Behavioural Nudges as a Water Savings Strategy

Report to the
Water Research Commission

by

Grant Smith and Martine Visser
(School of Economics, University of Cape Town)

WRC Report No. 2091/1/13
ISBN 978-1-4312-0508-0

February 2014

Obtainable from

Water Research Commission
Private Bag X03
GEZINA, 0031

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

Executive Summary

Many municipal areas in South Africa are water stressed for portions of the year and find themselves needing to manage the demand on their water resources. This need is likely to become more acute for large urban municipalities as South Africa seeks to facilitate economic growth and as the population continues to concentrate within the larger urban municipalities where household water demand already constitutes a significant share of total water demand.

Historically municipalities have largely relied upon tariffs, technical interventions (such as leakage control) and customer education campaigns to manage the demand for water. In recent years however municipalities, mainly in the developed world, have begun experimenting with strategies that are informed by the field of behavioural economics in order to manage the demand upon their utilities. These strategies have yielded some encouraging results and are gaining traction around the world. The attraction of these methods is that they are generally very cheap to deploy, require limited infrastructure and offer few, if any, opportunities for corruption.

Behavioural Economics is a discipline within economics that seeks to expand the classic model of rational human decision-making within economics by allowing for the various biases and heuristic foibles that characterise people as they really are. This is helpful for policy design insofar as that policy needs to interact with human decision-makers. Insights from the discipline are being increasingly used to craft policy in order to achieve policy goals through encouraging behaviour change amongst the target population. The attraction of these policies for utilities is twofold. First they tend to place a relatively little extra burden upon a utility since they tend to not be very expensive to run and require very little additional administrative engagement upon the part of the utility. Second they work through the way people choose rather than by removing choice, which makes them constitutionally attractive.

In the utility demand management space behavioural economics has been applied mainly to managing electricity demand (Alcott, 2011 for example) although there have been some efforts made with water demand specifically (Ferraro and Price, 2011). So far these sort of programs have been constrained to the developed world and typically these projects have employed three behavioural mechanisms to manage demand; these are: raising the salience of the consumer's consumption to themselves, comparing the consumer's consumption to some social norm and providing information about how to consume less. In all previous studies we are aware of these elements have not been tested separately so as to determine their individual effectiveness.

This study assessed whether there was scope to use feedback, informed by several principles derived from the behavioural economics literature, delivered in the post with the water bill to reduce household water consumption within the city of Cape Town.

Eight different sorts of feedback were assessed against each other and a further control group (receiving no feedback apart from their normal water bill) by means of a Randomised Control Trial (RCT) run in Cape Town in November 2012 and involving over 280 000 households. Households were allocated to the control or one of the eight feedback groups in a random fashion within meter reading areas. Only households that used more than Free Basic Water or less than two kilolitres per

day on average for the past year were considered for this study and randomly allocated to a feedback or control group.

The interventions aimed to test the effectiveness of three major behavioural elements: increasing the salience (understanding as well as awareness) of the household's own consumption, comparative social norms of consumption (by comparing households to their average neighbour) and tips about how to save water.

In the study, raising the salience of a household's water consumption, either by reporting their own previous consumption in a bar graph or by comparing their consumption to their neighbour via a bar graph, was found to lower water consumption significantly. The study was novel in that it is the first of its kind to isolate the effects of social norms and raising salience through separate treatments. It found the use of social norms does not induce significantly greater behavioral adjustments than raising salience does. For the 3 month post-treatment period for which data was available, it was found that reporting the last month's consumption was more effective than reporting the entire last year's consumption, indicating that it might be more effective to simplify the message being communicated as much as possible. The results of the analysis further strongly suggest that merely reporting tips or information about how to save water will not result in a noticeable reduction in household water usage.

In addition, this is the first study that that the research team know of that illustrates the impact of raising salience and social norms on demand for water for a *developing country*. Assessing such strategies in the developing world is important since the infrastructure within developing world municipalities is likely to be different to that in the developed world in that it is more heterogeneous and may allow less control overall.

The findings are encouraging in that they strongly suggest that behavioral elements may be leveraged in order to reduce household water consumption at the municipal level. Although the savings are roughly 1% (4 l per household per day in average) of total water consumption for this sample, they are significant. Considering the total annual consumption of potable water in Cape Town is 300 million Kl, this constitutes about a 0,5% of total demand. In the post treatment month of January the savings of water directly attributable to the 278 681 non-control households analysed in this project was 27 753 683 litres; for the shorter month of February this was a direct savings of 23 845 021 litres.

The benefits of rolling out a behavioural intervention such as this, seems cost effective when one considers that the field roll-out costs of the study came to about R100 000, whereas the total savings to households, over the two months for which we measured the impact of the once-off intervention, came to R586 000 at the R10,60/kl (the 2012 tariff for households consuming between 10-20 kl per month) tariff rate. Across a municipality, such savings would certainly be meaningful in relatively water scarce months. Considering that this was only a once off intervention and that this study identifies strategies which performed much better than others, it is very likely that greater savings are possible if an appropriate mix of these strategies is employed and if the intervention is repeated for several months. Overall this study suggests that behavioural economics offers a new and powerful set of tools with which municipalities may manage their domestic water demand.

Acknowledgements

The authors wish to acknowledge the financial support of the Water Research Commission for this project, the support of the project reference group – in particular Mr Jay Bhagwan. We also wish to thank the City of Cape Town, in particular the Director of Water and the Director of Billing, for allowing us permission to access the data required to set up this project as well as access to the billing process in order to conduct this fieldwork for this study.

Table of Contents

List of Figures

List of Tables

Introduction and Overview

The need for municipalities to develop programs that manage household and industry water demand is widely attested. As South Africa continues to grow in her urban centres the resources which municipalities draw on to satisfy household and industrial demand will only become further stressed. Traditional water demand strategies (outside of crisis periods) have tended to employ price, pressure management or information/awareness campaigns individually or in combination. Recently however the discipline of Behavioural Economics has added some further options that seem to hold promise for water demand management.

Behavioural Economics is a discipline within economics that has grown phenomenally within the last three decades. It seeks to expand the classic model of rational human decision-making within economics by allowing for the various biases and heuristic foibles that characterise people as they really are. This is helpful for policy design insofar as that policy needs to interact with human decision-makers. Insights from the discipline are being increasingly used to craft policy in order to achieve policy goals through encouraging behaviour change amongst the target population.

The attraction of these policies for utilities is twofold. First they tend to place a relatively little extra burden upon a utility since they tend to not be very expensive to run and require very little additional administrative engagement upon the part of the utility. Frequently a system can be set up that can run nearly automatically. Second they work through the way people choose rather than by removing choice, which makes them constitutionally attractive.

In the utility demand management space behavioural economics has been applied mainly to managing electricity demand (Alcott, 2011 for example) although there have been some efforts made with water demand specifically (Ferraro and Price, 2011). So far these sort of programs have been constrained to the developed world. Encouragingly these efforts seem to have realised small but significant reductions in utility consumption. Typically these projects have employed three behavioural mechanisms to manage demand; these are: raising the salience of the consumer's consumption to themselves, comparing the consumer's consumption to some social norm and providing information about how to consume less. In all previous studies we are aware of these elements have not been tested separately so as to determine their individual effectiveness.

This project set out to determine whether there was scope to apply the tools of behavioural economics to South African water bills in order to manage water demand. Furthermore, the researchers also wished to determine the effect of specific behavioural elements by testing each of these separately in addition to the more conventional combinations.

In order to investigate these questions a large-scale randomised control trial (RCT) was conducted in November 2012 in the City of Cape Town with the cooperation of the City of Cape Town's Department of Water and the City's Billing department. The study involved over 340 000 households.

The RCT had eight different treatment groups and a control group. The treatments all entailed including A4 inserts with households water bills, sometimes two A4 inserts. These inserts carried different sorts of information aimed at influencing the consumption of each household. Between the eight inserts the elements of salience (of the household's own water consumption), comparative social norms and information about how to save were delivered individually as well as in bundles of interest (such as social norms and salience). This allowed us to be able to identify the individual effect of each behavioural element. So far as the researchers are aware, this is the first project to systematically unbundle salience-of-own-consumption effect from social norms and thus the first study to be able to identify their respective individual effects. The control group received no inserts, but received only their usual monthly water bill.

In order to simplify analysis, the inserts were sent only once (in November to early December) to most treatment households[1]. In order to assess the effect of each different type of treatment the difference in average daily consumption for each month was compared between the various treatments. Comparisons were made in terms of simple comparisons of means as well as regression analysis. The validity of inferences drawn in this analysis depend critically upon the fact that all households stood an equal probability of being in any one of the eight different treatment groups or the control group and that no observable bias existed between the various treatment groups and the control before the mailing of the inserts.

The results of this analysis, drawn from such a comprehensive sample of households in a large municipality provide, in the researchers' minds, helpful guidance as to the possibility of incorporating consumption feedback informed by behavioural economics into the broader water demand management toolkit that municipal managers have at their disposal.

1 A relatively small minority of households (around 2000) received two or more mailings.

Chapter 1: Behavioural Economics and reducing water consumption

Behavioural Economics has found an increasing role in public policy in recent years. The utility of incorporating insights from the discipline is often described as the ability to produce policy interventions that "nudge"[2] people into a behaviour that is in their best interests, but which they, for some reason or another, consistently do not adopt. This project is an exercise in constructing various potential nudges, delivered via the water bill, and assessing their ability to influence household water consumption.

In this chapter we will examine various behavioural elements that inform these strategies. We group these elements into three major categories: Salience, Social norms and Mapping.

1.1 Salience

An important behavioural element to consider at the beginning of this report is that of salience. People seem to have limited cognitive resources, we cannot think about everything all at once. The implication for decision-making and behaviour is that we find it difficult to marshal all of the relevant facts and then weigh them up in order to make a decision. To do this takes time and a very conscious effort. For most day to day decisions, especially those that are taken quickly, people will tend to consider information that comes to mind with ease and to use that information to make their decisions. This process is widely noted in behavioural economics and psychology (see Kahneman, 2003 for an overview) and is often described as people incorporating that which is salient into their (intuitive or reflective) decision-making process, while excluding that which is not salient from their decision-making process. As a result, salience then influences behaviour.

There are many ways in which something could be made more salient, but perhaps the most simple would be to simply present information about that thing in a more noticeable manner. An intuitive example of this (from Levitt and Dubner, 2009) might be a person's relative fear of shark attacks keeping them from swimming in the sea, while they have no qualms driving, although the likelihood of injury while driving is much higher than from a shark attack while swimming. An explanation for this is that shark attacks, when they happen, are often front page news while road accidents are, mainly, not. As a result, shark attacks may be more salient to many people's minds than car accidents with the result that many people may be too afraid to swim while they feel fine about driving.

More formally, a substantial literature already exists within economics that explores the relationship between salience and behaviour. In simplicity this literature describes a relationship between salience and behaviour that is positive and strong. If something is more salient, it will be a more potent influence upon behaviour. This result has been confirmed several times in laboratory settings (Chartrand *et al.,* 2008; Sela and Shiv, 2009; LeBoeuf, Shafir, and Belyavsky Bayuk, 2010) as well as in the field.

Taxation has provided a rich source of data from the field on the relationship between salience and behaviour in recent years (Congdon et al., 2009 includes a useful overview). Chetty, Looney and Kroft (2009) find that customers respond much more to tax on goods that is included in the good's marked price compared to tax charged at the till. Finkelstein (2009) finds that with electronic toll collection drivers do not notice tolls as much as they do when they need to manually count out the toll, with the result that electronic tolls settle at significantly higher steady state levels. Further

2 The title of the paradigm-setting book by Thaler and Sunstein, 2008.

research by Saez (2009) as well as Chetty and Saez (2009) finds similar behaviour for normal income tax payers, who do not appear to respond with tactically optimal behaviour to a tax code that is fairly complex and difficult to understand and thus not very salient to normal income tax payers.

Further examples from the field include Stango and Zinman (2011) who find that the mere fact of answering surveys whose topic is an overdraft of a cheque account significantly lower the incidence of overdrafts amongst cheque account holders who took the survey compared to cheque account holders who did not take the survey. Similarly Zwane *et al.* (2011) find that being surveyed about health and finances significantly raises the likelihood of a person using water purification and purchasing insurance. Sexton (2011) finds that signing up to a program that automatically debits a person's account to pay for their electricity consumption actually increases consumption, whereas enrolment in a more participatory program reduces consumption by 5% on average.

With this in mind, if we wish to influence the water consumption behaviour of households and businesses a reasonable course of action would be to make strategies and concepts around water consumption generally more accessible.

However this does beg the question of might make water consumption less salient. Not reporting water consumption at all is certainly the most obvious way in which a low salience state of consumption could be encouraged. Apart from not reporting consumption, there appear to be other ways in which the salience of water consumption may be reduced. In this regard the literature on choice overload and information overload provide helpful insight. Since we are discussing water consumption rather than explicit choices between various products or other options (the main subject matter of the choice overload literature) only portions of that literature are potentially relevant here. The portions of the literature that are relevant are those which, like much of the information overload literature, deal with increasing or decreasing the cognitive burden of processing information. In simplicity, decreasing the cognitive burden of processing information is more likely than not to enable the subject of that information to become salient. In a meta-analysis of the choice overload literature Scheibehenne, Greifeneder and Todd (2010) note several aspects of choice and information presentation that are likely to influence the cognitive load of processing information and thereby the salience of that information's subject. Of these, two aspects are especially relevant to this research project: the amount of information presented and categorisation.

Earlier studies have shown that increasing information does not necessarily lead to more optimal choice (and behaviour) but rather seems to place an increasingly heavy cognitive burden upon the person by asking them to consider an increased data set from which to infer their best course of action. Put another way, increasing the amount of information supplied leads to a situation where at some point there is too much to consider in order to make a correct choice. In such cases, Lee and Lee (2004) and Lurie (2004) find that people make less informed choices. This state of affairs seems to lead to people feeling less satisfied in at least one study (Greifeneder, 2010). Overall, it seems that while presenting more information about water would entail more thinking, presumably related to water in some way, it would be tending away from clarity about one's own consumption and towards confusion and thus decreased salience of one's water consumption.

Unclear categorisation of information may also work against salience and optimal choice of action. A lack of clear categorisation of the information presented to people makes that information difficult to navigate and understand (Diehl, 2005; Kornish and Lynch, 2003, Huffman and Kahn, 1998; Russo, 1997). This would presumably make it more difficult to integrate the information into one's broader life and take appropriate action. In the choice overload literature, Mogilner et al. (2008) find a negative interaction between satisfaction and the amount of information only when that information was not clearly categorised.

Slabbert (2010a, 2010b) describes the state of many water bills in South Africa as being characterised by the factors described above. Categories are unclear, information is presented in a somewhat overwhelming manner and technical language is often used. This, as Slabbert contends, makes South African water bills difficult to understand. In this context it is quite likely that the amount of water consumed is something that is not salient to households. As a result, water consumption may well feature less prominently in household decision-making than it should.

Part of this research project thus aims to assess whether there is scope for mitigating water consumption by simply raising the salience of a household's or business' water consumption.

1.2 Social norms

The next behavioural topic we wish to discuss is that of social norms. There is a growing literature within behavioural economics which concerns itself with understanding the ability of social norms to prompt social change. Much of this literature draws on insights gained in earlier research in the field of social psychology. In this section we discuss four aspects of social norms that seem to influence their potency: the size of the referent group (section A), the social distance between the person/group and the referent group (section B) and the presence of pluralistic ignorance (section C) and the potential value of social norms as signals of important information generally (section D).

A. Size of the referent group

Since at least the work of Milgram, Bickman & Berkowitz (1969) the literature on social influence in social psychology has provided evidence that people are more likely to do something if more people are doing it than if fewer people are. This appears to be true whether there is a cooperative element to the behaviour, such as with contributing to public goods by paying income tax, or if straightforward adoption is what is at stake.

Milgram, Bickman & Berkowitz (1969) provide one of the canonical papers in regard to the influence of a crowd. Conducted upon a New York sidewalk, the experimenters got first one, then two and finally ten confederates to stand still and look up. In each case a crowd gathered before the original gazer left and the crowd dispersed, to be reconstituted later around a larger core of experimental confederates. The key result of this experiment was the fact that the size of the crowd that gathered around the initial gazing confederates was larger if more initial gazers were planted and smaller if fewer initial gazers were provided. This has echoes of earlier work still by Asch (1962) and Sherif (1936) who demonstrate the effect of others upon an individual (Asch) or group (Sherif) can extend to eliciting patently false judgements about reality. In more recent times this tradition of enquiry remains vibrant with various research projects.

Salganik, Dodds and Watts (2006) show, in a randomised field experiment, that the likelihood of people liking and downloading songs is significantly determined by how many people they can see have liked and downloaded each song before them.

Situational evidence from the marketing literature (notably McPhee, 1963; Ehrenberg, 1972; Ehrenberg et al., 1990) suggests a similar process may be at work in the shopping malls. These researchers report that brands with a greater initial market share are more likely to be bought more often (as opposed to just bought by more people) in a subsequent period than smaller brands.

Receiving a signal about the number of people engaged in an activity does not only influence consumption behaviour but also appears to influence contributions to public goods and compliance behaviour too. Essentially, behaviour that is observed can be described as "conditionally cooperative"; that is, people's contributions and compliance are conditional upon what they believe the levels of others' contributions and degree of compliance are.

This behaviour with respect to contributions has been observed in the field. Croson and Shang (2009) who test contributions to a public radio station in the USA. They find that a caller's contribution can be raised by reporting higher levels of contributions by other callers. Frey and Meier (2004a, 2004b) find a similar relationship with regard to the pattern of giving to two charities on the part of students at the University of Zurich. Similarly Heldt (2005) finds that cross country skiers in two ski resorts make contributions to a ski track in a fashion that is positively dependent upon the degree to which others contribute to the track.

