete everything to the best of your knowledge. Since every household is different, there are no "right" or "wrong" answers. All completed booklets will receive the reward.

We hope you are going to enjoy saving some water today!

3

Measure each activity and record approximately how many liters of water you used.

Washing hands

.................ℓ

Hand washing clothes

.................ℓ

Garden hose (litres per minute)

.................ℓ

Kettle

.................ℓ

brushing teeth

.................ℓ

filling up your kitchen sink

.................ℓ

People often drink with the tap running. In this way, fresh water is wasted. You can save water by using a cup. Try it both ways and measure the water you use while drinking.

Water used with the tap left running

.....................ℓ

Water used with a cup

.....................ℓ

6

CALCULATE YOUR WATER CONSUMPTION

Calculate your water consumption using both methods below.

Method 1: Copy the numbers from page 8 with the initial and final reading of your water meter. Calculate the difference.

____________________ Reading at the end of day 3

— ____________________ Reading at the beginning of day 1

= ____________________ Difference

The difference is your water consumption in kiloliters (1000 liters) throughout the 3 days.

Method 2: Add together the amounts of water used from your water diary. Use a calculator. Write down the result here.

3-day consumption from the diary:

________________________________ liters.

This is your water consumption in liters. The meter measures your water consumption in kiloliters, so you need to divide this amount by 1000 to compare. For example, if you used 5345 liters of water, then your water consumption in kiloliters is 5345 / 1000 = 5.345 kiloliters.

On the next page, we invite you to compute the cost of your water consumption.

Date	Time	What did I do?	How much water did I use?
Monday, June 1	7:30 am	Flush toilet	8 liters

3. PROJECT TEAMS AND FUNDING

This project was a cooperation between the University of Houston (UH), Odi Water Services (ODI), and Impact Research International (IRI).

Andrea Szabo and Gergely Ujhelyi (UH) were the principal investigators who developed and oversaw all aspects of the project (including the survey, education programme, and analysis). Patrick Chiroro (IRI) led the IRI team during the surveys. Fieldworkers, data collection and cleaning were supervised by Tronny Mawadzwa (IRI). Fieldworkers for the surveys were employees of IRI trained by the PIs. The education programme was developed by the PIs using materials from the Water Research Commission. Sipho Nkosi (ODI) supervised the education programme at Odi Water. The education officers were supervised in the field by Thabo Kau (ODI) and Ganatius M Loate (ODI). Pieter Avenant (ODI) oversaw the financial aspects of the education programme at Odi and provided the billing data.

The project was supported by the University of Houston, Odi Water, and the Water Research Commission. The principal investigators' involvement was funded entirely by the University of Houston. Odi Water employees were funded entirely by Odi Water. IRI's involvement was funded by UH for the baseline survey and by the Water Research Commission for the follow-up survey and for specific tasks related to the education programme.