Under the slightly more controlled conditions of lab experiments contributions seem to follow the same pattern. Using a linear public good game design Fischbacher *et al.* (2001) test how much each subject is prepared to give in a public goods game for varying indicated levels of contributions from other members of the group. Given the structure of the game and its payoffs, complete free riding was predicted. Fischbacher *et al.* find in their experiment that 50% of the subjects contribute in a conditionally cooperative fashion, while a third free ride. Croson (2007) conducted a series of experiments that aimed to test the comparative static predictions of linear public goods games based on assumptions about how individuals would contribute to the public good. The results supported a model of giving that suggested reciprocity sensitive behaviour on the part of subjects. Subjects were significantly and positively influenced by their beliefs about how much others were giving, in particular, subjects overall seemed to try to match the median contribution of their group.

Compliance behaviour may be influenced too. Fellner *et al.* (2009) find that belief amongst respondents about the general prevalence of paying TV licences is a strong, positive, predictor of whether a person will pay their TV licence. Spicer and Becker (1980) found that tax evasion was higher amongst people in higher tax brackets, who know that the average person is contributing less tax than them. Conversely they find tax evasion to be lower amongst people in lower tax brackets amongst those who know that others are contributing more than they are. Frey and Torgler (2007) estimate tax morale across a 30 country European survey[3]. They find very strong support in this data for tax behaviour and tax morale to be strongly determined by an individual's perception of the level of tax evasion in their country.

B. Social Distance

Social distance, or the degree to which a person or group feels that they are similar to a norm or a member of a larger group of which the norm is representative is another important way in which the influence of norms may be modified. Simply put: it may not be enough to simply report a norm to someone or some group, the reported norm probably needs to reflect behaviour of a person or group that is in some way close or similar in order for it to have an effect.

The idea is certainly not a new one. Early research in social psychology seemed to suggest that an individual was influenced by another or by a group because that individual wanted, in some way, to be part of and like by that one other or that group (Schultz *et al.*, 2008, Bandura, Ross & Ross, 1961,

3 The European Values Survey. European Values, 1999. Questionnaire. Tilburg University, The Netherlands

Deutsch and Gerard, 1955). Current research also suggests that social distance can play a role in mediating influence.

Within the education literature Sacerdote (2001) examines Grade Point Average (GPA) and fraternity data for freshmen students at Dartmouth and finds that roommates effect each other's GPA. In a Mexican study, Lalive and Catteneo (2009) find that subsidising the attendance of a portion of children at a school, thus ensuring their attendance, raises the attendance of those children who were *not* subsidised as well as those who were. In fact Lalive and Catteneo (2009) find this secondary, seemingly peer-influence effect is nearly as important as the direct funding effect.

Christakis and Fowler (2007) analyse 12 067 people in the Framingham Heart Study and find that the number of obese persons (with a body mass of 30 or greater) were to be found in discernible clusters. The chance of a person being obese in the data increased by a considerable 57% if they had a friend that became obese. This is similar to Leatherdale and Papadakis (2011) who obtain a similar result for obesity among adolescents as well as studies dealing with adolescent smoking (Cameron, 1999; Leatherdale *et al.*, 2006). The effect of those close to us thus appears to extend to our waistline and lungs too.

Within the neighbourhood effects literature there are uncovered role model effects, where people are shown to imitate others as well as peer influence effects, where people contemporaneously influence each other (Durlauf, 2004).

In terms of morality; research by Ariely *et al.* (2009) indicates that when people are confronted with a moral norm, such as for cheating, within their own group they are more likely to conform to that norm (be it honesty or cheating). However, when confronted with a moral norm in a group of which they are not a member, people are likely to be behave in the opposite fashion to the norm.

Although this body of research suggests that social distance plays a role in mediating the effect of a reported norm, it is unlikely that this happens in a fashion that is necessarily the same, or even straightforward, in all contexts. Experimental research in the laboratory by Croson, Handy & Shang (2006), at least, suggests that social distance is modified by the cultural context of the respondents.

It thus appears, from the literature that the reaction to norms may be significantly modified by the social distance between the person, or group, and the norm concerned. The implication for norms based messaging which aim to reduce consumption is that it would be better if the norms (and potentially other communication) was framed in socially close terms, however, the marginal effect of this is likely to differ across groups, at least with respect to culture but likely in interaction with other elements too.

C. Pluralistic ignorance and social norms

One of the most important ways in which reporting social norms is thought to work is by dispelling what is known as pluralistic ignorance. Pluralistic ignorance is said to be present when most or all of a group are ignorant about what others in the group do or think. The concept of pluralistic ignorance has particular bearing upon the investigation of normative social influence in regard to college drinking in the USA. The literature on normative influence and college drinking is predicated upon the thought that college students are, at least partly (but still), significantly motivated to drink too much because they think that overindulgence is the norm at college (Prentice& Miller, 1993; Perkins, 2002; Perkins *et al.,* 2005).

Several studies have demonstrated that college students routinely overestimate how much it is normal to consume by way of alcohol as a student (Berkowitz, 2005; Perkins, 2002 provide overviews of studies in this area) believing that college students consume, on average, more alcohol than they, in fact, do. Interventions in this area have thus been targeted at informing college students about what the normal level of alcohol consumption for their group really is and so reducing pluralistic ignorance on this issue.

The results of applying norms reporting to college drinking have, however, been mixed. With over-consumers reporting the norm appears to have had the effect of moderating consumption (Agostinelli, Brown, & Miller, 1995; Haines & Spear, 1996; Neighbors *et al.*, 2004) some of the time; for under-consumers particularly, reporting the normal level of consumption seems to have raised consumption toward the norm in a number of cases (Perkins *et al.*, 2005; Wechsler *et al.*,2003; Werch *et al.*, 2000).

D. Social norms and learning efficient behaviours

Another reason why reported norms may influence behaviour is that they could, for many, be particularly compelling ways of presenting important information about an optimal behaviour. This may be true especially in cases where it would be difficult and costly (in terms of time, money or some other resource) for people or households to obtain the information necessary to make a decision. In these cases norms may prove attractive as signals that a person/household can follow. They may serve as a signal about a state of information that is difficult to observe individually. Simply put, people or households may conform to norms when they do not know the best course of action and it is difficult to discover what that course is. The rationale would be that norms usually indicate the results of the costly trial-and-error discovery on the part of other people or households.

That people will use norms as a proxy for information itself is a result obtained experimentally (Wärneryd, 1994; Roth, 1985) and demonstrated theoretically (Potters *et al.,* 2005) in the literature. This is found to be true even to the extent that a stable norm may persist, with people conforming to it, even if it is not the most efficient behaviour. The evidence is both theoretical (Akerlof, 1980; Romer, 1984) and from the field (Young and Burke., 2001).

1.3 Information which maps behaviours to outcomes

The final behavioural element that is of relevance to this project that we wish to examine is that of information which maps behaviours to outcomes. By mapping we mean the relationship between an action and a consequence that is perceived by a person. Mappings between cause and effect are important because they are the basis upon which we act if we wish to achieve some particular outcome, such as reducing water consumption in our home. As such, establishing the correct mapping between action and consequence is exceptionally important if we wish people to adopt efficient levels of water consumption, consistently. Within the water conservation space it is a deficit of this mapping which information campaigns assume to be the root cause of a lack of conservation behaviour.

Several factors may hinder the establishment of a correct mapping between behaviour and consequence. Three factors that in particular militate against establishing a correct mapping are: complexity, lack of transparency and infrequency, or delay, of feedback (Thaler & Sunstein, 2009; Kahneman, 2003).

Complex situations often arise when there are several elements interacting with each other to produce an outcome. The more complex a situation, the more difficult it will be to separate out what is causing what and how. In situations of increasing complexity it is less and less likely that a person will perceive the correct mapping between behaviour and consequence. The water consumption of a household may occur in obvious and non-complex manners, such as pouring a bath. However, much of a household's water consumption will occur in a fashion that is highly mediated by appliances (such as washing machines) and hence be complex. [4]

Lack of transparency in a system will also mitigate against the establishing of a good mapping (Kahneman, 2003). By lack of transparency we mean that the manner in which an outcome is arrived at has an unduly low level of salience. There are two ways in which this will usually happen in the case of a household's water consumption, the first is with respect to partially observable consumption, such as leaks, showering and watering the garden. The second is the manner in which consumption is presented on the water bill.

Leaks are a significant source of water loss in a household, one of the reasons for this is that they are often not observable. Dramatic leaks are quickly noticeable. However many leaks, such as cistern leaks are slow and difficult to notice. Most leaks around the household are of this nature. As a result many households will not be aware of the fact that they are likely to have leaks that contribute to their consumption. Similarly activities such as watering the garden and showering are not amenable to accounting, at the time of the activity, of how much water is being used (unlike filling a bath or a sink). Households may well not know their flow-rate and hence will not know how many litres of water they use in such circumstance.

The water bill may not necessarily help to improve transparency either. A water bill that only presents a final billable figure in Rands but does not clearly show how this figure was determined lacks transparency. This does seem to be the case for the majority of water bills in South Africa; as stressed again and again by Slabbert (2010a, 2010b) in her analysis of water bills in South Africa.

Finally, the greater the delay between actions and true feedback, the less likely it will be that a person will perceive the correct mapping between their actions and the consequences of those actions (Thaler and Sunstein, 2009). Serially estimated billings punctuated by billings based on true meter readings are an example of a lack of feedback. In these cases a customer may perceive the estimated consumption as an invalid reflection of the reality and only regard the true readings. However, the interval of time between true readings almost certainly hides too much activity for specific behaviours to be remembered and linked to consumption.

If we hope to use feedback to help consumers establish efficient consumption behaviours that persist over the long term, it will almost certainly be necessary that we use this feedback to establish a clear mapping between consumption behaviour and consequences. The sort of feedback that will achieve this will need to simplify complex processes, bring transparency and be frequently given with as small a delay as possible.

4 The relationship between a household's water consumption and the well-being of a large system such as the municipality's natural environment is also an example of a complex system. It is likely that households will not be able to, on their own, establish a mapping between their own behaviours and consequences at these levels. This would potentially interfere with appeals to households to conserve water that employ a motive to conserve the natural environment.

1.4 Summary

The Behavioural Economics literature indicates that there are various elements which can influence our decision-making and behaviour. These elements do not necessarily always lead to an outcome that is in our own interests, nevertheless it seems that, more often than not, their influence can be decisive. The implication for policy is that policy should seek to incorporate these behavioural realities in order to work with people as they are, rather than as we think they should be. To do this the most straightforward recommendation would be to construct policy in terms of these behavioural elements.

This project will focus on three particular classes of behavioural elements: Salience, Social Norms and Information that enables better Mapping. These elements inform the design of the particular mailed interventions, or "treatments", that were used in this study. Via their expression in these mailed treatments, we will assess the ability of these various elements to influence a household's water consumption.

Chapter 2: Designing the experiment

This chapter describes the design of the experiment at the heart of this study. Before proceeding however; a restatement of the objective of this study is probably apposite. Given that Cape Town, along with many of South Africa's large cities, has meaningfully limited water resources for significant portions of the year it needs to find ways of managing the demand of that supply. One possible strategy may be to use feedback to households about their water consumption that employs the behavioural elements of either social norms, salience, tips about how to save water or some combination of them. Similar strategies seem to have met with success in the USA. This study aims to assess which of these behavioural elements (or combinations), if any, is most effective in eliciting a reduction in water consumption from households. We now describe the experiment.

This experiment was a randomised control trial (RCT) that had a total of nine treatments. Each of the treatments, apart from the control, utilises one or more of the behavioural elements described in the previous chapter. The sample is randomised at the neighbourhood level and is balanced in terms of the pre-treatment annual daily average water consumption of the households involved. In keeping with the emerging best practice around reporting RCTs we supply a detailed description of the sample selection process and method of randomisation as well as addressing ethical concerns associated with a study such as this. We supply evidence that the sample being mailed is not biased in any significant fashion.

2.1 Treatments used in the field

Our experimental design involved a total of nine different treatments. Eight of the treatments involve households receiving an insert, or combination of two inserts, with their water bill while the ninth is a control group of households who receive no inserts with their water bills. These inserts are presented in the language in which the household is billed. Examples of each type of insert are provided in the appendix at the end of this report.

Each of the inserts attempts to leverage a particular behavioural response to various elements that can be found in choice architectures, such as those that a bill may contribute to. These elements are: social norms and the increased salience of a household's own consumption, the salience of ways in which water (or electricity) could be saved and an increased salience of the fact that another is scrutinising the household's consumption.

Some of the treatments leverage one of these elements, we have termed these "single element treatments", while other treatments leverage several, we have termed these "combination element treatments". Some single-insert treatments are combination element treatments and all of the treatments that utilise two types of inserts are combination element treatments. Table 2.1 summarises the various treatments that are being administered in this project and sections 2.1 and 2.2 provide a slightly more detailed overview of each.

Table 2.1. Inserts mailed to selected households

	Treatment: name and description	Components of treatment	Non-pecuniary, behavioural elements present in each treatment
1	**Control** (No insert is included with water bill)	Water account	Status quo
2	**DHM** (Household's average daily consumption for last month)	DHM (expressed as a horizontal bar graph) + Water Account	Raised salience of own consumption
3	**DHMA** (Household's average daily consumption for last month compared to the median neighbour's average daily consumption for the past month)	DHMA (expressed as a horizontal bar graph) + Water account	Raised salience of own consumption+ comparative social norm
4	**DHMAT** (Household's average daily consumption for last month compared to the median neighbour's average daily consumption for the past month; with a separate page of water savings tips included)	DHMA (expressed as a horizontal bar graph) + DT + water account	Raised salience of own consumption comparative social norm + Tips (mapping inputs to outputs)
5	**DHMAE** (Household's average daily consumption for last month compared to the median efficient neighbour's average daily consumption for the past month)	DHMAE (expressed as a horizontal bar graph) + water account	Raised salience of own consumption + comparative (efficient) social norm
6	**DHY** (Household's average daily consumption for each of the last 12 months for which data is available)	DHY (expressed as a vertical bar graph) + water account	Raised salience of own consumption
7	**DHYT** (Household's average daily consumption for each of the last 12 months for which data is available; with a separate page of water savings tips included)	DHY (expressed as a vertical bar graph) + DT + water account	Raised salience of own consumption + Tips (mapping inputs to outputs)
8	**DHYA** (Household's average daily consumption for each of the last 12 months for which data is available compared to a single figure average daily household consumption for their neighbourhood)	DHYA (expressed as a vertical bar graph with a straight line drawn through it at the level of the annual household average for the neighbourhood) + water account	Raised salience of own consumption + comparative social norm
9	**DT** (A page of water savings tips)	DT + water account	Tips (mapping inputs to outputs)

2.1.1 Single element treatments

Mapping

The DT treatment offers advice to the household about ways in which they can reduce their water consumption. This treatment aims to raise the salience of the link between various actions and the desired outcome of lower water consumption. The savings tips are derived from those offered by the City of Cape Town in their Smart Living Handbook (2011) and on the Water Demand Management section of the City of Cape Town's website.[5]

Salience and scrutiny of household's water consumption

There are two treatments that aim to raise the salience of the household's own consumption, the DHY and DHM treatment. Both of these treatments may very well also raise the salience of the household's consumption being under scrutiny. It seems that it is impossible to do one without the other in the context of a mailed treatment. The DHY treatment delivers to the household a history of their average daily water consumption for each month in the last twelve months while the DHM reports only the daily average for last month. The history is represented as a histogram in the case of the DHY treatment and as a bar graph in the case of the DHM treatment.[6] The average daily water consumption figure for the most recent month is drawn from the water bill itself. The average daily water consumption figures for the eleven previous months that are reported in the DHY inserts are drawn from a database set up by ourselves based upon consumption data for each account obtained from the City of Cape Town.

2.1.2 Combination element treatments

There are several treatments that leverage several behavioural elements at once. Sometimes this was accomplished by combining the insert bearing tips about how to save water with another treatment. The primary manner in which elements were combined were in the comparative social norm treatments where a household's consumption was compared with some normal level of consumption. This section explains the three major ways in which behavioural elements were combined and through which treatments they were delivered.

Comparative Social Norms

The "DHMA" treatment is a treatment that leverages social norms in a comparative fashion. It does this by comparing the household's average daily water consumption for the last month to the average household in their neighbourhood's consumption for the same past month. The comparison is made by way of a labelled bar graph and a short sentence which states whether the household consumed the same, more or less than their average neighbour.

---

5 http://www.capetown.gov.za/en/Water/Pages/WaterDemandManagement.aspx

6 Readings for some months may be estimated while others are true reading the graph reports. We do not differentiate between these two sorts of possible reading types on the graph of the DHY, DHM, DHA, for two reasons. Firstly, this information on the water bill itself. Secondly, it would frustrate our efforts to identify the influence of raising the salience of consumption if we also raised the salience of the estimation status of a reading.

To obtain the figures for the estimated average household in the neighbourhood we took the median of the constrained sample for the neighbourhood for the relevant month in the previous year. We used the Municipal Meter Reading Units (MRUs) to define the extent of the various neighbourhoods in our sample. A MRU is an area that a group of meter readers can cover in a day.

The "DHYA" treatment also leverages social norms in a comparative fashion. In this case the household receives a histogram of their average daily consumption for each of the last twelve months in exactly the same fashion as the DHY treatment. Included on the same graph is a flat line which indicates the annual daily average water consumption for an average household in the neighbourhood. This average is the median consumption over the period 01 October 2011 to 30 September 2012 (although the estimation period is not communicated to the household receiving this insert).

Comparative Social Norms of efficiency

The DHMAE treatment is a treatment that also leverages social norms in a comparative fashion. Unlike the other treatments that involve comparative social norms (DHMA, DHMAT and DHYA) it does not do this by comparing the average daily consumption to the median average daily consumption for their neighbourhood, but rather to an estimated efficient household in their neighbourhood. We discuss the calculation of the "estimated efficient household" below.

Deriving a figure to use as the average daily consumption for the estimated average efficient household posed challenges. Since we could not observe individual household characteristics and consumption behaviour, it is very difficult to know, on a case by case basis, to what degree each house was consuming water efficiently or not. As a result our approach was to reference the average of efficient households. This changed the problem to finding an estimate of the degree to which the average household was efficient. This problem is still a complicated one. Our approach was to find a reasonable estimate of reductions in consumption that households would normally make in response to behavioural interventions (as opposed to interventions which leverage changing to more efficient in-home capital) and then reduce the median for the neighbourhood by that percentage.

The literature reviewing results of various behavioural interventions in water consumption presents a wide range of values. We decided to employ a figure from the conservative end of the spectrum. Markowitz and Doppelt (2009) review various interventions and present a range of results between 6.7% and 25%. We decided to use the conservative 6.7% figure. This figure was from a study by Staats *et al.* (2004) and, from that study, seemed to be a rate of saving that was sustained over the long term.

Mapping plus Comparative Social Norms

There are two composite treatments involving the informational treatment, DT. In both instances we are hoping to explore the link between two possible behavioural motivations to save water and knowledge about how to save water. For both inserts, the DT treatment is printed on the reverse side of the page.

In one case DHMAT, the DHMA treatment comparing the household's water consumption for the last month with the average household in the neighbourhood is augmented with the informational treatment. In this case we are looking at the link between comparative social norms and potentially improved information about how to save water. In the DHYT treatment the DHY insert is printed on one side of the page with the DT treatment printed on the reverse. In this treatment we are exploring the link between raised salience of a person's own consumption and potentially improved information about how to save water.

2.2 Final Sample

The sample for this study involves 398 423 households in 2443 neighbourhoods[7]. These households were randomly allocated to the 9 treatments as described in the table below. The control group was randomly allocated just over 20% of the sample, with the remaining treatments receiving roughly 10% of the sample each. These proportions were followed in order to stay within our budget which allowed us to mail 315 418 households, or 80% of our sample, only. We now move to describing how this sample was drawn.

2.2.1 Method of sample selection, randomisation and balancing

In this section we describe our sample selection and randomisation methodology. Using the methods this sections describes we achieved a sample where treatments were randomised by neighbourhood. Randomisation was done in a re-randomising fashion, common to the Randomised Control Trial literature (Bruhn and Mckenzie, 2009) such that balance was achieved in terms of pre-treatment annual daily average water consumption. Balance in this regard means that there is no bias between treatment groups in terms of pre-treatment annual daily average water consumption. In the spirit of Bruhn and Mckenzie (2009) we describe our method of random allocation explicitly. We also discuss the ethical considerations attending to this experiment's design in the spirit of Barrett and Carter (2010).

The water consumption history of all households that received a water bill from the City of Cape Town municipality, between 01 January 2011 and 30 September 2012, was obtained from the City of Cape Town's water billing database. Based upon each household's average daily consumption for the year period 01 October 2011 to 30 September 2012, households that consumed more than 2 000 litres per day, on average, were excluded as outliers. Households that consumed, on average, the free basic water (FBW) allocation of roughly 202 kilolitres per day or less were excluded. These households were excluded on the basis of ethical considerations (discussed in its own section below).

The households from which the sample is selected was further restricted to households that have not moved in the last year and which have had their meters actually read or estimated by the City's automatic estimation algorithm. Households of which this is not true are difficult to measure, with reasonable accuracy in terms of consumption. We include automatically estimated consumption in order to achieve a sufficiently large sample size and also because the estimation algorithm is mostly quite accurate. We exclude manual estimations, since they are not as accurate as the automatic algorithmic ones. We exclude re-estimations due to over-estimation as well as interpolation since it is impossible to tell which month these re-estimations apply. Re-estimations may apply to one or several months. In the case of interpolation the range of dates which the interpolation is based upon are not specified, so it is also difficult to know which consumption figure/s to adjust.

Amongst the remaining group of 398 423 households we randomly allocate the 9 treatments (including the control group) to households as follows. In order to remain within budget constraints

7 Neighbourhoods are meter reading units. Meter Reading Units are the areas covered by a group of meter readers in a day.

we randomly allocated 20% of households to the control group with the remaining 80% allocated to each of the remaining, print-based, treatments at 10% each. To do this we sorted accounts by ascending consumption within each neighbourhood. We then ascribed a randomly generated number to each account. These random numbers were generated by a uniform distribution function that was itself generated using as a seed a random number generated from MS. Excel's random number function. Once each account had a random number, these accounts were sorted in ascending order of their random values within each neighbourhood. This would effectively order households randomly with each neighbourhood with respect to observable characteristics, such as consumption, and unobservable characteristics such as household size. Households were then given an index number according to the rank of their random number's value for their neighbourhood.[8] Households were then ascribed a treatment based on which percentile of the index, for the neighbourhood, they occupied according to the table below. Once treatments had been ascribed to accounts for each suburb according to this method we then checked that there was balance between treatments in terms of annual daily average consumption. Balance between treatments in terms of consumption was assessed at the overall level as well as in terms of three pre-treatment consumption bands as defined below.

Low : 202.24 litres ≤ annual daily average consumption ≤ 500 litres

Medium: 500 litres ≤ annual daily average consumption ≤ 1 000 litres

High: 1 000 litres ≤ annual daily average consumption ≤ 2 000 litres

Balance in terms of consumption was deemed to occur if there was no significant difference at the 5% level between treatments in terms of their pre-treatment annual daily average consumption at the overall level and at any consumption band level (low, medium or high). If balance was found to have not been achieved then the process was repeated by obtaining a new seed number from MS Excel's random number function. The process was terminated as soon as balance was achieved. The final sample is thus one in which there is no difference between treatments in terms of pre-treatment annual daily average consumption either at the overall level or within the three specified consumption bands. The end result of the randomisation process is a sample that is distributed as follows

8 Thus the household with the random number that had the lowest value of all random numbers in the household would be ranked "1", the next highest valued random number for the neighbourhood would be ranked "2" and so on.

Table 2.2. Number of Accounts allocated to each treatment

		Pre-treatment annual daily average water consumption bands. Expressed in kilolitres (kl)		
	Total Sample	High Consumption Sub-sample (1-2 kl)	Medium Consumption Sub-sample (0.5-1 kl)	Low Consumption Sub-sample (0.2-0.5 kl)
Treatments				
DHM	39 286	6 423	16 170	16 693
DHMA	38 638	6 318	16 032	16 288
DHMAE	38 638	6 345	15 890	16 403
DHMAT	39 286	6 517	16 019	16 750
DHY	39 337	6 465	16 306	16 566
DHYA	39 285	6 489	16 052	16 744
DHYT	37 966	6 191	15 570	16 205
DT	39 337	6 465	16 200	16 672
Control	86 650	14 453	35 507	36 690
Total	398 423	65 666	163 746	169 011

2.2.2 Ethical considerations

Since human subjects were involved the ethical considerations of this experiment's methodology were important. As a result we will spend some time discussing them here.

The ethics of administering treatments designed to induce lower water consumption to households that are consuming low levels of water are complicated. On the one hand, since the Free Basic Water (FBW) allocation is calculated from WHO guidelines for minimum healthy water consumption, it would be wrong to try to get people consuming this amount of water to consume still less. However, since the amount of FBW is calculated at the household level and not at the individual level it is not necessarily true that all households consuming less than the FBW allocation will have individuals consuming less than what is healthy. Conversely, households consuming FBW only may not necessarily have adequate provision for all individuals in them. It is thus true that for households with individuals consuming in excess of what is needed for a healthy life consuming less would be beneficial for them in terms of paying less as well as for the community of which they are apart, by freeing up more water (and arguably keeping the cost down by lowering the demand for it).

Modifying the above concerns is the expectation on our part about the nature of the causal mechanism underlying any effect of the various treatments. Essentially, comparison to norms, raising the salience of consumption or providing information about how to save water targets the reflective and non-reflective decision-making process (Kahneman, 2003) of household members, via the bill (and/or whoever reads the bill). These influences tend to be small and at the margin. These influences upon the decision-making process will work in a context where other influences are at work too, such as the need for water on the part of household members. In situations where the household is already consuming as little as possible and members are at their daily minimum, we think that more primal concerns will dominate and we do not think it likely that such households will be swayed very much by our treatments. We expect that the behavioural elements embodied in our treatments will be more effective amongst households where there is consumption of water in excess of what is required for health and the decision process is not dominated by more primal, survival-like concerns. In short we expect that our treatments are likely to only be effective in situations where this IS water that can be saved.

Since we cannot observe household size and particular household circumstances it is difficult to know with certainty how the above considerations would earth in each particular household account number and consumption billing history (all that was visible to the researchers when preparing the sample). Considering the ethical concerns as a whole in a context where we cannot know the specific household characteristics and given the likely nature of the causal mechanism, we thought that it would be best the exclude households consuming the FBW amount or less. This seemed to balance the concern about targeting households with minimal healthy consumption with targeting households where we were most likely to be reaching situations where there was water to be saved.

2.2.3 Power calculations

Power calculations were made in order to determine the level of statistical power that would be achieved at the conventional 5% confidence levels when comparing two treatments groups. Since all of the results would essentially be comparisons made between two groups (either different active treatments or an active treatment to control)[9] this calculation was based on the formula for the comparison of two means set out below.

$$n_1 = \frac{\left(\sigma_1 + \frac{\sigma_2}{r}\right)\left(z_{1-\alpha/_2} + z_{1-\beta}\right)^2}{(\mu_1 - \mu_2)^2}$$

which we can rearrange to give us a calculation in terms of power in the following way

$$\sqrt{\frac{n_1(\mu_1 - \mu_2)}{(\sigma_1 + \frac{\sigma_2}{r})}} - z_{1-\alpha/_2} = z_{1-\beta}$$

Where n_2 is one treatment group's sample size, call it "group 2" while n_1 is the other treatment group's sample size, call it "group 1". Similarly σ_2 denotes the estimated likely variance for group 2 and σ_1 for group 1, while μ_2 denotes the mean proportional consumption for group 2 and μ_1 for group 1. The term "r " refers to the ratio of one group's sample size to the other and is calculated as simply $\frac{n_2}{n_1}$. Finally, the term $z_{1-\alpha/_2}$ is the confidence level for a conventional confidence level "α" of 5% when large numbers are involved (1.96) and $z_{1-\beta}$ is the critical statistic for the level of power desired, here 80%, when large numbers are involved (1.29). This is consistent with the conventional level of statistical power for these exercises in the literature (Ferraro and Price, 2011). Estimates for the likely mean and variance in consumption for a treatment group in the field period of December 2012 were drawn from consumption figures for December 2011. Power calculations are conducted for several different effect sizes.

Since we are unsure how the various treatments may differ in their effect sizes (there is almost no literature comparing treatments of the sort we are using here to each other) we report the

9. Although households were to be measured each month, since they were to be mailed each month also, this rendered each successive after a mailing month a different treatment to the one before it since we would be observing the effect, potentially, of two different amounts of mailings received.

minimum percentage difference between two groups that can be detected at the 5% confidence level with a power level of 80%. We calculate this by rearranging the equation above to obtain.

$$\sqrt{\frac{\left(\frac{\sigma_2^2 + \sigma_1^2}{\frac{n_2}{n_1}}\right)\left(z_{1-\alpha/2} + z_{1-\beta}\right)^2}{n_1}} = \mu_1 - \mu_2$$

The tables below report the minimum detectable effect $\mu_1 - \mu_2$ as a percentage of μ_1.

Table 2.3. Minimum detectable percentage difference between the Control group and other treatment groups at the 5% confidence level with 80% power

	Overall sample	High consumption band	Medium consumption band	Low consumption band
Control	1.37%	2.60%	1.32%	1.35%

Table 2.4. Minimum detectable percentage difference between treatment groups at the 5% confidence level with 80% power

	Overall sample	High consumption band	Medium consumption band	Low consumption band
DHM	2.04%	3.90%	1.96%	2.00%
DHMA	2.05%	3.93%	1.97%	2.03%
DHMAE	2.05%	3.92%	1.98%	2.02%
DHMAT	2.04%	3.87%	1.97%	2.00%
DHY	2.04%	3.88%	1.95%	2.01%
DHYA	2.04%	3.88%	1.97%	2.00%
DHYT	2.07%	3.97%	2.00%	2.03%
DT	2.00%	3.80%	1.92%	1.97%
mean	628.27 litres	1296.14 litres	669.85 litres	331.95 litres
std. Dev	552.12 litres	880.70 litres	363.64 litres	186.73 litres

2.3 Mailing methodology

The inserts were delivered to households as A4 pages that were included along with their water bills. The inclusion of inserts with their accounts happened automatically at the printers. We describe this process briefly.

Inserts are generated based upon a database sent to the printing company by ourselves. This database indicates, for each account, the type of treatment that has been allocated to that account as well as any of historical data or neighbourhood average consumption figures required to generate the specific insert that are discussed above. This database then drives the generation of inserts at the printers.

When the accounts are printed the account printing program references the City of Cape Town's bill-printing file for the day as well as the database for this experiment. This allows inserts to be generated concurrently with their accounts as part of the account print run process rather than being added in later. Inserts are thus generated, from a print run point of view, as part of the bill. Once the print run is complete a spreadsheet from this run is emailed to the researchers. On the spreadsheet are all of the accounts in our sample that were printed that day as well as relevant information about the account from the print run, such as period for which the account is being billed and the average daily consumption over that period.

In addition to the concurrent generation there is a back-up process to verify that the correct insert has been included with the correct account. Each account's water bill has a unique optical binary code to identify it. Each of the inserts used in this project has a similar optical binary code that automatically uniquely associates it with a specific account. During the automatic folding and envelope insertion process, the envelope inserter reads the optical binary codes on each page being inserted into a particular envelope to confirm that the correct pages are being included together in the same envelope. If the pairing is incorrect then the process halts until the pairing is restored.

Chapter 3: Fieldwork: Implementation of the Randomised Control Trial

This chapter provides a brief description of the course of fieldwork for the project. As described earlier, this project is a randomised controlled trial that aimed to test the relative effectiveness of 8 different strategies (outlined in table 3.1 below) to induce households to consume less water than they otherwise would. All of the strategies involve mailing households inserts, or combinations of inserts, with their water bills. Households would be mailed over the course of one month. This document describes that mailing process and how it proceeded. In doing so this brief chapter also describes the necessary administrative infrastructure that would be required to run such an experiment or implement one of the strategies tested.

3.1 Mailing

After a set up period of programming and testing this project entered its fieldwork phase on 30 October 2012 when the printing company (Mailtronic) began printing inserts to match to accounts. These first accounts, with inserts matched to them, were posted out on 01 November 2012.

The City of Cape Town's billing system is divided into 20 portions of printing per month. Since printing for this project began at the end of October, the first accounts to be mailed our inserts were accounts that were found in portion 20 of the City's printing for the month. By the 28th of November the final portion, portion 19, had been printed and was due to be mailed by 30 November.

This project is an exercise that has been conducted at great scale. By the time the final portion (portion 19) had been printed and mailed 286 598 unique accounts will have been mailed inserts. There were only 2 697 cases of duplicate mailings in the entire print run. A duplicate mailing is possible if a household is billed more than once by the City of Cape Town during the month of this print run. In terms of the amount of accounts printed or mailed on any day, the smallest batch was 12,353 accounts (portion 11) while the largest was the first batch mailed (portion 20) with 20,313 accounts. Table 4.2 and table 4.3 provides a portion by portion breakdown of how many accounts were mailed in each treatment for each portion.

This is slightly less than the total selected population described in chapter 2. Attrition from the selected sample is due to mainly the fact that in any month a percentage of municipal bills will be delayed in terms of their mailing. Delays may occur for a number of reasons, but most can be attributed to a billing that is somehow atypical and which the City needs to resolve or confirm to be accurate. These atypical billings are flagged by an algorithm in the City's billing system and are then referred to human managers for checking, which may generate a delay. A small minority of households will receive billings in periods less frequent than monthly (e.g. quarterly or bi-annually), some such households may have fallen outside of this project's mailing window period. As described in the next chapter this attrition appears to have happened in a random fashion across all treatment groups (including the control group) with the result that there was no bias between the various treatment groups in terms of pre-treatment consumption.

Table 3.1 Number of Accounts mailed (duplicate mailings included)

		Treatments								
Meter Reading Portion	**Date Posted (as distinct from date printed)**	DHM	DHMA	DHMAE	DHMAT	DHY	DHYA	DHYT	DT	**Total**
20	01/11/2012	2,539	2,565	2,555	2,471	2,585	2,599	2,498	2,501	**20,313**
1	05/11/2012	2,140	2,084	2,080	2,185	2,125	2,160	2,095	2,091	**16,960**
2	06/11/2012	1,781	1,764	1,741	1,754	1,752	1,797	1,643	1,798	**14,030**
3	08/11/2012	1,803	1,749	1,788	1,812	1,777	1,819	1,755	1,777	**14,280**
4	09/11/2012	1,672	1,630	1,635	1,645	1,700	1,686	1,600	1,700	**13,268**
5	12/11/2012	2,084	2,002	2,001	2,063	2,018	2,018	1,978	2,021	**16,185**
6	13/11/2012	1,708	1,641	1,708	1,751	1,746	1,744	1,670	1,703	**13,671**
7	14/11/2012	1,738	1,717	1,661	1,736	1,724	1,688	1,666	1,690	**13,620**
8	15/11/2012	2,033	2,017	2,066	2,041	2,046	2,080	1,932	2,074	**16,289**
9	16/11/2012	1,564	1,498	1,570	1,592	1,544	1,588	1,509	1,554	**12,419**
10	19/11/2012	1,630	1,605	1,545	1,582	1,612	1,583	1,584	1,586	**12,727**
11	19/11/2012	1,558	1,494	1,506	1,605	1,575	1,571	1,507	1,537	**12,353**
12	21/11/2012	1,637	1,588	1,569	1,631	1,619	1,578	1,615	1,625	**12,862**
13	21/11/2012	1,803	1,762	1,798	1,802	1,832	1,808	1,773	1,842	**14,420**
14	22/11/2012	1,700	1,712	1,660	1,716	1,714	1,733	1,674	1,728	**13,637**
15	23/11/2012	1,721	1,716	1,683	1,687	1,726	1,717	1,663	1,749	**13,662**
16	26/11/2012	1,736	1,693	1,651	1,710	1,698	1,660	1,646	1,696	**13,490**
17	To be posted	1,862	1,845	1,901	1,896	1,918	1,891	1,863	1,921	**15,097**
18	To be posted	2,007	1,974	1,945	1,974	1,977	2,006	1,916	1,960	**15,759**
19	To be posted	1,788	1,760	1,772	1,750	1,759	1,824	1,752	1,848	**14,253**
Total		**36,504**	**35,816**	**35,835**	**36,403**	**36,447**	**36,550**	**35,339**	**36,401**	**289,295**

Table 3.2. Number of Account mailed (duplicate mailings excluded)

		Treatments								
Meter Reading Portion	**Date Posted (as distinct from date printed)**	DHM	DHMA	DHMAE	DHMAT	DHY	DHYA	DHYT	DT	**Total**
20	01/11/2012	2,539	2,565	2,555	2,471	2,585	2,599	2,498	2,501	**20,313**
1	05/11/2012	2,139	2,083	2,079	2,183	2,123	2,158	2,088	2,088	**16,941**
2	06/11/2012	1,777	1,761	1,735	1,752	1,751	1,791	1,640	1,795	**14,002**
3	08/11/2012	1,794	1,741	1,784	1,805	1,773	1,813	1,751	1,771	**14,232**
4	09/11/2012	1,669	1,628	1,633	1,637	1,699	1,683	1,597	1,696	**13,242**
5	12/11/2012	2,078	1,996	2,000	2,058	2,007	2,010	1,974	2,016	**16,139**
6	13/11/2012	1,696	1,636	1,701	1,745	1,736	1,738	1,666	1,697	**13,615**
7	14/11/2012	1,733	1,711	1,652	1,729	1,718	1,686	1,664	1,689	**13,582**
8	15/11/2012	2,022	2,006	2,052	2,028	2,037	2,073	1,924	2,065	**16,207**
9	16/11/2012	1,555	1,490	1,559	1,583	1,534	1,583	1,500	1,539	**12,343**
10	19/11/2012	1,615	1,579	1,534	1,563	1,596	1,565	1,570	1,570	**12,592**
11	19/11/2012	1,540	1,479	1,493	1,588	1,562	1,554	1,486	1,520	**12,222**
12	21/11/2012	1,618	1,576	1,556	1,618	1,599	1,563	1,599	1,599	**12,728**
13	21/11/2012	1,783	1,747	1,777	1,785	1,815	1,788	1,750	1,816	**14,261**
14	22/11/2012	1,672	1,688	1,639	1,699	1,694	1,702	1,644	1,701	**13,439**
15	23/11/2012	1,696	1,683	1,646	1,661	1,704	1,699	1,636	1,731	**13,456**
16	26/11/2012	1,711	1,671	1,623	1,691	1,682	1,634	1,619	1,670	**13,301**
17	To be posted	1,835	1,815	1,870	1,874	1,898	1,872	1,833	1,895	**14,892**
18	To be posted	1,977	1,924	1,909	1,932	1,930	1,960	1,876	1,923	**15,431**
19	To be posted	1,706	1,682	1,694	1,691	1,682	1,750	1,680	1,775	**13,660**
Total		**36,155**	**35,461**	**35,491**	**36,093**	**36,125**	**36,221**	**34,995**	**36,057**	**286,598**

Chapter 4: Analysis and Results

In this chapter we present an analysis of the various treatments in this study. The outcome variable of interest was the average daily water consumption of households in the nine different treatment cells. In particular we are interested in assessing whether there was any difference between households receiving the various mailed treatments and households in the control group, who received no treatment.

Since this study was constructed as a randomised control trial (RCT) there was very little chance of bias being present in the various treatment cells. As described in earlier reports, there was no bias in terms of historical daily average water consumption between households in the various treatment cells.

Treatment and post-treatment period

Our treatments were posted to households over the month of November. Allowing for an average of five days postal time, they would have arrived at households sometime between the second week of November and the second week of December. Households would have thus been able to react to the treatment from those times. In terms of the conventions of dating water data collection, this means that our first post-treatment month is December. Our post treatment period thus becomes December, January and February (data was only available up to the end of February at the time of writing this report).

4.1 Data

The data used in this project to select the sample and then to measure the impact of the various treatments is the account level consumption data captured by the City of Cape Town's billing database. On its billing database the City records, for each reading on each account: the reason that reading was made, the type of reading it was and the status of that reading in terms of being billed. These three elements have a significant bearing on calculating the treatment effect and so we will discuss each in turn.

4.1.1 Reason for reading

There are several reasons why a reading may be made. The most common reason for a meter reading to enter the City of Cape Town's billing database is for the purpose of a periodic billing. This is the reading which is done in order to allow the City to assess how much water the account has consumed since their last bill and therefore how much to bill that account. However, this is not the only reason that a meter reading may enter the billing database; there are several more.

Although the billing database's code frame makes provision for 31 different reasons for a meter to be read readings only enter the database for five reasons beyond that of a periodic meter reading in the treatment and post treatment period. The frequency with which meters were read for these reasons over the treatment and post-treatment period are set out in table 1 below. Of the five non-periodic billing reasons, two relate to readings taken when a meter is disconnected (such as it might be for non-payment for the account) and when it is connected again. Two further reasons why meters were read in our sample over the period of interest relate to properties changing ownership and moving out of the property. When this occurs a final reading should be taken in order to conclude the former owner's account and a new reading should then be taken in order to initiate the new owner's account, where a new owner takes occupancy. The frequency with which these readings occur seem unreasonably low for a city as large as Cape Town, suggesting perhaps that these readings are frequently not made when ownership of a property changes. The final additional reason that a meter may have been read is for the purpose of an interim reading being done in order to establish some aspect of billing.

Table 4.1. Reasons for meter being read: Number of meter readings in each category for each month

Reason for meter being read	Treatment period	Post treatment period			Total
	November 2012	December 2012	January 2013	February 2013	
Periodic Meter Reading	356 647	347 750	408 178	344 080	1 456 655
Interim meter reading with billing	97	53	290	72	512
Final meter reading at move-in/out	768	75	207	75	1 125
Meter reading at move-in	1	0	2	0	3
Meter reading on disconnection	2 492	507	1 692	1 995	6 686
Reconnection meter reading	1 253	526	792	721	3 292
Total Accounts	361 258	348 911	411 161	346 943	1 468 273

In the analysis of the treatments we exclude all readings that are not periodic readings done for the sake of billing. This is done because for the non-periodic reasons it is not clear that a link between the treatment and consumption can be made. This is fairly obvious when we are dealing with a change of ownership somewhere in the period of interest. The household characteristics change and those in the household will not necessarily consider a comparison with previous tenants to be valid. Disconnection provides a structural break in consumption. Finally interim readings are redundant since they are folded into the figure billed in the periodic meter reading, including them would be to double count that consumption.

4.1.2 The way in which a meter reading is made

Consumption is measured, for each account, by means of the water meter for that account. In the billing database there are two main types of meter readings, each identifiable by a code. The majority of meter readings are readings that are classified as having been actually made by a meter reader looking at a physical meter and recording the value.[10] These are often referred to as "actual meter readings" by the people we worked with in the municipality although they are referred to in the billing database itself, and in table 2 below, as a meter reading "done by the utility".

The other type of meter readings are estimated meter readings and readings made by the customer. Estimated meter readings arise when a meter reader cannot observe the meter and a reading for that meter must be estimated. The City of Cape Town has an algorithm which achieves

10. The meter reading value along with the account number associated with it are entered onto a handheld device that the meter reader carries with them. At the end of the meter reading round the meter reader will upload the data from the handheld device to the City's billing database.

the bulk of this estimation. A meaningful portion of readings in each month, between 12.4% and 16.2%, are estimates. Meter readings done by customers very seldom enter the billings database.

Table 4.2. The way in which a meter reading was made: Number of meter readings in each category for each month

Reason for meter being read	Treatment period	Post treatment period			Total
	November 2012	December 2012	January 2013	February 2013	
Meter reading by utility	313 809	292 172	357 005	303 598	1 266 584
Meter reading by customer	150	100	142	198	590
Automatic estimation	47 299	56 639	54 014	43 147	201 099
Total Accounts	361 258	348 911	411 161	346 943	1 468 273

In the analysis we generally report estimates that include all the various ways of reading a meter as well as estimates that include only meter readings that were done by the utility, that is meter readings where it is claimed that a meter reader has actually read the meter. Given that the estimated readings clearly would not pick up on any behavioural change or adjustment as a result of the treatment, we do not include such readings in the analysis.

4.1.3 Status of the meter reading in the billing database

Finally, the figures recorded for each account's consumption are not all necessarily of billable status, although – as can be seen from table 3 below – the vast majority are. In some cases queries may arise as to the plausibility of the reading; in these cases the system will lock the reading and a manual investigation by an agent is required to ascertain whether the reading is legitimate or not. Once an agent has determined that a reading is legitimate, or replaced it with a correct reading, then the reading is released onto the system for billing. We include only readings that have been released onto the system or which were considered billable at the time.

Table 4.3. Status of the meter reading in the billing database: Number of meter readings in each category for each month

Status of the meter reading in the billing database	Treatment period	Post treatment period			Total
	November 2012	December 2012	January 2013	February 2013	
Billable	334 246	322 338	380 398	322 466	1 359 448
Locked by Agent	1	8	5	3	17
Released by Agent	27 010	26 564	30 757	24 473	108 804
Unspecified	1	1	1	1	4
Total	361 258	348 911	411 161	346 943	1 468 273

To summarize, our analysis includes only data that were from periodic billings, were actually read and were considered billable at the time entering the city's billing database.

4.1.4 Consumption over time

Cape Town's water consumption varies over the year, being high in summer and low in winter. The highest consumption months are January and February. These months are thus an ideal time to target from the point of view of mitigating water demand.

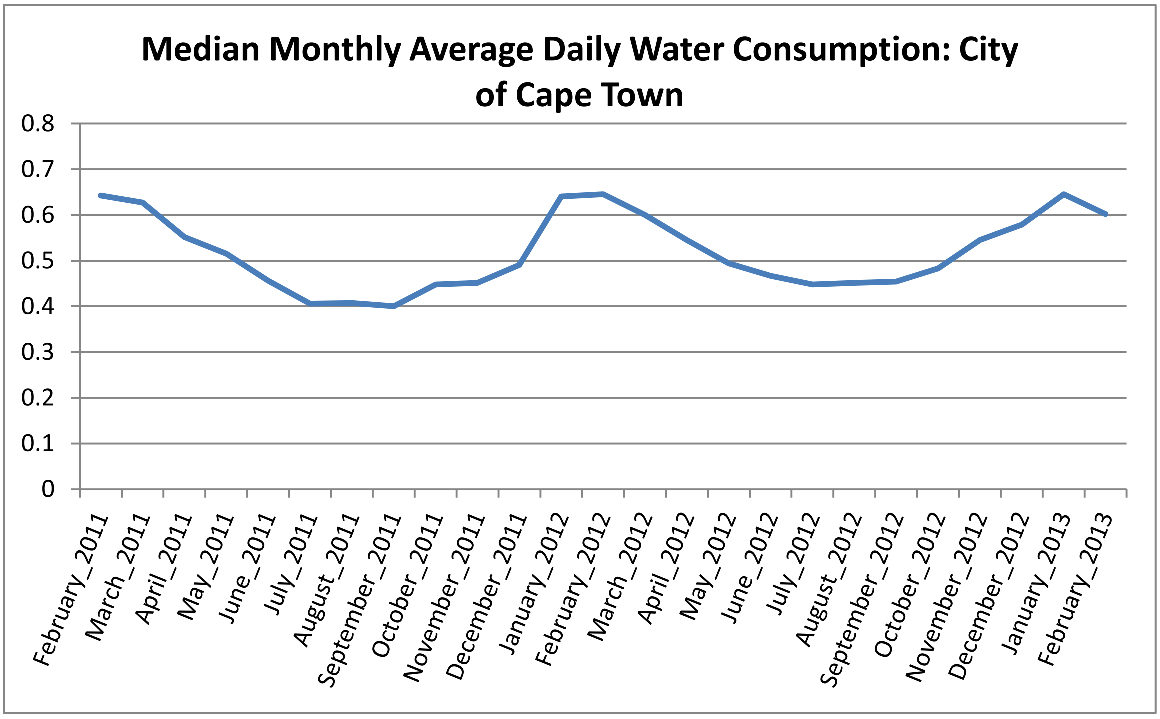

Figure 4.1. Median consumption pattern over the year for households in the City of Cape Town

Within each month there is a lot of variation in consumption across households and it is important to control for outliers, particular in the upper end of the consumption spectrum. The Kernel density plots overleaf demonstrate graphically the effect upon the consumption distribution of including and excluding consumption outliers. For estimation purposes we consider accounts whose consumption is below the 4th percentile of the consumption distribution or above the 96th percentile in a particular month to be outliers for that month. We exclude outliers in all of our estimates. Since we have percentile cut-off points to identify outliers, the exact cut-off values of average consumption will vary month to month, as describe in the table below.

Table 4.4. Outlier values for monthly average daily consumption

Outlier values	November 2012	December 2012	January 2013	February 2013
Lower bound (4th percentile)	0.15 kl	0.16 kl	0.16 kl	0.16 kl
Upper bound (96th percentile)	1.60 kl	1.90 kl	2.10 kl	2.10 kl

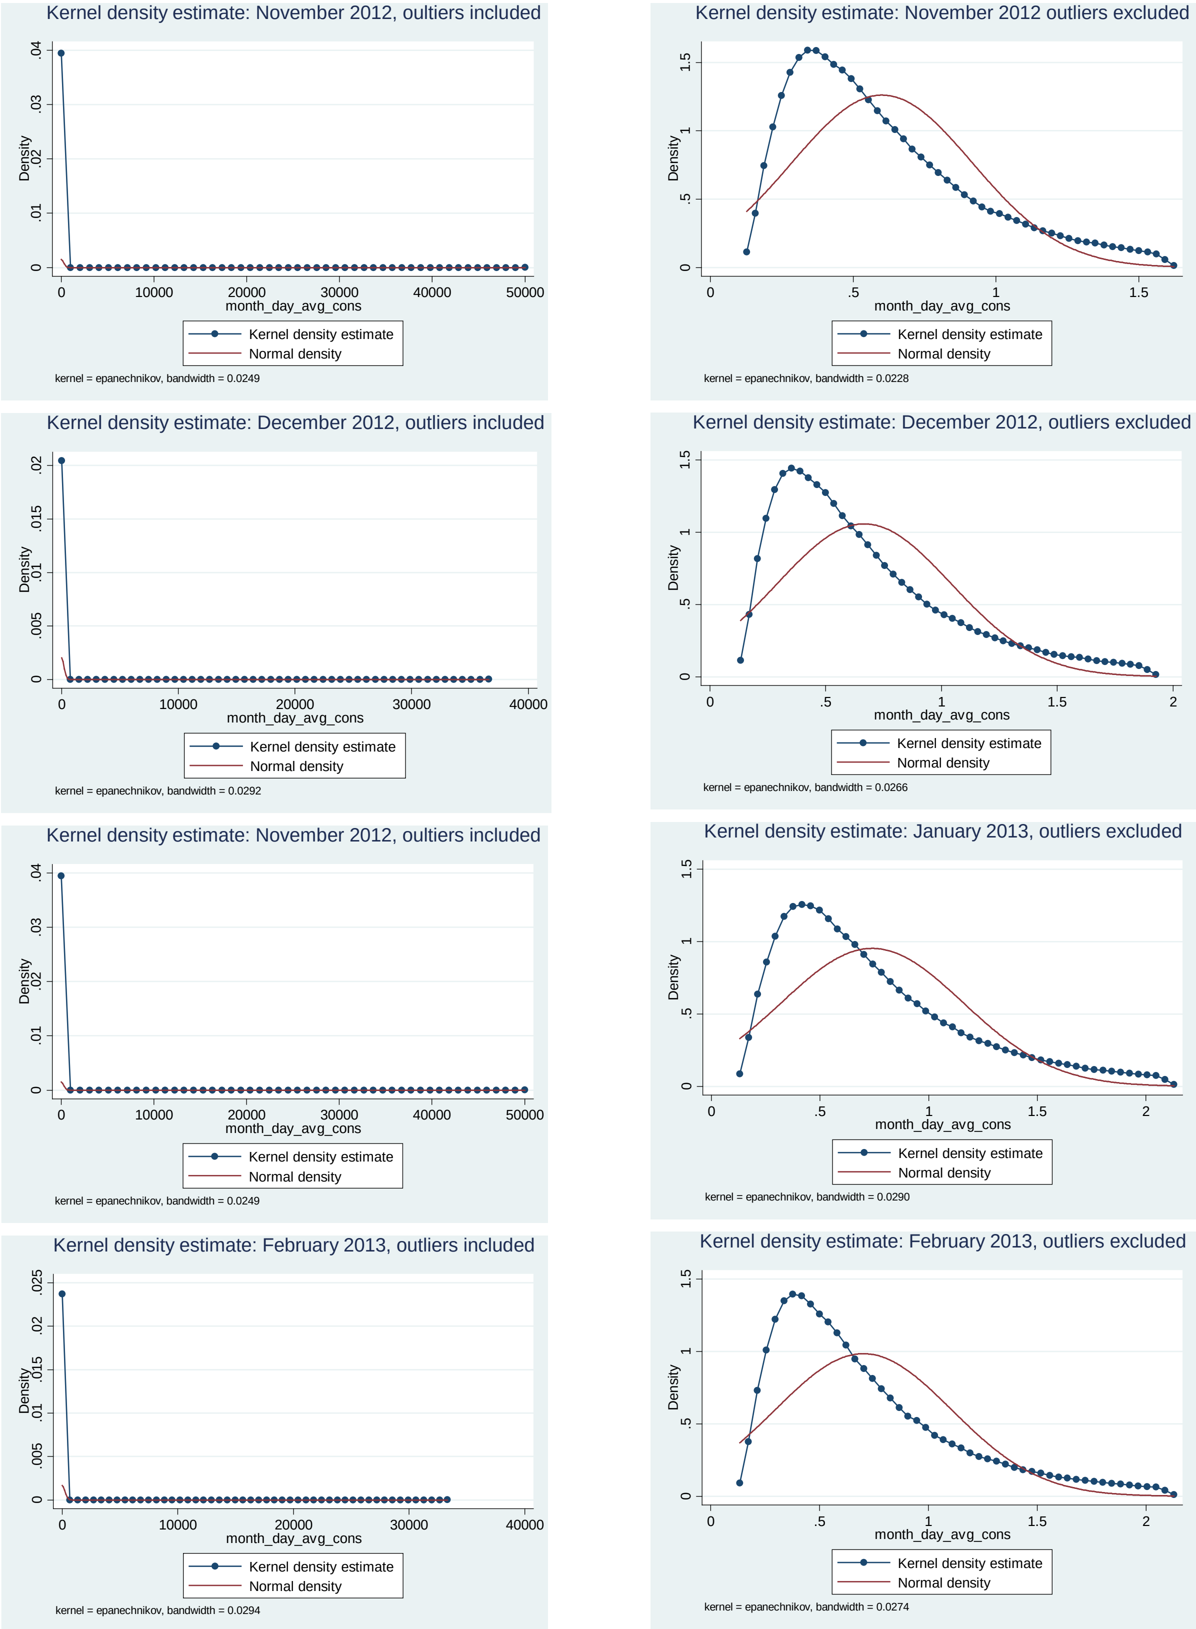

Figure 4.2. Median consumption pattern over the year for households in the City of Cape Town

Examining the distribution of household consumption with outliers removed reveals a distribution that is much more reasonable, although still not normally distributed. The bulk of households cluster at the lower end of the consumption spectrum, at around 0.5 kl or per day or below during the post-treatment months. This is to be expected however, since it corresponds broadly to the income and wealth distribution in the city of Cape Town, which is itself skewed in a similar fashion towards the lower end of the consumption spectrum.

4.2 Treatment effect within the total municipal sample: preliminaries

In order to assess whether our treatments had an effect or not we examine the consumption levels of households receiving each of the different sorts of treatment and look for systematic differences between them. As discussed in chapter 2 apportioning households to each treatment in a random fashion minimises the risk of systematic bias between these various treatment cells.

By way of reminder, there were nine different treatments (including the control) included in this project. One of these was the control group, who received no inserts, while the rest of the eight were inserts included with the water bill. The various treatments are described in the table below. Rather than describe an entire treatment each time it is discussed we make use of the abbreviations for each treatment that are included in the table below. Chapter 2 describes each in more detail.

Table 4.5. Inserts mailed to selected households

	Abbreviation	Components of treatment	Non-pecuniary elements added to the choice architecture of households.
1	Control	Water account	Status quo
2	DHM	DHM+ Water Account	Raised salience of own consumption (last month)
3	DHMA	(DHMA) + Water account	Raised salience of own consumption (last month)+ Social norm
4	DHMAT	(DHMA) + DT + water account	Raised salience of own consumption (last month)+ Social norm + Tips
5	DHMAE	(DHMAE) + water account	Raised salience of own consumption (last month)+ Efficient Social norm
6	DHY	DHY + water account	Raised salience of own consumption (last 12 months)
7	DHYT	DHY + DT + water account	Raised salience of own consumption (last 12 months) + Tips
8	DHYA	DHYA + water account	Raised salience of own consumption (last 12 months) + Social norm
9	DT	DT + water account	Tips (mapping inputs to outputs)

Considering the post treatment period.

The final thing we need to consider before assessing the effect of the mailed treatments is to consider the period during which we are likely to pick up the treatment effect.

From a household's meter being read to its bill being printed a period of around five working days should elapse. Once the bill has been printed it takes about a day to reach the post office. Once posted, the bill should spend between three to five working days in the mail before reaching its household. Thus, a bill is likely to reach a household roughly half way through the month following the month it is being billed for. For instance, if your meter was read on 01 November 2012, which we consider "November consumption" since the meter reading date was in November, you were likely to receive your bill, (together with its insert – unless you were in the control group or outside of the experimental sample) on roughly the 15th of November 2012. Considering that a household's meter will usually be read at about the same time every month, this would mean that you would have received your bill (and its insert, if you were in a treated group) about half a month before your next meter reading. In other words, treatments posted in November 2012 would have reached households in the middle of what is considered their December consumption period. The implication is that households would have had only half a month to change their behaviour (or not change their behaviour) in response to any insert they received.

Given this, we would expect any effect of the inserts to be muted in December. The first full month during which we could observe a treatment effect, assuming that it persisted for that month, is January 2013. As a result we consider December 2012, January 2013 and February 2013 as our post-treatment period. We do this separately for each month.

Calculating the effects of different treatments

As discussed in chapter 2, the randomisation strategy used to achieve balance between the various treatment cells renders the chance of systematic bias existing between different treatment cells very small. This means that the only likely systematic difference between the treatment cells will be due to the treatment itself. This in turn means that the analytic strategy to identify a treatment effect is fairly straightforward. What is required is to measure the mean consumption for each treatment cell and then to test whether these differences are significant or not. The differences will then give us an idea as to what the treatment effect was. We achieve this by means comparisons as well as, regression analysis. Although estimates of the treatment effect calculated by each of these two approaches will not necessarily be precisely the same, they should certainly be consistent with each other.

Table 5 reports the mean consumption for each post treatment month for each treatment cell, within the constraints described in sections 1 and 2 above. Each of the means for the treatment cells that received an insert are compared to the mean for the control group and tested, using a simple t-test, to determine whether they are different to the control group's mean at conventional significance levels.

Table 6 reports the results for our regression analysis. The regression are simple Ordinary Least Squares regressions. For each regression the insert-based treatments are included as dummy variables with the control treatment omitted as the reference category. In such a set up the regression allows us to assess what the average treatment effect of each insert was. Two sorts of regressions are run: one which calculates standard errors in the straightforward fashion and another which calculates them in a robust fashion.

4.3 Treatment effect within the total municipal sample: results

In this section we discuss the treatment effects that seem to be evident in the sample as a whole. This discussion of results is broken up into four parts. We begin in section 4.3.1 with the effect of receiving any insert whatsoever; here we compare the consumption of the pooled sample of households who received any of the treatments against the consumption of households in the control sample who received no sort of insert at all. The discussion of results then moves on in section 4.3.2. to consider the relative effectiveness of reporting social norms. This discussion includes an examination of the effect of receiving norms feedback as distinct from only raising the salience of consumption. This was one of the key behavioural economics' questions posed by the experimental design. This particular experiment is the first instance of these two effects being considered separately at scale that the researchers are aware of. Section 4.3.3 then moves the discussion to considering the relative effectiveness of providing information about how to save water. Results reported in this section provide a commentary on what is perhaps the most traditional method of behavioural water demand management. Section 4.4.4 discusses whether there is any relative benefit of treatments that reported a year's worth of consumption or if more was gained in terms of savings from inserts reporting only last month's consumption.

4.3.1 The effect of receiving any insert

A reasonable initial question to ask is whether receiving any insert at all made an impact upon household consumption. At the very least each of the inserts, as has been argued, raised the salience of water consumption in their own way. Accordingly we would expect that households that received any of the inserts would reduce their consumption.

Encouragingly it does seem to be the case that small but significant reductions in consumption are associated with some of the treatment inserts. Overall, it seems that treatment effects are strongest in January, the first full post-treatment month, potentially beginning to diminish in February. Treatment effects are least noticeable in the December period when households would have had little time to react to any inserts they had received before their next meter reading.

Table 4.6. Kilolitres saved by households receiving any type of insert compared to the control group

	December 2012		January 2013		February 2013	
	OLS	Robust std. errors	OLS	Robust std. errors	OLS	Robust std. errors
Receiving any insert	-0.0008 (0.0018	-0.0008 (0.0018)	-0.0041** (0.0018)	-0.0041** (0.0018)	-0.0039** (0.0019)	-0.0039** (0.0019)
N	267 749	267 749	325 572	325 572	278 681	278 681
F	0.2108	0.2107	5.1309	5.0843	4.2889	4.2486
P (F=0)	0.6462	0.6462	0.0235	0.0241	0.0384	0.0393

Significantly different from zero at the 10% level *, 5% level **, 1% level *** (standard errors within brackets)

As reported in table 4.6, regression estimates indicate that the average reduction for a household receiving any of the inserts relative to the control appears to hold steady at around 4 litres per day on average.

This may seem like a small amount of water per household per day, but we should not dismiss such a reduction too quickly. The following should be borne in mind. First, these reductions are averages over a very large number of households. The sheer scale of this experiment means that, if we exclude outliers, the direct savings of this experiment alone were 32,436,426 litres in January and 23,845,021 litres in February. That is a little over 55 million litres of water saved in the two highest water consumption months for the City of Cape Town. Second, there is likely to be some heterogeneity of the treatment effect in several planes, but most especially between treatment types. Different treatments may perform differently overall, in which case it would be better to deploy the more effective treatments only. The mean consumption for each treatment cell reported in Table 4.7 and the regression results reported in Table 4.8 suggest there is heterogeneity of treatment effect. We discuss this in the remainder of this report. Finally different treatments may meet with different success amongst various sub-samples of households.

Table 4.7: Mean average daily consumption for each month (kilolitres) for each treatment

		Control	DHM	DHMA	DHMAE	DHMAT	DHY	DHYA	DHYT	DT
	Mean	0.6729	0.6696	0.6708	0.6704	0.6700	0.6791	0.6725	0.6716	0.6721
December	*Significance*						**			
2012	*(Std. dev.)*	*(0.3812)*	*(0.3795)*	*(0.3779)*	*(0.3786)*	*(0.3795)*	*(0.3858)*	*(0.3828)*	*(0.3817)*	*(0.3823)*
	n	57 866	26 398	25 940	26 098	26 500	26 492	26 469	25 628	26 358
	Mean	0.7499	0.7428	0.7425	0.7438	0.7462	0.7524	0.7451	0.7427	0.7508
January	*Significance*		***	***	**			*	***	
2013	*(Std. dev.)*	*(0.4255)*	*(0.4197)*	*(0.4193)*	*(0.4199)*	*(0.4233)*	*(0.4241)*	*(0.4224)*	*(0.4218)*	*(0.4260)*
	N	70 368	32 241	31 464	31 804	32 121	32 212	32 144	31 143	32 075
	Mean	0.7105	0.7064	0.7050	0.7013	0.7028	0.7145	0.7060	0.7075	0.7090
February	*Significance*			*	***	***				
2013	*(Std. dev.)*	*(0.4128)*	*(0.4097)*	*(0.4077)*	*(0.4034)*	*(0.4089)*	*(0.4120)*	*(0.4103)*	*(0.4127)*	*(0.4104)*
	n	60 320	27 483	27 035	27 111	27 520	27 539	27 521	26 685	27 467

Significantly different to the control group at the 10% level *, 5% level **, 1% level ***

Table 5.8. Regression results: Estimated treatment effects

Estimated difference between the average daily consumption (in kilolitres) for Control and the various treatment cells – December, January and February 2013

For each regression, the estimate of the treatment effect is reported in the following fashion

Coefficient estimate of the treatment effect (significantly different from zero)
(standard error)

Significantly different from zero at the 10% level *, 5% level **, 1% level ***

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	Robust std. errors	OLS	Robust std. errors	OLS	Robust std. errors
DHM	-0.0033 (0.0028)	-0.0033 (0.0028)	-0.0071** (0.0028)	-0.0071** (0.0028)	-0.0041 (0.003)	-0.0041 (0.003)
DHMA	-0.002 (0.0028)	-0.002 (0.0028)	-0.0074* (0.0029)	-0.0074* (0.0029)	-0.0054*** (0.003)	-0.0054*** (0.003)
DHMAE	-0.0024 (0.0028)	-0.0024 (0.0028)	-0.0060** (0.0029)	-0.0060** (0.0028)	-0.0092* (0.003)	-0.0092* (0.003)
DHMAT	-0.0028 (0.0028)	-0.0028 (0.0028)	-0.0037 (0.0028)	-0.0037 (0.0029)	-0.0077* (0.003)	-0.0077* (0.003)
DHY	0.0062** (0.0028)	0.0062** (0.0029)	0.0025 (0.0028)	0.0025 (0.0029)	0.004 (0.003)	0.004 (0.003)
DHYA	-0.0004 (0.0028)	-0.0004 (0.0028)	-0.0048*** (0.0028)	-0.0048*** (0.0029)	-0.0045 (0.003)	-0.0045 (0.003)
DHYT	-0.0012 (0.0029)	-0.0012 (0.0029)	-0.0072** (0.0029)	-0.0072** (0.0029)	-0.003 (0.003)	-0.003 (0.003)
DT	-0.0007 (0.0028)	-0.0007 (0.0028)	0.0009 (0.0028)	0.0009 (0.0029)	-0.0015 (0.003)	-0.0015 (0.003)
Summary statistics						
Number of accounts	*267 749*	*267 749*	*325 572*	*325 572*	*278 681*	*278 681*
F statistic	*1.4784*	*1.456*	*2.9299*	*2.9295*	*2.864*	*2.8781*
Significance of the Overall regression (P value of the F-Test)	*0.1591*	*0.1676*	*0.0028*	*0.0028*	*0.0035*	*0.0033*

4.3.2 The effect of reporting social norms

Reporting social norms is a strategy of managing consumption which seems to have met with encouraging results in the developing world and, as a results, is gaining popularity there. The researchers were interested to investigate whether similar strategies would meet with similar results in South Africa and as a result four such strategies were included for evaluation in this project (these being the treatments abbreviated as DHMA, DHMAE, DHMAT and DHYA).[11]

The results of comparing the norm treatments with the control group in table 4.8 above (highlighted for convenience in table 4.9 below) are encouraging. These results show that households receiving norm treatments consumed significantly less, on average than households in the control group who received nothing.

Table 4.9. Regression results: Estimated treatment effects for Norms treatments relative to the control group (average kilolitres per day). Reproduced from table 4.8.

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	*Robust std. errors*	*OLS*	*Robust std. errors*	*OLS*	*Robust std. errors*
DHMA	-0.002 (0.0028)	-0.002 (0.0028)	-0.0074* (0.0029)	-0.0074* (0.0029)	-0.0054*** (0.003)	-0.0054*** (0.003)
DHMAE	-0.0024 (0.0028)	-0.0024 (0.0028)	-0.0060** (0.0029)	-0.0060** (0.0028)	-0.0092* (0.003)	-0.0092* (0.003)
DHMAT	-0.0028 (0.0028)	-0.0028 (0.0028)	-0.0037 (0.0028)	-0.0037 (0.0029)	-0.0077* (0.003)	-0.0077* (0.003)
DHYA	-0.0004 (0.0028)	-0.0004 (0.0028)	-0.0048*** (0.0028)	-0.0048*** (0.0029)	-0.0045 (0.003)	-0.0045 (0.003)

Significantly different from zero at the 10% level *, 5% level **, 1% level ***

It is important to note that these results reported below are comparisons between the control group and norms treatments as a whole. This is a worthwhile comparison to make, since it gives us an idea of the overall effect of reporting norms comparatively. However, comparisons of this nature are not really tests of the hypothesis about how norms treatments work.

Informed by behavioural economics and social psychology the central hypothesis about how reporting norms will change behaviour is that consumers will alter their consumption towards the norm. This means that those reported as having consumed more than the norm will reduce their consumption more than they would otherwise have done as they are prompted to move their consumption towards the norm. The same logic implies that households who were reported as having consumed less than the norm will raise their consumption towards the norm more than they would otherwise have done.

We can express this formally as

$$\Delta HH\_Consumption_{it} = \beta_i(HH\_Consumption_{it-1} - \mu_{it-1})$$

11 See Table 4.5 for a summary of each of these treatments.

where $(HH\_Consumption_{it-1})$ is the household's consumption and the norm reported to that household's neighbourhood is (μ_{it-1}) . In this simple expression we would expect that $\beta_{it} < 0$ since we expect that households would tend towards the norm.

It is clear then that if we are to test the hypothesis about how reporting a norm may affect a household's consumption that we will need to assess the consumption response of those reported as having been above the norm separately to the consumption response of those reported as having consumed less than the norm.

Before we can move on to identify the treatment effect of a comparative norms strategy we must deal with one more obstacle that stands in the way of that identification. In the natural course of events it is plausible that households will display perturbations around some expected consumption level for that month. That is, there is a level of consumption which we would expect a normal household in the neighbourhood to display. It is plausible that in some months households may overshoot this mean consumption and some months they may undershoot it. Over a reasonable period of time however, such as a year, we may see them adjusting consumption naturally toward this dynamic mean.

In trying to identify whether a household responded to their reported consumption relative to the consumption of an average household in their neighbourhood (the norm in these treatments) it is necessary to untangle this potential "natural" tendency of a household to revert to this mean consumption level[12] for the month in question from a response that sees the household move towards a mean when presented with it in the sort of feedback provided in the norms treatments evaluated in this project.

In order to explain the identification strategy we develop a simple but sufficient model of household water consumption below.

If households naturally change their consumption month to month in response (ρ_i) to changes in several factors (such as climate, holidays) from month t-1 to month t, captured in the vector X_{it}, plus some unintentional error (e_{it}) of over-shooting and undershooting the natural change then the change in the ith household's consumption from month from month t-1 to month t, $\Delta HH\_Consumption_{it}$ is expressed as

$$\Delta HH\_Consumption_{it} = \rho_i X_{it} + e_{it}$$

The general form of how we would expect households that were exposed, in the prior month, to a norm treatment that reported their monthly consumption $(HH\_Consumption_{it-1})$ relative to the monthly consumption for a normal household for their neighbourhood (μ_{it-1}) to change their consumption can then be expressed in the following fashion

$$\Delta HH\_Consumption_{it} = \beta_i (HH\_Consumption_{it-1} - \mu_{it-1}) + \rho_i X_{it} + e_{it}$$

with the expectation that $\beta_{it} < 0$ since we expect that households would tend towards the norm.

Our general model with respect to households exposed to the norm treatment $(N_{it} = 1)$ and households not exposed to a norms treatment $(N_{it} = 0)$ becomes

$$\Delta HH\_Consumption_{it} = \beta_i (N_{it})(HH\_Consumption_{it-1} - \mu_{it-1}) + \rho_i X_{it} + e_{it}$$

12 Most likely the mean would be dynamic.

wherein we still expect that $\beta_{it} < 0$.

Given the random allocation of accounts to treatments and the lack of pre-consumption bias between treatments we may assume that $\rho_i X_{it}$ and e_{it} are independently distributed of N_{it} and that they also follow distributions that are not significantly different across values of N_{it}. As a result we the likely difference in household consumption ($HH\_Consumption_{it}$) across values of N_{it} as

$$HH\_Consumption_{it} = HH\_Consumption_{it-n} + \left[\sum_{n=1}^{j} \Delta HH\_Consumption_{it-n+1}\right]$$

$$= HH\_Consumption_{it-n} + \left\{\sum_{n=1}^{j} [\beta_i(N_{it-n+1})(HH\_Consumption_{it-n} - \mu_{it-n}) + \rho_i X_{it-n+1} + e_{it-n+1}]\right\}$$

thus: $(HH\_Consumption_{it-n}|N_{it-n+1} = 1) - (HH\_Consumption_{it-n}|N_{it-n+1} = 0)$

$$= \sum_{n=1}^{j} [\beta_i(HH\_Consumption_{it-n} - \mu_{it-n})]$$

which is the treatment effect of a norms treatment that compared the prior month's consumption of a household to a normal household for the neighbourhood.

What is clear from the above model is that identifying the treatment effect of the comparative norms insert depends crucially upon the assumption that $\rho_i X_{it-n+1}$ and e_{it-n+1} follow the same distribution between treatments, are not biased between treatments and are independant of all treatments. Randomisation has taken care of the independance, distribution concern and bias concern across all treatments as a whole. However, as discussed above, we would introduce bias in $\rho_i X_{it-n+1}$ and e_{it-n+1} if we were to compare households that had consumed more than the average household in their neighbourhood (receiving norms treatments) with the rest of the entire distribution of households in other treatments in the neighbourhood.

Put more simply, the implication for this project and assessing the effect of the normative treatments is that if we compare the tendencies of households that were reported as having consumed more/less than an average household in their neighbourhood against the aggregate control group we run the risk of comparing a group made up predominately of households that were systematically consuming more than their "natural" mean consumption for that month with a group reflects a more balanced composition of households. There will thus be a bias between the two groups that is unrelated, causally, to the norms treatments.

Since distributions between treatments in respect to pre-treatment consumption were not significantly biased and were independent of the treatments we may overcome this difficulty simply by splitting the distribution of all households in all treatments according to their consumption relative to the reported norm for that month (here, the average daily consumption reported in November 2013) and comparing the groups above/below that average with other households in other treatments that were above/below that norm. Doing so removes systematic pre-treatment bias and allows us to identify the consumption response to the norms treatment as distinct from a natural reversion to some mean level of consumption.

For the purposes of this report we only compare households that were above the November norm reported to the DHMA treatment in the November 2012 mailing period. This period is when the overwhelming majority of households were mailed. Very few households were mailed the October 2012 norms.[13] Results off such a small base would not be significant.

Table 4.10 reports the results of a comparison between the Monthly Norm treatment DHMA household's consumption and control group households in terms of average kilolitres consumed per day. It reports results for regressions that compared households that were reported as being above the November norm with households in the control group that would have been reported as having been above the November reported norm, had they been placed in the DHMA treatment which reported this.[14]

Table 4.10. Monthly Norm treatment DHMA household's consumption compared to control group households (average kilolitres consumed per day)

Treatment	Above the norm			Below the norm		
	December	January	February	December	January	February
DHMA	-0.0097* (0.0033)	-0.0133* (0.0034)	-0.0124* (0.0035)	-0.0007 (0.0038)	0.0024 (0.004)	0.0052 (0.0043)
N	52 385	63 625	53 146	21 851	26 038	23 407
F	8.671	15.6159	12.5637	0.0344	0.3426	1.4943
P (F=0)	0.0032	0.0001	0.0004	0.8529	0.5583	0.2216

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

From the results in table 4.10 we can see that households which were reported as having consumed more than the norm in the November treatment ended up consuming significantly less, at the 10% confidence level, than similar households in the control group in all of the post-treatment months (between 9.7 litres and 13.3 litres per household per day on average). This is consistent with expectations around how households would tend to alter their consumption towards norms.

In contrast, for households that were reported as having consumed less than the norm in the DHMA treatment, we were unable to detect any significant difference between their consumption and the consumption of households in the control group. This is heartening since it means that for households reported as being below the norm the tendency to increase their consumption more than they would have otherwise is either not present, or is only weakly present.

13 Recall from chapter 2 that households whose reading date fell into October 2012 would be compared to the norm for their neighbourhood in October and households with their reading date in November 2012 would be compared to the November average household for their neighbourhood.

14 Since randomisation was not conducted to obtain balance with respect to households' relative position to the November reported norm in the pre-treatment phase (that would have been impossible) we had to control for outliers, especially in the upper ranges of the consumption distribution when running these regressions. These regressions were all run over a sample that excluded the highest 5 percentiles of consumption for 2012's consumption up to the last pre-treatment month (October 2012 for these households). Households consuming less than free basic water are excluded.

The average treatment effect for a household in the DHMA treatment that was reported as having consumed more than the norm appears to be between 12.4 litres and 13.3 litres per household per day in January and February. These months are the months during which water consumption in Cape Town tends to peak, they are also the first full months that households have to respond to this project's experimental treatments. Both of these factors may explain why the treatment effect appears to have been larger in January and February (2013) than it was in December 2012.

The results above suggest that it would be a viable water demand management strategy to report to households that were consuming more than the median household in their neighbourhood this fact.

We now move our discussion to consider whether the effect of comparing a household's water consumption to a norm is significantly different to merely raising the salience of a household's water consumption. For the reasons already outlined our identification strategy in this case again necessitates that we consider households reported above a norm separately to households reported as having consumed less than the norm. To allow us to do this we designed this experiment with two treatments which both raised the salience of a household's water consumption in the prior month (DHM and DHMA) but differed only in the fact that one compared that consumption to an average household for their neighbourhood (DHMA). Any difference between the two treatments could thus be attributed to this comparison to a norm rather than to the increased salience of a household's own water consumption. Table 4.11 reports the results of these comparisons.

Table 4.11. The effect of a comparison to the "average household in your neighbourhood". Monthly Norm treatment DHMA household's consumption compared to monthly consumption for salience treatment households in the DHM treatment (average kilolitres consumed per day).

Treatment	Above the norm			Below the norm		
	December	January	February	December	January	February
DHMA	-0.0041 (0.0039)	-0.0033 (0.0039)	-0.0090** (0.0042)	0.0066 (0.0044)	0.0078*** (0.0047)	0.0135* (0.005)
N	30656	37174	30857	12832	15258	13637
F	1.0893	0.7143	4.674	2.2456	2.7101	7.3075
P (F=0)	0.2966	0.398	0.0306	0.134	0.0997	0.0069

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

From the results in table 4.11 there is weak evidence that the effect of reporting that a household had consumed more than the norm reduces their consumption more than raising the salience of their previous month's consumption would. The difference in consumption between households reported as being above the norm in the DHMA treatment and households in the salience-only DHM treatment who would have been reported as having consumed more than the norm is only ever significant for one post-treatment month, February. Thus, although the signs on the estimated differences are as we would expect, the treatment effects of a comparison to a norm and merely raising salience appear to be not very different for households that would be reported as having consumed more than the norm.

For households that would have been reported as having consumed less than the norm the difference between the salience-only treatment DHM and the norms treatment DHMA, would allow us to assess if there was a tendency for households to raise their consumption toward the norm by controlling for the effect of raising salience. From table 4.11 we see that the consumption of households reported as having consumed less than the average household in their neighbourhood is significantly higher in January and February than that of similar households who received the salience only treatment. The results in table 4.11 thus suggest that when households had a full month or more to respond to the treatments those households did indeed, on average, raise their consumption in response to being reported as having consumed less than the norm.

Taken together, results for the post treatment consumption of households reported as having consumed more than the norm and this reported as having consumed less indicate that households did adjust their consumption towards the norm. For households reported as having consumed more than the norm, this meant a reduction in consumption. This reduction in consumption was not generally, however, of sufficient size to be consistently greater than the reduction resulting from merely raising the salience of consumption. So, while it seems that there is benefit to employing a comparative norm treatment, this benefit is not very much larger than that which is gained by raising the salience of consumption. For households reported as having consumed less than the norm an undesirable increase in consumption was observed relative to the salience-only treatment. For these households it seems that it would be best to not compare them to a norm but to merely raise the salience of their own consumption.

We now turn our attention to whether there was any benefit to be gained by trying to indicate an efficient norm. The DHMAE treatment was the same as the DHMA norm in all respects except that the consumption figure it reported would have been 6% lower than the DHMA norm for the same neighbourhood and that the norm was referred to as the "estimated average efficient household in your neighbourhood" rather than the "estimated average household in your neighbourhood".

As with the DHMA norm, estimating the effect of the DHMAE norm requires that we split the sample of households in these two treatments into two groups according into households that were reported as having been above the estimated efficient norm, or would have been, and households that were reported as having consumed less than the estimated efficient norm, or would have been.

Since the DHMAE norm was calculated as being 6.7% less than the DHMA norm for each neighbourhood[15] it is different to the DHMA norm for each neighbourhood. This also means that it would be possible for some households to have been reported as having been below the DHMA norm while households with similar pre-treatment consumption in the DHMAE treatment would have been reported as having consumed more than the norm. In this case it is plausible that one group would respond by increasing their consumption towards the norm they were below while the other group would respond by decreasing their consumption toward the norm they were above. In both cases the norm would have been effective in the manner we expected, the difference between the two groups would potentially only be due to the fact that the kilolitre value of one norm was lower than the kilolitre value of the other norm. In order to avoid this sort of situation we define our "above the norm" group as households that were above both the DHMA and DHMAE norm, and our "below the norm" group as households that were below both the DHMA and DHMAE norm. Doing this places us in a position where we are comparing households that have similar pre-treatment consumption, both are above/below a norm with the major difference being that one is the efficient

[15] See section 2.1.2 "Comparative Social Norms of Efficiency".

norm. This allows us to test whether the combination of a slightly lower norm labelled as "the estimated average efficient household" has any effect that is significantly different to the DHMA norm which is described as "the estimated average household in your neighbourhood".

Table 4.12 reports the results of these comparisons. From the table we can see that there were no significant consumption differences between the DHMA and DHMAE groups in any of the post-treatment months. This result does not differ by whether the households were reported as having been above or below the norm. The clear implication is that describing a norm as "efficient" does not seem to make a significant difference to household water consumption in this context.

Table 4.12. The effect of a comparison to the "estimated average efficient household in your neighbourhood" in the DHMAE treatment rather than to the "estimated average household in your household" of the DHMA treatment. DHMA compared to DHMAE (average kilolitres consumed per day). Regressions with robust errors.

Treatment	Above the norm			Below the norm		
	December	January	February	December	January	February
DHMAE	-0.005 (0.0040)	-0.004 (0.0040)	-0.0026 (0.0043)	-0.0011 (0.0048)	-0.0034 (0.0051)	-0.0076 (0.0055)
N	28 870	35 288	28 936	10 705	12 685	11 446
F	1.5585	0.9595	0.3605	0.0509	0.4332	1.9618
P (F=0)	0.2119	0.3273	0.5482	0.8216	0.5105	0.1613

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

Reporting norms and annual consumption

Two treatments (DHY and DHYA) within this RCT received inserts with their water bill that reported their household's average daily consumption, as a histogram, for each of the past twelve months for which the household's meter was read. One of these two treatments (DHY) received an insert that only reported consumption for the past twelve months in this fashion. Another treatment, DHYA, received an insert that placed a straight line over the twelve month consumption report. The line indicated a single figure that was the daily average consumption for the past year for an average household in the neighbourhood and was a form of social norm. The WRC was particularly interested in assessing this sort of norm. Since this line was also the only difference between the DHYA and the DHY treatment comparing DHYA with DHY allows us to identify what the effect of adding this comparative social norm was.

Unlike the month based norm treatments (DHMA and DHMAE) it is not feasible in the case of the DHYA treatments to separate out accounts on the basis of being reported as being above or below the norm. The reason for this is the multitude of above/below combinations between the represented annual household consumption and the annual average to allow for a meaningful consideration of each separately in this report. The value of this comparison is mainly in seeing whether introducing a norm as a reference point was able to induce systematically lower levels of consumption. That said, the results do potentially suggest a hypothesis from the information overload literature that is interesting from the point of view of behavioural economics.

Table 4.13 reports the results of the regression analyses comparing DHY and DHYA. From the results in table 4.13 it seems that adding in a comparative norm is associated with a level of consumption that is significantly lower than reporting a monthly history for the past twelve months.

Table 4.13. Regression results: estimated average difference in average daily consumption between DHY and DHYA

	December 2012		January 2013		February 2013	
	OLS	Robust std. errors	OLS	Robust std. errors	OLS	Robust std. errors
DHYA	-0.0066** (0.0033)	-0.0066** (0.0033)	-0.0073** (0.0033)	-0.0073** (0.0033)	-0.0085** (0.0035)	-0.0085** (0.0035)
N	52 961	52 961	64 356	64 356	55 060	55 060
F	3.9216	3.9217	4.7368	4.7368	5.8677	5.8677
P (F=0)	0.0477	0.0477	0.0295	0.0295	0.0154	0.0154

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

It is incredibly interesting that these two treatment groups should differ in consumption when the only difference between them was a straight line added to the graph. Although it can only be speculative, we attempt an explanation for why this may be the case below.

First, it seems unreasonable that a norm could have an effect that would not be at least imperfectly additive with increased salience in the month case. In other words if there was an effect flowing from receiving a comparative norm to consumption then we would have expected to find a significant difference between the salience-only and comparative norm treatments in *both* the month reporting and twelve month reporting scenarios. Since we do not observe this something else must be going on.

Given this, the researchers suspect that the causal mechanism underlying the difference between the DHY and DHYA treatment group is similar in nature to that noted in the information and choice overload literature.[16] As noted earlier in the literature review, the benefit of increasing the amount of information presented to a person often seems to decline beyond a certain amount. One of the reasons offered for this in the literature, which has relevance to the present case, is that increasing the amount of data offered to a person increases the cognitive burden that needs to be borne in order to process the data, figure out what the data means and what action to take. This explanation would be consistent with the lower consumption observed amongst groups receiving a report of last month's consumption only (DHM) rather than the last month's consumption (DHY).

This still begs the question, however, of why the addition of a single line generates a difference in consumption when twelve months are reported. An element, from the choice overload literature, which may be at work here is that of information categorisation. Some research within consumer choice (Scheibehenne, Greifeneder & Todd, 2010; Mogilner et al., 2008; Diehl, 2005; Kornish and Lynch, 2003, Huffman and Kahn, 1998; Russo, 1997) reports that presenting options in clear categories and orders may increase customer satisfaction, while the lack of clear categorisation and ordering of information may create confusion, dissatisfaction and ultimately, one assumes, lack of

[16] The authors are grateful to Rachel Glennerster for first suggesting this explanation to us.

action or sub-optimal action. By adding in a line what may have happened is that a potentially rich, but somewhat overwhelming, set of data (the last twelve months of water consumption) may have been transformed into a data set consisting of two clear categories: my household and my neighbour. Within this environment the cognitive load may have been lightened by allowing a quick, if fuzzy, comparison between my household and my neighbour thereby raising the salience of the household's water consumption and eliciting a slightly more careful use of water than before.

4.3.3 Providing information about how to save water

We turn now to examine the effect of providing information about how a household can save water. This information was expressed in a series of tips. The tips used in this study were derived from the City of Cape Town's Smart Living Handbook and from the Water Demand Management section of the City of Cape Town's website. The tips themselves are expressed in the "DT" treatment included in the appendix to this report.

Considering a treatment that provides tips, especially tips alone, is relevant to this study for two reasons. First, households may not know how to save water and thus, they do not save water unless by accident. Second, and extremely relevant for water demand management policy, is that providing information about how to save water seems to have been the most frequently used approach in the past.

The treatment which delivered tips alone was the DT treatment. It was an A4 sheet included with the water bill and upon it were the tips described above. In addition two other treatments combined the tips sheet with other treatments. One was the DHMAT treatment, which combined the DT treatment with the monthly comparative social norm treatment DHMA. The other treatment with which the DT sheet was combined was the DHY treatment that reported average daily consumption for each month that had been billed out of the past twelve months. We begin by considering the post-treatment consumption of households that received only tips (DT) to the control group. We will then consider the marginal benefit of adding tips to the month-based norm treatment (DHMAT versus DHMA) and to the treatment reporting annual consumption (DHYT versus DHY).

Table 4.14. Regression results: estimated average difference in average daily consumption (kilolitres) between DT and the Control group (reproduced from table 4.8).

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	*Robust std. errors*	*OLS*	*Robust std. errors*	*OLS*	*Robust std. errors*
DT	-0.0007 (0.0028)	-0.0007 (0.0028)	0.0009 (0.0028)	0.0009 (0.0029)	-0.0015 (0.003)	-0.0015 (0.003)

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

Table 4.14 reproduces, for convenience, the estimated difference between DT and the control group in terms of each group's average daily water consumption reported earlier in table 4.8. The size of the estimated differences in each post treatment month are quite small, around 1 litre per day on average) and for no month are they significantly different to zero. It is thus extremely unlikely that households who received the information-only treatment (DT) systematically consumed less than they would have otherwise. This experiment thus failed to identify any benefit of reporting tip alone on an information sheet.

Although there seems to have been no effect upon consumption of receiving tips alone it is plausible that there could have been a significant interaction between a relatively strong, salience-raising treatment such as DHMA and receiving information about how to save water. One can imagine a household paying more attention to the tips if they pay more attention to their water consumption in general. To investigate this we compare the post-treatment consumption of households that received the month normative treatment DHMA and households that received the DHMA insert and the DT insert (the DHMAT treatment). In this comparison, the only difference between the two groups of households is the receipt of the tips. A comparison of these two treatment is therefore able to identify whether there is any marginal benefit of adding tips to the month-based norm treatment. For the reasons discussed in section 4.3.2 we split the analysis by considering households reported as above or below the norm separately.[17]

Table 4.15. The effect of adding tips to the month-based norm treatment a comparison to the "average household in your neighbourhood". Regressions comparing the post treatment average daily consumption (average kilolitres consumed per day) of households in the DHMA and DHMAT treatments. Regressions with robust errors.

Treatment	Above the norm			Below the norm		
	December	January	February	December	January	February
DHMAT	-0.0018 (0.0039)	0.0022 (0.0040)	-0.0016 (0.0041)	0.0000 (0.0045)	-0.0026 (0.0048)	-0.0057 (0.0051)
N	30 856	37 230	30 974	12 703	15 108	13 514
F	0.2142	0.2957	0.1458	0.0000	0.2933	1.2494
P (F=0)	0.6435	0.5866	0.7026	0.9984	0.5881	0.2637

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

From table 4.15 we can see the results of this comparison and it is clear that there was no significant difference in the post treatment consumption of households that received the monthly norm treatment DHMA and households that received the monthly norm feedback augmented with tips (DHMAT). This was the case for groups that were reported as having been above the norm as well as those reported as having been below the norm. There thus appears to have been no additional savings gained by including tips along with monthly norms.

Another way in which the researchers tried to assess whether there was any marginal benefit to adding tips was in the context of by reporting monthly average daily consumption for all of the prior twelve months for which the household's meter was read. This was done by including, for a random subset of households, the "DT" information tip sheet along with the DHY histogram of annual

[17] Since randomisation was not conducted to obtain balance with respect to households' relative position to the November reported norm in the pre-treatment phase (that would have been impossible) we had to control for outliers, especially in the upper ranges of the consumption distribution when running these regressions. These regressions were all run over a sample that excluded the highest 5 percentiles of consumption for 2012's consumption up to the last pre-treatment month (October 2012 for these households). Households consuming less than free basic water are excluded.

consumption. This random subset of households was the DHYT treatment group. By comparing the post treatment consumption of these two groups, DHY and DHYT, we were able to discern whether any benefit was realised by including these tips. We have already seen that adding in a form of norm (DHYA) into this context elicited savings. It may be that adding in tips may yield savings in this context too.

Table 4.16. The effect of adding tips to the DHY treatment reporting annual consumption. Regressions comparing the post treatment average daily consumption (average kilolitres consumed per day) of households in the DHY and DHYT treatments. Regressions with and without robust errors.

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	*Robust std. errors*	*OLS*	*Robust std. errors*	*OLS*	*Robust std. errors*
DHYT	-0.0074** (0.0034)	-0.0074** (0.0034)	-0.0096* (0.0034)	-0.0096* (0.0034)	-0.0070** (0.0035)	-0.0070** (0.0035)
N	52 120	52 120	63 355	63 355	54 224	54 224
F	4.8916	4.8933	8.2149	8.2164	3.9334	3.9332
P (F=0)	0.027	0.027	0.0042	0.0042	0.0473	0.0473

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

From the results in table 4.16 it seems that there was some benefit to including tips when reporting monthly consumption over a year period. For all of the three post treatment months considered there was a significant difference in average daily consumption between the DHY and DHYT treatments. The average estimated savings ranges between 7 litres per day (in December 2012 and February 2013) and 9.6 litres per day in January 2013.

It is not clear why tips would make a difference to household water consumption in this context, but not anywhere else. Although unlikely, it is possible that this result may be the result of some unobservable statistical bias between the DHYT and DHY treatment groups that only became manifest in the post-treatment months. More plausible however is an explanation in the spirit of that used to explain the difference between the DHY and DHYA treatment earlier. Continuing in that vein it may be that the DHY treatment raised the salience of a household's water consumption, but it did so in such a manner that it was unclear to the household what they should do as a result. This led households receiving the DHY treatment only to take no special action. Adding tips about how to save water into this context may have made clear to households that when they were thinking about their water consumption they should think about reducing their water consumption. In short the overall subject matter of the tips – save water – may have served to clarify a salient but confusing message.

4.3.4 Raising the salience of water consumption by reporting a year's worth or a month's worth of consumption

The final strategy we will consider is that of simply raising the salience of a household's water consumption. We will consider and compare the two strategies for achieving this (the DHM and DHY treatments) that were tested in this experiment.

As the reader will be aware, the DHM treatment reported to the household their average daily consumption for the prior month by way of a single horizontal bar on a horizontal bar chart. An example of this sort of insert can be found in the appendix. The DHY treatment reported, by way of a histogram, the average daily consumption in each month of the prior twelve months for which a reading for the household was done. An example of this treatment can also be found in the appendix.

Table 4.17. The effect of receiving the DHM or DHY treatments consumption. Regressions comparing the post treatment average daily consumption (average kilolitres consumed per day) of households in the DHY and DHM treatments to households in the control group (reproduced from table 4.8). Regressions with and without robust errors.

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	*Robust std. errors*	*OLS*	*Robust std. errors*	*OLS*	*Robust std. errors*
DHM	-0.0033 (0.0028)	-0.0033 (0.0028)	-0.0071** (0.0028)	-0.0071** (0.0028)	-0.0041 (0.0030)	-0.0041 (0.0030)
DHY	0.0062** (0.0028)	0.0062** (0.0029)	0.0025 (0.0028)	0.0025 (0.0029)	0.0040 (0.0030)	0.0040 (0.0030)

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

The two salience-only treatments seem to have performed slightly differently to each other with respect to mitigating household water consumption. The DHM treatment, reporting only last month's consumption, was associated with significantly lower consumption than the control group, on average 7.1 litres less, in the post-treatment month of January. For the other months, the sign on the estimated effect also suggests a lower consumption level, however the difference for those months is not significant.

In contrast to the last-month-only-reporting treatment the signs on the estimated average difference between the DHY treatment and the control group is always positive, indicating higher average consumption, with this difference significant in fact significant in December. At best one can say that the effect of reporting a year's worth of consumption to a household in the DHY fashion is generally too small to be significant. At worst, it seems that sending households the DHY treatment served to nudge their consumption slightly above what it would otherwise have been.

The rationale behind the treatment which reported the household's average daily consumption for each of the past twelve months (DHY) was that reporting a household's consumption back to itself was likely to raise the salience of the household's consumption for that household. This may then lead to more efficient consumption habits and lower levels of consumption. The DHY treatment thus leverages the same behavioural element as the DHM treatment – salience alone – except it does so by reporting a different time period. By including both the DHM and the DHY treatments we wished to assess whether reporting different periods would make a difference to a household's consumption.

Regression analysis comparing the DHY treatment to the DHM treatment, reported in table 4.18 reveals that households in the DHY treatment consistently consumed significantly more water than households in the DHM treatment, in each post-treatment months. These estimates suggest that the size of this average difference in consumption ranged between 8.1 litres a day and 10 litres a day, depending on the month.

Table 4.18. Estimated average difference in average daily consumption between DHY and DHM (average kilolitres consumed per day). Regressions with and without robust errors.

Treatment	**December 2012**		**January 2013**		**February 2013**	
	OLS	*Robust std. errors*	*OLS*	*Robust std. errors*	*OLS*	*Robust std. errors*
DHY	0.0100* (0.0035)	0.0100* (0.0035)	0.0096* (0.0033)	0.0096* (0.0033)	0.0081** (0.0035)	0.0081** (0.0035)
N	52 890	52 890	64 453	64 453	55 022	55 022
F	8.1055	8.106	8.3731	8.3731	5.3226	5.3226
P (F=0)	0.0044	0.0044	0.0038	0.0038	0.0211	0.0211

Significantly different from zero at the 10% level *, 5% level **, 1% level ***
Standard errors reported in brackets below estimates

This study seems to suggest that, in terms of salience alone, it is best, on average, to report last month's figures. As we have already discussed it is not clear why reporting the last twelve months' worth of monthly daily averages performs so much poorer than just reporting last month's daily average consumption. One attractive explanation for why this may be the case, that is consistent with the balance of the other results, is that providing information about the last twelve months' worth of consumption may overcomplicate matters and impose an unhelpfully large cognitive burden upon the household for this sort of treatment. This may result in sub-optimal behaviour here as it has elsewhere in the information and choice overload literature.

4.4 Summary

Of the treatments assessed in this study, the most consistent savings relative to the control group are associated with treatments that employed a social norm comparison and did so using the prior month's consumption and strategies that raised the salience of consumption by reporting the prior month's consumption.

For the monthly norms results differ by whether the household was reported as having consumed more or less than the norm or less than the norm. Households that were reported as having consumed more than the norm demonstrate savings consistently. Although the estimated savings appear larger than the salience-only treatment the difference is not significant. This suggests that the most important reason for the effectiveness of norms may lie in their ability to raise the salience of consumption.

Households that were reported as having consumed less than the norm did not consume significantly less than the control group. These households however did consume significantly more than the households receiving the salience only treatment reporting only last month's consumption. This suggests that there is a tendency for households to move their consumption towards the norm, at least when reported as having consumed less than the norm, and that failure to observe this in comparisons to a control that received no feedback at all is due to the inability to separate the moderating effect that raising salience of consumption has upon consumption. It also suggests that the appropriate strategy to adopt with households consuming less than the norm is to raise the salience of their consumption but not compare them to a norm.

Reporting each of the months of the last twelve months of consumption for which there was a reading seems to have had an insignificant or undesired effect upon consumption. Keeping feedback simple and reporting only in terms of the prior month emerged as a superior strategy.

Providing information about how to save water by way of providing tips did not, by itself, induce water consumption savings and did not appear to add any benefit to treatments that were already effective.

Chapter 5: Concluding comments and recommendations

Many municipal areas in South Africa are water stressed for portions of the year and find themselves needing to manage the demand on their water resources. This need is likely to become more acute for large urban municipalities as South Africa seeks to facilitate economic growth and as her population continues to concentrate within the larger urban municipalities where household water demand already constitutes a significant share of total water demanded.

Historically municipalities have largely relied upon price, pressure management and information campaigns to manage the demand for water. In recent years however municipalities, mainly in the developed world, have begun experimenting with strategies that are informed by the field of behavioural economics in order to manage the demand upon their utilities. These strategies have yielded some encouraging results and, as a result, are gaining traction around the world.

This study sought to test several of the more popular behavioural economic strategies within a large South African municipality, the City of Cape Town. In doing so, this project moves the field of water studies, water studies in South Africa and behavioural economics forward in three ways.

Previous studies of utility demand management using the behavioural elements of social norms, salience and information have tended to bundle several or all of these together. As such none of these studies had been able to assess the effectiveness of each of these elements individually, even although hypotheses about how each may be working has motivated their inclusion in the policy bundle. This study is the first study that the researchers are aware of to assess the elements of norms, salience and information separately in real-world setting at scale. We think that as a result, this study will be of significant benefit to the discipline of behavioural economics at large.

This study is also of benefit to the field of water studies generally in that it is one of the first, if not the first, assessment at scale of water demand management by these strategies in the developing world. Assessing such strategies in the developing world is important since the infrastructure within developing world municipalities is likely to be different to that in the developed world in that it is more heterogeneous and may allow less control overall. The daily realities of developing world households area also likely to be different, on average, to their developed world counterparts.

Of benefit to the field of water studies in South Africa is the methodology used to test the various hypotheses in this study. Designing assessments of potential water policy as a Randomised Control Trial allows researchers to be able to isolate and identify the effect of specific policy elements in a very efficient manner while still being earthed in the "real world". It is likely that the field of water studies will benefit in many of its practical questions by adopting this methodology.

The researchers find that small but significant savings in household water consumption can be induced by several of the strategies assessed in this study.

The strategies which yielded the clearest savings were those that reported the household's water consumption for the prior month and compared it to the norm for their neighbourhood when that consumption exceeded the norm for the neighbourhood. Households in these cases lowered their consumption more than they would otherwise have done so.

Although households reported as having consumed more than the norm for their neighbourhood were more frequently demonstrated to have consumed less than the control group than households to whom was reported their prior month's consumption, these two groups never had significantly different levels of consumption. This suggests that salience is likely to be the most important reason for why norms treatments, which unavoidably bundle norms proper and salience, have an effect.

This also suggests an appropriate strategy for managing the water demand of households that consume less than the norm. Such households, if compared to a norm that they are below, demonstrated a tendency to raise consumption more than similar households who only received information about their prior month's consumption. For such households the best strategy appears to be to merely raise the salience of water consumption.

The way in which salience of water consumption is raised appears to be important, however. In this study two strategies were considered. One strategy reported only the prior month's consumption, while the other reported all of the previous twelve months for which there was a reading. Reporting the prior month only induced consumption that was clearly and consistently lower than the strategy reporting the past twelve months' worth of monthly consumption. Moreover, there are worrying indications from this study that raising salience by reporting the last twelve months of consumption may induce higher consumption than would otherwise have been the case. It seems that it is best to keep thing simple and report only the prior month's consumption.

Providing information about how to save water, in the form of tips, was also assessed. The main reason for the inclusion of this strategy was that this has perhaps been the most frequently employed non-price or non-pressure management water demand management strategy historically, outside of crisis periods. The balance of evidence from this study is that providing tips in themselves make very little difference to water consumption. If they are to be used they should at the very least be coupled with the sort of norms used in this study or interventions that raise the salience of a household's water consumption.

Overall this study required very little by way of active administration in order to deploy the various treatments themselves. All that was required was access to the administrative water billing data and clear instructions to the printers of the water bills in terms of algorithms to use. Direct set up time of the treatments themselves for this project did not exceed four months and once, set up, the system ran smoothly. The direct cost of reaching the scale of households involved in this study did not exceed R100 000. This is much cheaper than many other strategies of water demand management.

The results of this study are encouraging. In the post treatment month of January the savings of water directly attributable to the 278 681 non-control households analysed in this project was 27 753 683 litres; for the shorter month of February this was a direct savings of 23 845 021 litres. Considering that this study identifies strategies which performed much better than others, it is very likely that greater savings are possible if an appropriate mix of these strategies is employed.

Overall this study suggests a new set of tools with which municipalities may manage their domestic water demand.

References

* Agostinelli, G., Brown, J., & Miller, W. (1995). Effects of normative feedback on consumption among heavy drinking college students. Journal of Drug Education, 25, 31-40.

* Akerlof G. (1980). “A Theory of Social Custom, of Which Unemployment May be One Consequence.” *The Quarterly Journal of Economics*, Vol. 94, No. 4 (Jun., 1980), pp. 749-775

* Allcott H. (2011). "Social norms and energy conservation" J. Public Econ. Volume 95, Issues 9-10, October 2011, Pages 1082-1095

* Ariely D., Ayal S., & Gino F. (2009). “Contagion and Differentiation in Unethical Behavior : The Effect of One Bad Apple on the Barrel.” Psychological Science (20) 3. Pp. 393-8.

* Asch, Solomon. (1962). “Group forces in the modification and distortion of judgements”, Chapter 16 (p.450-501) in “Social Psychology” (7th printing) by Solomon Asch. Prentice Hall, New Jersey. 7th printing: June, 1962. First printing: September 1952.

* Bandura, A., Ross, Dorothea, Ross, Sheila A. (1961). "Transmission of aggression through imitation of aggressive models." The Journal of Abnormal and Social Psychology, Vol 63(3), Nov 1961, 575-582. doi: 10.1037/h0045925.

* Barrett , Christopher B. and Carter, Michael R. (2010). "The Power and Pitfalls of Experiments in Development Economics: Some Non-random Reflections." Applied Economic Perspectives and Policy (2010) volume 32, number 4, pp. 515-548.

* Berkowitz AD. (2005). "An overview of the social norms approach." In Lederman LC, Stewart LP, eds. *Changing the Culture of College Drinking: A Socially Situated Health Communication Campaign*. Cresskill, NJ: Hampton Press; 2005:193-214.

* Bruhn, Miriam and McKenzie, David. (2009). " In Pursuit of Balance: Randomization in Practice in Development Field Experiments." American Economic Journal: Applied Economics 2009, 1:4, 200-232.

* Chartrand, Tanya, Joel Huber, Baba Shiv, and Robin Tanner. (2008). “Nonconscious goals and consumer choice.” *Journal of Consumer Research* 35 (August): 189-201.

* Chetty, R., Looney, A., & Kroft, K. (2009). "Salience and Taxation: Theory and Evidence." The American Economic Review, 99(4), 1145-1177.

* Chetty, Raj, and Emmanuel Saez. “Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients.” NBER Working Paper No. 14836. Cambridge, MA: National Bureau of Economic Research, April, 2009.

* Christakis N. A. and Fowler J.H. (2007). “The Spread of Obesity in a Large Social Network over 32 Years.” New England Journal of Medicine. Volume 357 (4) pp. 370-379.

* City of Cape Town. (2011). "SMART LIVING HANDBOOK". Fourth edition, May 2011. Published by the City of Cape Town. ISBN 978-0-9802784-7-7. Available at: http://www.capetown.gov.za/en/EnvironmentalResourceManagement/Documents.

* City of Cape Town. (2006). "Water saving tips for the home". http://www.capetown.gov.za/en/Water/Documents/Publications/home_E.pdf

* Congdon, W., Kling, J. R., & Mullainathan, S. (2009). "Behavioral economics and tax policy" NBER working papers (No. w15328). National Bureau of Economic Research.

Croson R. (2007). “Theories of commitment, altruism and reciprocity: evidence from linear public goods games.” Economic Inquiry*. 2007 (45) no. 2. Pp. 199-216.

* Croson R., Shang J. (2009). “A Field Experiment in charitable contribution: the impact of social information on the voluntary provision of public goods.” *The Economic Journal.* 119 (October) pp.1422-39.

* Croson, R., Handy, F., & Shang, J. (2006). “The relationship between norms, social information and subsequent giving: results from a donor survey and a lab experiment .” (Working paper). University of Texas at Dallas.

* Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. Journal of Abnormal and Social Psychology, 51, 629-636.

* Diehl, Kristin (2005), “When Two Rights Make a Wrong: Searching Too Much in Ordered Environments,” *Journal of Marketing Research*, 42 (3), 313-22.

* Diehl, Kristin, Laura J. Kornish, and John G. Lynch Jr. (2003), “Smart Agents: When Lower Search Costs for Quality Information Increase Price Sensitivity,” *Journal of Consumer Research*, 30 (1), 56-71.

*Durlauf, S.E. (2004), “Neighborhood effects,” In: J.V. Henderson and J-F. Thisse (Eds.), Handbook of Regional and Urban Economics Vol. 4, Amsterdam: Elsevier Science, pp. 2173-2242.

* Ehrenberg A.S. (1972) "Repeat-Buying: Theory and Applications", 2nd ed. London: Charles Griffin; New York: Oxford University Press.

* Ehrenberg A.S., Barwise T.P. & Goodhardt G.J. (1990), "Double Jeopardy revisited". The Journal of Marketing, Vol. 54, No. 3 (Jul., 1990), pp. 82-91.

* Fellner G., Sausgruber R., Traxler C. (2009) “Legal Threat, Moral Appeal and Social Information: Testing Enforcement Strategies in the Field.” Unpublished Paper.

* Ferraro, Paul J., and Michael K. Price. (2011). “Using Non-Pecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment.” NBER Working Paper No. 17189. http://www.nber.org/papers/w17189

* Ferraro, P. J., Miranda, J.J., & Price, M. K. (2011). The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment. *American Economic Review*, *101*(3), 318.

* Finkelstein, A. (2009). "E-ztax: Tax salience and tax rates." The Quarterly Journal of Economics, 124(3), 969-1010.

* Fischbacher, U., Gaechter, S., & Fehr, E. (2001). “Are people conditionally cooperative? Evidence from a public goods experiment.” *Economics Letters*, *71*, 397-404.

* Frey, B., Meier S. (2004a). Pro-social behaviour in a natural setting. *Journal of Economic Behavior and Organization* 54, 65-88.

* Frey, B., & Meier, S. (2004b). "Social comparisons and pro-social behavior: testing 'conditional cooperation' in a field experiment." *American Economic Review*, *94*, 1717-1722.

* Frey B., & Torgler B. (2007). "Tax morale and conditional cooperation." *Journal of Comparative Economics.* Vol. 35 pp. 136-159.

* Greifeneder, Rainer, Benjamin Scheibehenne, and Nina Kleber (2010), "Less May Be More When Choosing Is Difficult: Choice Complexity and Too Much Choice," *Acta Psychologica*, 133, 45-50.

* Haines, M., & Spear, S. (1996). Changing the perception of the norm: A strategy to decrease binge drinking among college students. Journal of American College Health, 45, 134-140.

* Heldt, T. (2005). "Conditional Cooperation in the Field: Cross-Country Skiers' Behavior in Sweden". unpublished manuscript, Uppsala University.

* Huffman, Cynthia and Barbara E. Kahn (1998), "Variety for Sale: Mass Customization or Mass Confusion?" *Journal of Retailing*, 74 (4), 491-513.

* Kahneman, Daniel. (2003). "Maps of Bounded Rationality: Psychology for Behavioural Economics." American Economic Review, 93(5): 1449-75.

* Lalive R. & Catteneo A. (2009). "Social interactions and schooling decisions." *The Review of Economics and Statistics.* 91(3) pp. 457-77.

* Leatherdale, S.T., Papadakis, S. (2011). "A Multi-Level Examination of the Association between Older Social Models in the School Environment and Overweight and Obesity among Younger Students." *Journal of Youth and Adolescence* 40, 361-372.

* Leatherdale, S., Cameron, R., Brown, K., Jolin, M., Kroeker, C. (2006). "The Influence of Friends, Family, and Older Peers on Smoking among Elementary School Students: Low-Risk Students in High-Risk Schools." *Preventive Medicine* 42, 218-222.

*LeBoeuf, Robin, Eldar Shafir, and Julia Belyavsky Bayuk. (2010). "The conflicting choices of alternating selves." *Organizational Behavior and Human Decision Processes* 111: 48-61.

* Lee, Byung-Kwan and Wei-Na Lee (2004), "The Effect of Information Overload on Consumer Choice Quality in an On-line Environment," *Psychology and Marketing*, 21 (3), 159-83.

* Levitt, S. D., & Dubner, S. J. (2009)." Superfreakonomics." Penguin.

* Lurie, Nicholas H. (2004), "Decision Making in Information-Rich Environments: The Role of Information Structure," *Journal of Consumer Research*, 30 (4), 473-86.

* Markowitz, E.M. & Doppelt, R. (2009). "Reducing greenhouse gas emissions through behavioral change: An assessment of past research on energy use, transportation and water consumption." Climate Leadership Initiative. Retrieved the Scholarship bank at the University of Oregon website, 02 November 2012. https://scholarsbank.uoregon.edu/xmlui/bitstream/handle/1794/8373/behavioral_change.pdf?sequence=1

* McPhee, W. N. (1963), Formal Theories of Mass Behavior. New York: The Free Press.

* Milgram S., Bickman L., Berkowitz L. (1969). “Note on the drawing power crowds of different size.” *Journal of Personality and Social Psychology*. 1969, Vol. 13, No. 2, 79-82

* Neighbors, C., Larimer, M., & Lewis, M. (2004). "Targeting misperceptions of descriptive drinking norms: Efficacy of a computer delivered personalized normative feedback intervention." *Journal of Consulting and Clinical Psychology*, 73, 434-447.

* Perkins H. W (2002). " Social Norms and the Prevention of Alcohol Misuse in Collegiate Contexts." JOURNAL OF STUDIES ON ALCOHOL / SUPPLEMENT NO. 14,2002.

* Perkins H. W., Haines P.M. & Rice R. (2005). “Misperceiving the College Drinking Norm and Related Problems: A Nationwide Study of Exposure to Prevention Information, Perceived Norms and Student Alcohol Misuse.” Journal of Studies on Alcohol and Drugs. Vol. 66 (4), July 2005.

* Potters J., Sefton M., Vesterlund L. (2005). “After you—endogenous sequencing in voluntary contribution games.” *Journal of Public Economics.* (89) pp. 1399-1419.

* Prentice, D.A., & Miller, D.T. (1993). "Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm." Journal of Personality and Social Psychology, 64, 243-256.

* Romer, D. (1984). "The Theory of Social Custom: A Modification and Some Extensions. " *The Quarterly Journal of Economics (1984) 99(4): 717-727*

* Roth, Alvin E. 1985. Toward a focal point theory of bargaining. In Game-Theoretic Models of Bargaining, Alvin E. Roth, ed. Cambridge UK and New York: Cambridge University Press.

* Russo, J. Edward (1977), “The Value of Unit Price Information,” *Journal of Marketing Research*, 14 (2), 193-201.

* Sacerdote B. (2001). “Peer Effects with Random Assignment: Results for Dartmouth Roommates.” The Quarterly Journal of Economics, Vol. 116, No. 2 (May, 2001), pp. 681-704

* Saez, Emmanuel. “Do Tax Filers Bunch at Kink Points? Evidence, Elasticity Estimation, and Salience Effects.” University of California at Berkeley. Mimeo, 2009.

* Salganik M. J., Dodds P. J., & Watts D.J. (2006). "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." SCIENCE. Vol(311) 10 February 2006.

* Schultz, P.W., Khazian, A. M., & Zaleski, A. C. (2008). Using normative social influence to promote conservation among hotel guests. Social Influence, 3, 4-23.

* Sela, Aner, and Baba Shiv. (2009). “Unraveling priming: When does the same prime active a goal versus a trait?” *Journal of Consumer Research* 36 (October): 418-433.

* Sexton, S.E. (2011). " Automatic Bill Payment, Price Salience, and Consumption: Evidence from Residential Electricity Consumption." Draft.

* Sherif, M. (1936). "The psychology of social norms." Oxford, England: Harper.

* Slabbert S. (2010a). "Towards standards for municipal invoices in South Africa." Water Research Commission of South Africa. 2010/06/23; Research Report No.TT 458-10.

* Slabbert S. (2010b). "Guidelines on domestic water accounts – towards a consistent approach in South Africa." Water Research Commission of South Africa. 2010/06/14; Research Report No.TT 457/10.

* Smith G. (2013). "Social norms, social approval and social distance: A natural field experiment in electricity consumption in Cape Town." Unpublished masters' thesis. University of Cape Town.

* Spicer, Micheal W., Becker, Lee A., (1980). "Fiscal inequity and tax evasion: An experimental approach." National Tax Journal 33, 171-175.

* Staats H, Harland P & Wilke H. (2004). "Effecting durable change: a team approach to improve environmental behavior in the household." Environment and behavior, Vol. 36 No. 3, May 2004, 341-367.

* Stango, Victor, and Jonathan Zinman. (2011). "Fuzzy math, disclosure regulation, and credit market outcomes: Evidence from Truth-in-Lending reform." *Review of Financial Studies* 24 (2): 506-534.

* Thaler R. and Sunstein C. (2008) "Nudge". Penguin

* Wärneryd, K. 1994. Transaction cost, institutions, and evolution. Journal of Economic Behavior and Organization 25, 219-239.

* Wechsler, H., Nelson, T., Lee, J.E., Seiberg, M., Lewis, C., & Keeling,R. (2003). Perception and reality: A national evaluation of social norms marketing interventions to reduce college students' heavy alcohol use. Quarterly Journal of Studies on Alcohol, 64, 484-494.

* Werch, C.E., Pappas, D.M., Carlson, J.M., DiClemente, C.C., Chally, P.M., & Sinder, J.S. (2000). Results of a social norm intervention to prevent binge drinking among first-year residential college students. Journal of American College Health, 49, 85-92.

* Young, H. P. and Burke M.A. (2001). "Competition and custom in economic contracts: a case study of Illinois agriculture." American Economic Review 91, 559-573.

* Zwane, Alix, Jonathan Zinman, Eric Van Dusen, William Pariente, Clair Null, Edward Miguel, Michael Kremer, et al. 2011. "Being surveyed can change later behavior and related parameter estimates." *Proceedings of the National Academy of Science* 10 (1073): 1-6.

Appendix: Examples of inserts

DHM

Page 3 of 3

Your average daily water consumption for last month.

000014

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD

THIS CITY WORKS FOR YOU

Page 3 of 3

Your average daily water consumption for last month.

compared to

the estimated average household in your neighbourhood.

You consumed more than your average neighbour

000008

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD

THIS CITY WORKS FOR YOU

Your average daily water consumption for last month.

compared to

the estimated average efficient household in your neighbourhood.

You consumed more than your average efficient neighbour

000117

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD

THIS CITY WORKS FOR YOU

How to use less water.

Leaks: Find leaks and fix them. In particular, check your toilet cistern and taps for leaks and drips. Ask a plumber to help you fix any leaks you find if you do not think you are able to fix them yourself.

Washing: Flush your toilet only when necessary. Wash one large load of dishes instead of many small loads. Wash one large load of clothes rather than many small loads of clothes. When you shower: establish how long you shower for and then try take shorter showers. When you bath: establish how full you usually fill your bath and then try to fill it less the next time.

Outdoors: Clean driveways, patios or other hard surfaces with a broom and not with water. Water your garden when the sun is down to reduce water lost to evaporation. Cover your swimming pool to reduce water lost to evaporation. Collect rain water (in bottles or tanks attached to your gutters as well as other appropriate places on your property) to use for watering the garden (do not use this rain-water for drinking, since this may be unsafe and is also illegal).

Buy efficient appliances: Installing an efficient shower head that aerates the water can greatly reduce the amount of water you use when showering. If you need to buy new appliances that use water, be sure to buy water-efficient appliances. For appliances such as washing machines you can check on the packaging how much water they will require every time they are used.

000040

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD

THIS CITY WORKS FOR YOU

Your household's average daily water consumption each month over the past year.

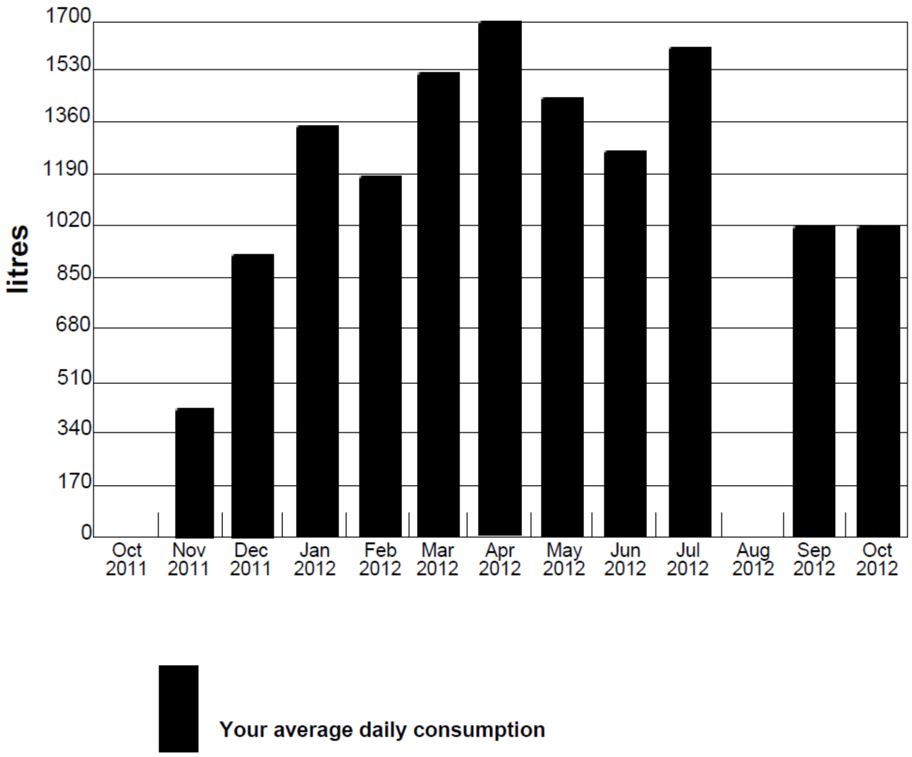

000010

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD
THIS CITY WORKS FOR YOU

Your household's average daily water consumption each month over the past year compared to the average daily water consumption over the last year for the estimated average household in your neighbourhood.

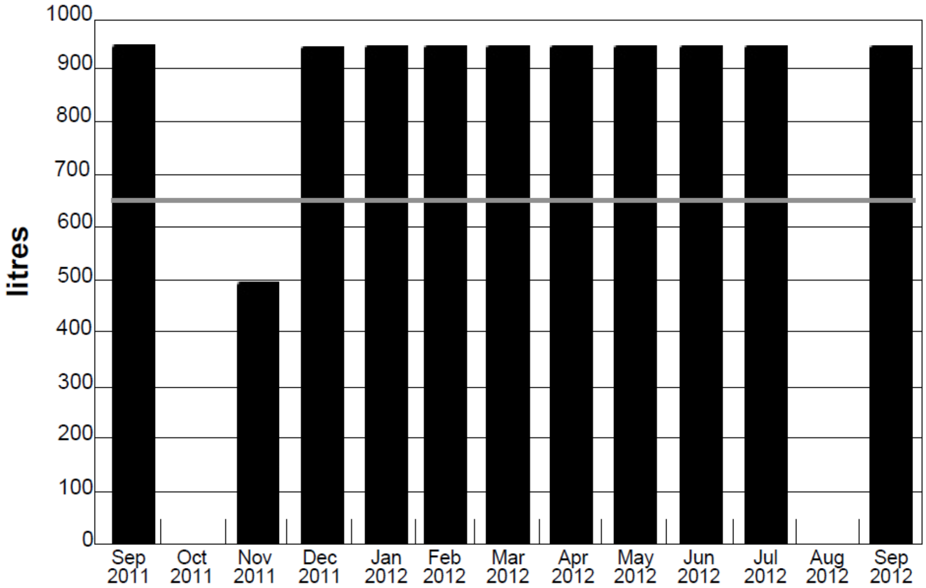

Your average daily consumption

Average neighbour's estimated daily average for last year

000002

CITY OF CAPE TOWN | ISIXEKO SASEKAPA | STAD KAAPSTAD

THIS CITY WORKS FOR YOU